RTView® Oracle® Coherence Monitor User's Guide

Version 6.4



$\mathsf{RTView}^{\mathbb{R}}$

Copyright © 1998-2017. All rights reserved.

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

LIMITATIONS ON USE

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

SL, SL-GMS, GMS, RTView, SL Corporation, and the SL logo are trademarks or registered trademarks of Sherrill-Lubinski Corporation in the United States and other countries. Copyright © 1998-2017 Sherrill-Lubinski Corporation. All Rights Reserved.

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

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

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



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

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

Contents

Contents	iii
Preface	1
About This Guide	1
Document Conventions	1
Additional Resources	1
Release Notes	2
SI Documentation	2
Support Knowledge Base	2
Contacting SL	2
Internet	2
Technical Support	2
Chapter 1 - Introduction to the Monitor	3
Overview	3
Multi-Cluster Configuration	4
Solution Package Version	5
System Requirements	5
Installation	5
Standalone Monitor	5
Setup	5
Verify System	6
Gather Information for Configuration	6
JMX Connection Information	6
Direct Connection Information	6
Create a Project Directory	7
Quick Start	8
Web Browser Deployment	8
Desktop Application Deployment 1	10
Chapter 2 - Configure a Data Connection	15
Overview 1	15
Initialize a Command Prompt or Terminal Window 1	16
Configuring a IMX Connection 1	16
Named JMX Connection	17
Password Encryption 1	18

Multi-Cluster Configuration	18
Configuring a Direct Connection	19
Getting Coherence Parameters	19
Using the OCM Agent	21
Verifying Your Configuration	22
Windows	22
UNIX/Linux	23
Chapter 3 - Deployment	25
Start the Monitor	25
Windows.	25
UNIX	26
Stop the Monitor	27
Windows	27
UNIX	28
Chapter 4 - Using the Monitor	29
Overview	29
Monitor Main Display	30
Heatmaps	32
Tables	33
Multiple Column Sorting	34
Column Visibility	34
Column Filtering	34
Column Locking	36
Column Reordering	36
Saving Settings	37
Row Paging	37
Row Color Code	38
Row Keyboard Selection	38
Trend Graphs	38
litle Bar	40
Context Menu	41
Multiple Windows	41
Cluster Selector	42
Cluster Selector	42
Cluster Views	44
Cluster - Overview	44
Caches / Nodes / Alerts	48
Memory/Network Health	51
Stability Metrics	53
All Services History	55
All Caches History.	58

All Nodes History	. 62
Proxy Services	. 64
Proxy / Extend Overview	. 65
Proxy / Extend Connections	. 66
Proxy / Extend Detail	. 72
Proxy Nodes History	. 76
Extend Connections History	. 78
Cache Services	. 79
Single Service Summary	. 80
Service Metrics Overview	. 84
Service Metric Heatmap	. 85
Single Service History	. 87
Cache Service Detail	. 89
Federated Clusters	. 91
Federated Destination Detail	. 92
Federated Destination Summary	. 95
Federated Origin Detail	. 97
Federated Origin Summary	100
All Caches	102
All Caches Heatmap	103
Storage Nodes Cache Map	104
Current Size Chart	106
Current Activity Chart	107
Single Cache	108
Single Cache Summary	109
Size Trends	112
Activity Trends	115
Cache Detail Tables	117
Storage Manager Detail	119
Node/Group Distribution	121
Front/Back Analysis	123
All Nodes	124
All Nodes by Type/Host/Memory	125
All Nodes CPU	126
All Nodes Grid View	127
Communication Issues	128
All Nodes - Detail	130
Invocation Service Detail	132
Single Node	133
Node Summary	134
Service Trends	137
Node Detail	139
JVM Summary	141
JVM Memory Trends	143
JVM GC Trends	145
System Properties	147
Time Range Analysis	148

	Service Comparison 1	148
	Cache Comparison 1	149
OC A	Administration 1	150
	OC Metrics Administration 1	151
	Cluster MBean Servers 1	153
	Management Settings 1	155
	Node Administration 1	157
	Cache Administration 1	159
	RTView Cache Tables 1	160
	About 1	162
Alert Vie		162
	t Detail Table	162
AICH	Administration	166
Alert	t Administration	166
AICH	Tabular Alert Administration	170
	Setting Override Alerts	171
	Metrics Administration	172
	DTViow Cacho Tablos	174
	About	174
		170
	amont Companyator	
Chapter 5 - Re	eport Generator	. / /
Overviev	w 1	177
Generat	ing Reports from the Monitor 1	177
Generat	ing Reports From a Command Line 1	178
Customi	izina Reports	178
The	Report Generator Configuration File 1	178
Repo	ort Generator Configuration File Parameters 1	179
Sam	ple Report Configuration File 1	181
Chapter 6 - A	dditional Configurations1	183
- Databas	e Configuration	183
	C	18/
	3C	185
Data	abase Canacity and Purging	185
Defi	ning Database Connections	185
Modi	ifving Database Connections	186
mou	Encrypting Passwords	186
Defir	ning Database Tables	187
Configur	ring Servlets	188
\\/inc	dows	188
	1 1	189
Configui	Ting Desktop Clients	190
Stan	ndaione Operation1	190
11010	A Server Lonnection	19()

Data Servlet Connection	191
Configure Alert Notification	192
Using a Batch File or Shell Script	192
Windows Batch File	193
UNIX/Linux Shell Script	193
Using the Java Command Handler	195
Customizing the Java Command Handler	195
Java Command Handler Substitutions	196
Notification Persistence	197
Log File Management on UNIX/Linux Systems	197
Role Based Security	197
Single Sign-On	197
High Availability Configurations	198
Custom Web Application Queries	199
Configure rtvquery Servlet	199
Install rtvquery Servlet	200
Deploy Multiple rtvquery Servlets	200
Using rivquery Serviet	201
SOL Table Query Format	202
Response Formats	204
Servlet Configuration Files	207
JavaScript Library	207
rtvQuerv Class Fields	207
rtvQuery Class Functions	208
Encrypting Passwords for users.xml	212
Invocation-only Configuration	214
Appendix A - Monitor Scripts	217
Scripts	217
rtvservers.dat	223
Single Configuration File	223
Multiple Configuration File	224
Annendix B - Monitor Properties	207
	221
	227
Property Format	228
Property Filters	228
Properties	229
Substitutions	231
Substitution Syntax	232
Database Substitutions	232

Contents

Alert Substitutions 2	35
Filter Substitutions 2	36
Cache Substitutions 2	37
Node Expiration Mode Substitutions 2-	43
Cluster Substitutions 24	44
Appendix C - JMX Connection Options	47
Connection to Cluster Using JMX Remote Port or RMI URL 24	47
Optimizing Data Retrieval Using JMX Tables 24	48
Direct Connection to Cluster as a Coherence Management Node	50
Appendix D - Alert Definitions	53
Alert Thresholds	53
Alert Types 2	53
Monitor Alerts 2	54
Appendix E - Limitations	61
iPad Safari Limitations 20	61
Appendix F - Third Party Notice Requirements	65

Preface

Welcome to the RTView® Oracle® Coherence 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" on page 1
- "Additional Resources" on page 1
- "Contacting SL" on page 2

About This Guide

The *RTView® Oracle® Coherence 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
italics	Within text, new terms and emphasized words appear in italic typeface.
boldface	Within text, directory paths, file names, commands and GUI controls appear in bold typeface.
Courier	Code examples appear in Courier font: amnesiac > enable amnesiac # configure terminal
< >	Values that you specify appear in angle brackets: interface <ipaddress></ipaddress>

Additional Resources

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

- "Release Notes" on page 2
- "SL Documentation" on page 2

Release Notes

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

SL Documentation

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

Support Knowledge Base

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

Contacting SL

This section describes how to contact departments within SL.

Internet

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

Technical Support

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

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

CHAPTER 1 Introduction to the Monitor

This section contains the following:

- "Overview" on page 3
- "System Requirements" on page 5
- "Installation" on page 5
- "Setup" on page 5

Overview

The RTView® Oracle® Coherence Monitor (also referred to as the *Monitor*) provides information about the health and configuration of your Oracle Coherence cluster elements, including caches, nodes, services and clients. The Monitor can be configured for a single Coherence cluster or multiple Coherence clusters (see below).

The Monitor collects metrics from all your Coherence elements simultaneously, and does so at frequent intervals (typically every 10 seconds). At each interval, the Monitor performs analytic calculations on the gathered metrics (on the Data Server rather than a database for optimal performance) in terms of the cluster as a whole. It then presents consistently updated health "snapshots" of your entire cluster, in real time, using a dashboard format and visually rich and legible graphics.

For example, history heatmaps, such as the following cache heatmap, show you utilization trends, over time, for your entire cluster.



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

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

The Monitor is also often used in pre-production environments for conducting load testing and performance tuning.



The following figure illustrates the main Monitor display.

Multi-Cluster Configuration

The Multi-Cluster Configuration is suitable for monitoring many small clusters with a single Monitor instance, or for monitoring a single large cluster. If you have more than one cluster to monitor, consider the multi-cluster configuration benefits:

- Centralizes the monitoring of Coherence clusters
- Metrics for all clusters is accessed from a single URL (rather than a URL for each cluster)
- Easy to configure
- Requires a single historical database instance
- Reduces hardware costs
- Simplifies Monitor configuration

For details, see "Multi-Cluster Configuration" on page 18.

Solution Package Version

The RTView Oracle Coherence Monitor can also be installed as a Solution Package within the RTView Enterprise Monitor® product. RTView Enterprise Monitor is an end-to-end monitoring platform that allows application support teams to understand how infrastructure, middleware and application performance data affect the availability and health of the entire application. When used as a Solution Package, the Coherence metrics and health state are but one source of information which determines the entire health state of the application.

For details about RTView Enterprise Monitor® and the Solution Package for Oracle Coherence, see the *RTView Enterprise Monitor® User's Guide*, available at http://www.sl.com/support/.

System Requirements

Please refer to the **README_sysreq.txt** from your product installation. A copy of this file is also available on the product download page.

Installation

This section describes how to install the RTView Oracle Coherence Monitor.

Standalone Monitor

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

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. Navigate to your installation directory, make the file fixperms.sh executable, then execute it in the current shell (first **chmod 755 fixperms.sh**, and then **./fixperms.sh**).

Proceed to "Setup," next.

Setup

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

- "Verify System" on page 6: This is required.
- "Gather Information for Configuration" on page 6: This is required.

- "Create a Project Directory" on page 7: This is required.
- "Quick Start" on page 8: This is optional.

Verify System

Verify Coherence Settings and JMX Compatibility, and make adjustments as needed.

Coherence Settings

Verify that your Coherence cluster is configured with unique Member Names for each node. For information, see Oracle Coherence documentation:

http://docs.oracle.com/cd/E18686_01/coh.37/e18677/cluster_setup.htm#COHDG5446.

JMX Compatibility

Verify that your system is able to monitor Coherence using JMX. See Using JMX to Manage Coherence at the following link: http://docs.oracle.com/cd/E18686_01/coh.37/e18682/jmx.htm. Specifically consider section **2.2 Accessing Coherence MBeans**: http://docs.oracle.com/cd/E18686_01/coh.37/e18682/jmx.htm#BABDIEJG.

Gather Information for Configuration

After installing the Monitor you configure it to monitor your Coherence Cluster(s). We recommend that you gather the necessary configuration information prior to installing the Monitor. Choose the appropriate connection and monitoring mechanism:

- "JMX Connection Information," next: JMX Connection requires knowledge of how JMX monitoring is enabled and configured on the cluster.
- "Direct Connection Information" on page 6: Direct Connection requires knowledge of how the cluster is configured so the Monitor can join the cluster directly as a node.

JMX Connection Information

Use the table below to gather information you need for configuring a JMX Connection to an existing MBean Server.

Information Needed

JMX Remote Port Connection: • The IP or host name.

• The port number.

JMX RMI Adapter Connection:

• The URL for the connection.

Authenticated JMX Remote Port Connection:

• JMX Authentication and authorization mechanisms (for example, Username and password).

Direct Connection Information

Use the table below to gather information you need for configuring the Monitor application. If the Java properties and class path used by the cluster you want to monitor are not readily available from cluster design documents, cluster launch scripts, or knowledgeable personnel, the following might assist with discovery of the required settings:

- Use the Coherence MBeanConnector to add a management node to your cluster as described in section 2.2.3 Setting Up the Coherence MBean Connector at the following link: http://docs.oracle.com/cd/E18686_01/coh.37/e18682/jmx.htm#CEGBECFH.
- Adjust JVM options and the class path as needed so that the MBean Connector joins the cluster.
- Use jconsole to connect to the MBean Connector node, and verify that all Coherence MBeans appear in the jconsole MBeans tab (Cache,Cluster,Connection,ConnectionManager,Node,Platform,Service,StorageM anager).
- Save the JVM options and class path settings that work with the MBean Connector for configuring the Monitor.

Information Needed

JAR File Information:

- Paths to Coherence JARs and patch JARs.
- Paths to all JAR files that facilitate deserialization of MBeans, including JAR files for custom and third party MBeans.

Paths to database JDBC JAR files.

Direct Connection - Cluster Discovery:

- **Override File** The name of the override file, if one is used, that contains all cluster discovery parameters.
- WKA Connection -
 - The name of the cluster.
 - WKA IP or host, or WKA list.
 - WKA port if not using default.

Multicast Parameters Used By Existing Cluster Nodes

- Cluster name
- Cluster address
- Cluster port
- Local port
- Edition
- Mode

Java Properties:

- All Coherence command line override properties used by existing cluster nodes.
- All Java properties that effect communication with cluster nodes such as network protocol properties.
- JMX authentication properties if applicable.

Create a Project Directory

Create a project directory by copying the default settings files into your own project settings directory. Creating this project directory ensures that your projects are not overwritten when the Monitor software is upgraded. Instructions in this documentation also assume you created this project directory. All examples (of configurations, property settings, command execution and so forth) refer to the project directory.

NOTE: Copy these default settings files only once and do so *before* you begin configuring the Monitor.

Create a project settings directory in the **ocmon/projects** directory and name it **mysample**. For example:

ocmon/projects/mysample

Copy all files from the **ocmon/projects/sample** directory into the **mysample** project directory you just created.

Proceed to "Quick Start," next to install, configure and start the Monitor using default settings.

Quick Start

This section describes how to install, configure and start the Monitor using default settings. These instructions are intended for deployments in which you connect to a single Coherence cluster, use the Direct Connection data connection method (in which the Monitor joins the Coherence cluster as a node) and the default HSQLDB database.

NOTE: The HSQLDB database is sufficient for evaluation and testing, it is not recommended for use in production deployments.

This section includes instructions to install as a:

- "Web Browser Deployment" on page 8
- "Desktop Application Deployment" on page 10

Web Browser Deployment

To deploy the Monitor as a web browser:

- "Windows" on page 8
- "UNIX/Linux" on page 9

Windows

1. Download and unzip the contents to your local server:

Unzip the OCMonitor_<version>.zip file into a directory of your choosing:

unzip -a OCMonitor_<version>.zip

2. Create a project directory as follows:

In the **rtvapm\ocmon\projects** directory, create a directory named mysample. Copy the contents of the **rtvapm\ocmon\sample** directory to your mysample directory.

3. Open the **rtview.properties** file, located in your project directory, and edit as follows: Specify that the Monitor connect as a node:

sl.rtvapm.ocmon.node=true

Define values for the tangosol properties that your cluster nodes use to join the cluster:

tangosol.coherence.cluster=MyClusterName tangosol.coherence.wka= tangosol.coherence.override= tangosol.coherence.cacheconfig=

Set the value of the **sl.rtvapm.ocmon.jmxconn** property to the name of the cluster:

sl.rtvapm.ocmon.jmxconn=MyClusterName

Add the path to the Coherence jar to the java class path: sl.rtview.cp=c:/coherence/lib/coherence.jar

NOTE: You must also add any Coherence patch jars, and all jar files that facilitate deserialization of MBeans, including jar files for custom and third party MBeans.

4. Start the Monitor. "Initialize a Command Prompt or Terminal Window", and perform the following steps:

Navigate to the **rtvapm** directory and type:

rtvapm_init

Navigate to the **rtvapm\ocmon\projects\mysample** directory and type:

start_rtv default

For details about start_rtv, see "Scripts" on page 217.

- 5. Review log files for errors, located in the **rtvapm\ocmon\projects\mysample\logs** directory.
- 6. Prepare the application server by deploying **ocmon.war**, located in the **ocmon\webapps** directory. For example, in the case of Apache Tomcat this could be done by copying **ocmon.war** to the Tomcat **webapps** directory.
- **7.** Open the Monitor: Open a browser and go to the URL of the deployed ocmon servlet. For example, http://localhost:8080/ocmon. If authentication is required:

Login. The default user name and password are:

User Name: demo

Password: demo

The Monitor Cluster Overview display opens.

See "Using the Monitor" on page 29 for details about using the Monitor.

UNIX/Linux

1. Download and unzip the contents to your local server:

Unzip the **OCMonitor_<version>.zip** file into a directory of your choosing:

```
unzip -a OCMonitor_<version>.zip
```

- 2. Navigate to your installation directory, make the file fixperms.sh executable, then execute it in the current shell (first chmod 755 fixperms.sh, and then ./fixperms.sh).
- 3. Create a project directory as follows:

In the **rtvapm/ocmon/projects** directory, create a directory named mysample. Copy the contents of the **rtvapm/ocmon/sample** directory to your **mysample** directory.

4. Open the **rtview.properties** file, located in your project directory, and edit as follows:

Specify that the Monitor connect as a node:

sl.rtvapm.ocmon.node=true

Define values for the tangosol properties that your cluster nodes use to join the cluster:

tangosol.coherence.cluster=MyClusterName tangosol.coherence.wka= tangosol.coherence.override= tangosol.coherence.cacheconfig=

Set the value of the **sl.rtvapm.ocmon.jmxconn** property to the name of the cluster:

sl.rtvapm.ocmon.jmxconn=MyClusterName

Add the path to the Coherence jar to the java class path:

sl.rtview.cp=coherence/lib/coherence.jar

NOTE: You must also add any Coherence patch jars, and all jar files that facilitate deserialization of MBeans, including jar files for custom and third party MBeans.

Save your changes.

5. Start the Monitor. "Initialize a Command Prompt or Terminal Window", and perform the following steps:

Navigate to the **rtvapm** directory and type:

../rtvapm_init.sh

Navigate to the **rtvapm/ocmon/projects/mysample** directory and type:

start_rtv.sh default

For details about start_rtv, see "Scripts" on page 217.

- 6. Review log files for errors, located in the **rtvapm/ocmon/projects/mysample/logs** directory.
- 7. Prepare the application server by deploying **ocmon.war**, located in the **ocmon\webapps** directory. In the case of Apache Tomcat this could be done by copying **ocmon.war** file to the Tomcat webapps directory.
- **8.** Open the Monitor: Open a browser and go to the URL of the deployed ocmon servlet. For example, http://localhost:8080/ocmon. If authentication is required:

Login. The default user name and password are:

User Name: demo

Password: **demo**

The the Monitor **Cluster Overview** display opens.

See "Using the Monitor" on page 29 for details about how to use the Monitor.

Desktop Application Deployment

To deploy the Monitor as a desktop application:

- "Windows" on page 11
- "UNIX" on page 12

Windows

1. Download and unzip the contents to your local server:

Unzip the **OCMonitor_<version>.zip** file into a directory of your choosing: unzip –a **OCMonitor_<version>.zip**

2. Create a project directory as follows:

In the **rtvapm\ocmon\projects** directory, create a directory named mysample. Copy the contents of the r**tvapm\ocmon\sample** directory to your mysample directory.

3. Open the **rtview.properties** file, located in your project directory, and edit as follows: Specify that the Monitor connect as a node:

sl.rtvapm.ocmon.node=true

Define values for the tangosol properties that your cluster nodes use to join the cluster:

tangosol.coherence.cluster=MyClusterName tangosol.coherence.wka= tangosol.coherence.override= tangosol.coherence.cacheconfig=

Set the value of the **sl.rtvapm.ocmon.jmxconn** property to the name of the cluster:

sl.rtvapm.ocmon.jmxconn=MyClusterName

Add the path to Coherence jar to the java class path:

sl.rtview.cp=c:/coherence/lib/coherence.jar

NOTE: You must also add any Coherence patch jars, and all jar files that facilitate deserialization of MBeans, including jar files for custom and third party MBeans.

Save your changes.

4. Start the Monitor. "Initialize a Command Prompt or Terminal Window", and perform the following steps:

Navigate to the **rtvapm** directory and type:

rtvapm_init

Navigate to the **rtvapm\ocmon\projects\mysample** directory and type:

start_rtv default

For details about start_rtv, see "Scripts" on page 217.

- 5. Review log files for errors, located in the **rtvapm\ocmon\projects\mysample\logs** directory.
- 6. Open the Monitor: Start the Viewer by typing: start_rtv viewer_ds

Login. The default user name and password are:

User Name: demo

Password: demo

The Monitor Cluster Overview display opens.

See Using the Monitor for details about using the Monitor.

UNIX

1. Download and unzip the contents to your local server:

Unzip the **OCMonitor_<version>.zip** file into a directory of your choosing: unzip –a **OCMonitor_<version>.zip**

- 2. Navigate to your installation directory, make the file fixperms.sh executable, then execute it in the current shell (first chmod 755 fixperms.sh, and then ./fixperms.sh).
- **3.** Create a project directory as follows:

In the **rtvapm/ocmon/projects** directory, create a directory named mysample. Copy the contents of the **rtvapm/ocmon/sample** directory to your **mysample** directory.

4. Open the **rtview.properties** file, located in your project directory, and edit as follows:

Specify that the Monitor connect as a node:

sl.rtvapm.ocmon.node=true

Define values for the tangosol properties that your cluster nodes use to join the cluster:

tangosol.coherence.cluster=MyClusterName tangosol.coherence.wka= tangosol.coherence.override= tangosol.coherence.cacheconfig=

Set the value of the **sl.rtvapm.ocmon.jmxconn** property to the name of the cluster:

sl.rtvapm.ocmon.jmxconn=MyClusterName

Add the path to Coherence jar to the java class path:

sl.rtview.cp=coherence/lib/coherence.jar

NOTE: You must also add any Coherence patch jars, and all jar files that facilitate deserialization of MBeans, including jar files for custom and third party MBeans.

Save your changes.

5. Start the Monitor. "Initialize a Command Prompt or Terminal Window", and perform the following steps:

Navigate to the rtvapm directory and type:

../rtvapm_init.sh

Navigate to the **rtvapm/ocmon/projects/mysample** directory and type:

start_rtv.sh default

For details about start_rtv, see "Scripts" on page 217.

- 6. Review log files for errors, located in the **rtvapm/ocmon/projects/mysample/logs** directory.
- 7. Open the Monitor: Start the Viewer by typing: start_rtv.sh viewer_ds

Login. The default user name and password are:

User Name: demo

Password: demo

The Monitor Cluster Overview display opens.

See "Using the Monitor" on page 29 for details about using the Monitor.

CHAPTER 2 Configure a Data Connection

This section describes how to configure the RTView Oracle Coherence Monitor data connection. This section includes:

- "Overview," next
- "Configuring a JMX Connection" on page 16
- "Configuring a Direct Connection" on page 19
- "Using the OCM Agent" on page 21
- "Verifying Your Configuration" on page 22

Overview

This section provides step-by-step instructions for configuring a connection to acquire data from the cluster. You configure the data connection using property values in conjunction with property filters. For details about properties and property filters, see "Monitor Properties" on page 227.

The data connection method options are: a named JMX connection, a multi-cluster configuration, a direct connection and an OCM agent. JMX connection methods are generally used for small clusters, and direct or OCM agent connection methods are generally used for large clusters. NOTE: To configure the Monitor for multiple Coherence clusters a JMX connection method is required.

If you have an existing management node in your cluster, choose the JMX connection variant that applies. If you do not have an existing management node in your cluster, see the following requirements.

See "JMX Connection Options" on page 247 for an overview of data connection methods.

The data connection method options are:

- Named JMX Connection (see "Configuring a JMX Connection" on page 16): This method connects to the cluster via a named JMX connection. The JMX connection name is used to identify the cluster in the database tables.
- Multi-Cluster Configuration (see "Configuring a JMX Connection" on page 16): This
 method is suitable for monitoring many small clusters with a single Monitor instance, or
 for monitoring a single large cluster. This method consolidates the monitoring of
 Coherence clusters (rather than having a monitoring system for each cluster), is easy to
 configure (using one of the JMX Connection methods) and requires a single historical
 database instance.
- Direct Connection (see "Configuring a Direct Connection" on page 19): This method is generally used for large clusters.
- OCM Agent (see "Using the OCM Agent" on page 21): This method is generally used for large clusters to minimize the amount of garbage collection associated with collecting and processing JMX monitoring data.

Assumptions

This document assumes that:

- you installed the Monitor per instructions in "Installation" on page 5.
- you have a project directory, as described in "Create a Project Directory" on page 7.
- you use the configuration files provided and retain their file names. If you change a .properties file name, you must specify the name on the command line.

Initialize a Command Prompt or Terminal Window

To start any RTView process (Data Server, Historian, Viewer and so forth), you must first initialize a command line window on the host. Unix scripts are Bourne shell compatible.

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

Windows

Go to your Monitor installation directory and type:

rtvapm_init

UNIX

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

../rtvapm_init.sh

(Note that there are two periods with a space between them.)

Configuring a JMX Connection

This section provides step-by-step instructions for configuring a JMX connection to acquire data from the cluster. NOTE: To configure the Monitor for multiple Coherence clusters a JMX connection method is required.

If you have an existing management node in your cluster, choose the JMX connection variant that applies. If you do not have an existing management node in your cluster, see the following guidelines.

Also see "JMX Connection Options" on page 247 for further details about JMX connection options.

The data connection method options are:

- "Named JMX Connection" on page 17: This method is can be used with both the JMX remote port and the JMX RMI URL connection methods. Use this method when the JMX connection requires a user name and password.
- "Multi-Cluster Configuration" on page 18: This method is for using the Monitor for multiple Coherence clusters. This method consolidates the monitoring of Coherence clusters (rather than having a monitoring system for each cluster), is easy to configure (using one of the JMX Connection methods) and requires a single historical database instance.

Named JMX Connection

This section describes how to create an RTView JMX connection in the **rtview.properties** file using a text editor and the encode_string utility. These instructions assume you have a project directory, as described in Monitor Setup. The **rtview.properties** file is located in your project directory. For details, see "Create a Project Directory" on page 7.

 Open the rtview.properties file in a text editor and add the following line (below the Collector named JMX Connections JMX connections comment) to set the value for the maincollector.sl.rtview.jmx.jmxconn property:

maincollector.sl.rtview.jmx.jmxconn=<conn_name> <host> <port> URL:- - - 'false'
maincollector.sl.rtview.jmx.jmxconn=<conn_name> <host> <port> URL:- - - 'false'
Where:

<conn_name> is the name of the connection

<host> is the hostname of the machine with the Coherence management node

ort> is the port number used by the management node

(specified by **-Dcom.sun.management.jmxremote.port=xxxx**, as described above) For example:

maincollector.sl.rtview.jmx.jmxconn=MyCluster localhost 9971 URL:- - - 'false'

2. Set the **<username>** and **<password>** as desired for the connection.

For example, when no username or password are required:

maincollector.sl.rtview.jmx.jmxconn=MyCluster localhost 9971 URL:- - - 'false'

For example, when a username and password are required:

maincollector.sl.rtview.jmx.jmxconn=MyCluster localhost 9971 URL:myusername mypassword 'false'

3. Specify the name of the JMX connection you just created by setting the **sl.rtvapm.ocmon.jmxconn** property value as follows:

- use a named jmx connection

sl.rtvapm.ocmon.jmxconn=<conn_name>

Where:

<conn_name> is the name of the connection you created

 Specify to not connect as a node (so we can use JMX) by setting the sl.rtvapm.ocmon.node property value to false:

Specify whether the Monitor should connect as a node or not

sl.rtvapm.ocmon.node=false

5. Verify that all other JMX connection properties are comments:

sl.rtvapm.ocmon.jmxhost
sl.rtvapm.ocmon.jmxport
sl.rtvapm.ocmon.jmxurl

6. Save the rtview.properties file and exit the text editor.

See Password Encryption, next.

Or proceed to "Verifying Your Configuration" on page 22.

Password Encryption

If you create a JMX connection by editing the **rtview.properties** file in a text editor, the connection password will be in plain text. To encrypt the password perform the following steps:

1. "Initialize a Command Prompt or Terminal Window", and in the **ocmon/projects/sample** directory, type:

encode_string jmx <password>

where <password> is your password to be encrypted.

For example:

encode_string jmx newpassword

The encrypted value, a series of numbers, is returned.

2. Copy and paste the encrypted value into the **<password>** field of the JMX connection definition in the **rtview.properties** file. For example:

maincollector.sl.rtview.jmx.jmxconn=MyCluster localhost 9971 URL:myusername 01343013550134601331013490134901353013450134801334 'false'

- 3. Edit the rtview.properties file as needed for authentication:
- Add all necessary JMX options to the JVM property sl.rtview.jvm=
- Add all necessary class paths to the classpath property as **sl.rtview.cp=property** values
- Add keystore

Proceed to "Verifying Your Configuration" on page 22.

Multi-Cluster Configuration

This section describes how to configure the Monitor for multiple Coherence clusters. To configure the Monitor for multiple Coherence clusters JMX connections are required (a direct connection is not compatible, since there can only be a direct connection to a single cluster). These instructions assume you have a project directory, as described in "Create a Project Directory" on page 7. The **rtview.properties** file is located in your project directory.

NOTE: Multi-cluster monitoring requires sufficient resources to monitor all the clusters you intend to monitor. Verify that you have sufficient resources for the clusters you intend to monitor.

To configure the Monitor for multiple Coherence clusters:

- 1. Configure named JMX connections as described in "Named JMX Connection" on page 17.
- **2.** Verify that you can connect to each cluster you wish to monitor using an explicit named JMX connection.
- **3.** Ensure that the following property values are set in the **rtview.properties** file that are used for multi-cluster monitoring:

- sl.rtvapm.ocmon.node=false (where **false** specifies not to use a single direct connection)
- sl.rtvapm.ocmon.jmxconn=* (where * specifies to use all named JMX connections)

NOTE: The ocmon command line scripts use **rtview.properties** as the default **.properties** file. Other **<user>.properties** files can be named explicitly as a command line argument to the **ocmon** scripts. For example, you can use a properties file named

MultiCluster.properties for monitoring multiple clusters, and use it to configure the Monitor processes. For example: **start_rtv default all -properties:MultiCluster**.

Also, a multi-cluster **.properties** file can refer to more than one cluster. Therefore the name of the file should describe the group of clusters monitored (for example, DemoClusters, DevClusters). For a sample **rtview.properties** file, see "Monitor Properties" on page 227.

- 4. Open the rtview.properties file and ensure that all clusters you intend to monitor have correct, unique and meaningfully named JMX connection definitions. Edit as needed. The cluster name should be descriptive as it is used in Monitor displays and alert messages. For example, DEV1 and SALES1. For a sample rtview.properties file, see "Monitor Properties" on page 227.
- **5.** If additional clusters need to be monitored that are not specified as a named connection in the **rtview.properties** file, add an entry for the cluster (using unique and meaningfully named JMX connection definitions).
- 6. Save the file.
- 7. Restart the Data Server.

Proceed to "Verifying Your Configuration" on page 22.

Configuring a Direct Connection

This section provides step-by-step instructions for configuring a direct connection to acquire data from the cluster. This method requires the Monitor Data Server to join the cluster as a node. For an overview, see the Direct Connection system diagram.

All Coherence parameters, options, configuration files and JAR files used by existing cluster nodes must be made available for Data Server configuration. This includes JARs that contain custom MBeans, patches, cluster configuration file, cache configuration file, POF configuration file and any configuration files that they reference.

These instructions assume you have a project directory, as described in "Create a Project Directory" on page 7. The **rtview.properties** file is located in your project directory.

Getting Coherence Parameters

If the Java properties and class path used by the cluster you want to monitor are not readily available from cluster design documents, cluster launch scripts, or knowledgeable personnel, the following steps might assist with discovery of the required settings.

 Use the Coherence MBeanConnector to add a management node to your cluster as described in section 2.2.3 Setting Up the Coherence MBean Connector at the following link:

http://docs.oracle.com/cd/E18686_01/coh.37/e18682/jmx.htm#CEGBECFH.

- Adjust JVM options and the class path as needed so that the MBean Connector joins the cluster.
- Use jconsole to connect to the MBean Connector node, and verify that all Coherence MBeans appear in the jconsole MBeans tab. (Cache,Cluster,Connection,ConnectionManager,Node,Platform,Service,StorageManager).
- Save the JVM options, and class path settings that work with the MBean Connector for later use in Monitor configuration.

To configure a direct connection:

- 1. Open the **rtview.properties** file, located in the **ocmon/projects/sample** directory, in a text editor and make the following changes:
- sl.rtvapm.ocmon.node=true
- Set Coherence Properties for Cluster Node configuration. The Monitor property files support the following Coherence command line override properties (and any other property that begins with tangosol):

tangosol.coherence.cluster tangosol.coherence.clusteraddress tangosol.coherence.clusterport tangosol.coherence.edition tangosol.coherence.mode tangosol.coherence.wka tangosol.coherence.wka.port tangosol.coherence.localhost tangosol.coherence.localport tangosol.coherence.override tangosol.coherence.override tangosol.coherence.cacheconfig tangosol.coherence.management.refresh.policy tangosol.coherence.management.refresh.expiry

- Define values for the properties that your cluster nodes use to join the cluster.
- Uncomment, and set the value of the sl.rtvapm.ocmon.jmxconn property to the desired name of the direct connection. Typically, the name of the cluster is used for this value (or, if the cluster does not have a Coherence cluster name, a unique name is used).

Example:

sl.rtvapm.ocmon.jmxconn=DevClusterA

NOTE: The name specified for the **sl.rtvapm.ocmon.jmxconn** property (in Direct Connection mode) is also used to identify the cluster in alerts and data persisted in the database. Because data from multiple databases might subsequently be combined, best practices dictate that a globally unique identifier be specified for cluster names to easily identify data from each cluster.

- 2. Using the Jar Properties information you previously gathered (see "Gather Information for Configuration"):
- Add the path to the Coherence JAR files required by the existing cluster to sl.rtview.cp=. Include any patch JARs used by cluster nodes, as well as the paths to JAR files for POF serialization. Use one entry per sl.rtview.cp property instance as they are combined into a cumulative classpath.

Windows Example:

sl.rtview.cp=c:\coherence352\lib\coh-352-patch-01.jar sl.rtview.cp=c:\coherence352\lib\coherence.jar

- Add the path to the database JDBC driver, and other required JAR files, using additional sl.rtview.cp= property values.
- Add JVM options as sl.rtview.jvm= property values.

Example:

sl.rtview.jvm=-Dtangosol.coherence.mbeans=/sl-custom-mbeans.xml

- **3.** If none of the following are true, skip this Step. If any of the following are true, complete this Step:
- the existing cluster uses Java authentication features.
- additional Tangosol/Coherence properties or options are needed.
- additional JAR files are needed to connect to the cluster.
- additional JAR files are needed to support custom and third party MBeans.
- additional JAR files are needed for database access.
- additional file paths are needed for Coherence configuration files such as POF configuration.

If any of the above are true, edit the **rtview.properties** file as follows:

- Add the paths to JAR files required to join the cluster as sl.rtview.cp=property values.
- Add the paths needed for JAR and Coherence configuration file directories required by the existing cluster, third party applications or database as sl.rtview.cp=property values.
- Add additional Coherence overrides and Java security properties used by the existing cluster as sl.rtview.jvm=property values.

4. Save the file.

Proceed to "Verifying Your Configuration" on page 22.

Using the OCM Agent

This section describes how to use the OCM Agent to gather data from the Coherence cluster. The OCM Agent is a method for acquiring JMX data from a Coherence cluster that reduces the occurrence of packet loss errors sometimes seen in large clusters. This method is especially useful for monitoring very large clusters containing many MBeans.

Normally, and by default, the Data Server acquires data from a Coherence cluster and supplies it to the Monitor. The Data Server also aggregates and processes raw data into caches, and in the process creates large amounts of transient data. This extra overhead can cause long garbage collection pauses on a node in the cluster, and increase both communication delays and other garbage collection activity in the cluster.

The OCM Agent method employs an OCM Agent Data Server that joins the cluster and is dedicated solely to acquiring JMX data and forwarding it to the Data Server. The Data Server does not join the cluster and is dedicated solely to aggregating and processing raw data into caches. This configuration minimizes garbage collection, reducing the overhead for optimal monitoring.

To summarize the OCM Agent method process:

- The Data Server receives a request for a new data sample.
- The Data Server forwards the request to the Monitor Agent Data Server.
- The Monitor Agent Data Server provides the data to the Data Server.
- The Data Server updates its caches with the new data sample.

The OCM Agent method entails an additional Java process (the OCM Agent) and an additional connection (between the Agent and the Data Server).

These instructions assume you have a project directory, as described in "Create a Project Directory" on page 7. The **rtview.properties** file is located in your project directory.

To use the OCM Agent

- **1.** Verify that all Monitor applications and the cluster are stopped. See Stopping the Monitor for details about stopping Monitor applications.
- 2. "Initialize a Command Prompt or Terminal Window", and navigate to the ocmon/ projects/sample directory.
- 3. In the **projects/sample** directory, start the Monitor applications by typing:

start_rtv agentmode all

NOTE: If you do not wish to start all the Monitor applications at once, you must start each application (or "server") individually and in the following order:

start_rtv agentmode database

start_rtv agentmode dataserver

start_rtv agentmode agentsender

start_rtv agentmode historian

start_rtv agentmode displayserver

NOTE: The **start_rtv** script starts processes in a Monitor configuration as specified in the "rtvservers.dat" configuration file. For details about **start_rtv**, see "Scripts" on page 217.

The OCM Agent is now configured to gather data from the Coherence cluster.

Proceed to "Verifying Your Configuration," next.

Verifying Your Configuration

The purpose of this section is to verify that your Monitor configuration (performed in earlier sections of this documentation) is operating properly before performing a full Web deployment.

These instructions assume you have a project directory, as described in "Create a Project Directory" on page 7. The **rtview.properties** file is located in your project directory.

Windows

On Windows, use the standalone Monitor to test data acquisition from the cluster. For details about **start_rtv**, see "Scripts" on page 217.

- 1. "Initialize a Command Prompt or Terminal Window", and navigate to the **projects\mysample** directory of the Monitor installation directory.
- 2. Start the HSQLDB database by typing:

start_rtv default database

A command prompt window appears for the HSQLDB database.

3. Start the Viewer by typing:

start_rtv viewer -console

A command prompt window appears, followed by the Monitor application.

- 4. Inspect the messages that appear in the command prompt window for errors.
- If there are no errors, the Cluster Overview display appears. Your Monitor configuration is operating properly. It takes at least two JMX retrieval cycles to obtain data and fully populate the displays. Proceed to full deployment as described in the section Starting the Monitor.
- If there are errors they appear in the command prompt window, or display fields remain empty. Proceed to the next Step.
- 5. Perform the following steps:
- Note the errors that occurred.
- Terminate the Monitor.
- Modify the configuration files as needed.
- Restart the Monitor.
- 6. Repeat Steps 2 5 as needed to resolve any data acquisition issues.

Proceed to Starting the Monitor.

UNIX/Linux

For details about start_rtv, see "Scripts" on page 217.

1. "Initialize a Command Prompt or Terminal Window", and navigate to the mysample directory.

Example: cd projects/mysample

2. Start HSQLDB.

Example: start_rtv.sh default database

3. Check the HSQLDB log file for errors. Example without errors:

```
cat hsqldb.log
[Thread[main,5,main]]: checkRunning(false) entered
[Thread[main,5,main]]: checkRunning(false) exited
Startup sequence initiated from main() method
Loaded properties from [/home/m/SLTest/rtvoc_55c1/projects/myocm/server.properties]
Initiating startup sequence...
Server socket opened successfully in 8 ms.
Database [index=0, id=0, db=file:DATA/alertdefs, alias=alertdefs] opened successfully in
452 ms.
```

Database [index=1, id=1, db=file:DATA/rtvhistory, alias=rtvhistory] opened sucessfully in 878 ms. Startup sequence completed in 1341 ms. 2009-11-20 11:16:56.800 HSQLDB server 1.8.0 is online To close normally, connect and execute SHUTDOWN SQL From command line, use [Ctrl]+[C] to abort abruptly

4. Start the Data Server.

Example: start_rtv.sh default dataserver

5. Carefully inspect the Data Server log file, **logs/dataserver.log**, for Monitor and Coherence errors.

NOTE: Coherence can produce a large number of long messages making errors difficult to notice. Common errors at this stage are "class not found" errors due to missing JAR files, and incorrect or blocked ports.

- If there are no errors, proceed to Starting the Monitor.
- If there are errors proceed to Step 8.
- 6. Terminate the Data Server and correct any configuration errors found.

Example: stop_rtv.sh default dataserver

- 7. Repeat steps 6 8 as needed until the Data Server output produces no errors.
- 8. Terminate the data server and HSQLDB and use "ps –ef | grep hsqldb" to find the HSQLDB process.

CHAPTER 3 Deployment

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

- "Start the Monitor" on page 25
- "Stop the Monitor" on page 27

Start the Monitor

This section describes how to start Thin Client Browser deployment components, and includes steps for:

- Windows," next
- "UNIX" on page 26

Windows

- 1. Copy the .war files, located in the rtvapm\ocmon\webapps directory, and deploy them to your Application Server.
- 2. Start your Application Server.
- **3.** "Initialize a Command Prompt or Terminal Window" and change directory (cd) to the ocmon\projects\mysample directory.
- 4. In the **projects\mysample** directory, start the Monitor applications by typing: **start_rtv default**.

NOTE: The **start_rtv** default command starts all the Monitor applications at once. To start each Monitor application individually use the following commands:

start_rtv default database Starts the HSQLDB database.

start_rtv default dataserver Starts the Data Server.

start_rtv default historian Starts the Historian.

start_rtv default displayserver Starts the Display Server.

A log file is created for each application and stored in the logs subdirectory.

5. Open a Web browser and go to http://host:port/ocmon.

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

6. Login. The default user name and password are:

User Name: demo

Password: demo

The Cluster Overview display opens.

✔ n/a:DemoCluster	•	Cluster -	Overview	30-Jun-2
Total Nodes: 60	Storage: 24	Client: 36	Caches: 15	Cohere
Senior Node: 20		Cluster Memory	y Usage Totals	Aler
Client Nodes:	Max MB: 10,543	Used MB: 361	Used %: 3.4	
Storage Nodes:	Max MB: 6,061	Used MB: 1,83	38 Used %: 30.3	3 🧣
		Cache S	Services	
Service Name	StatusHA	Total Nodes	Storage Nodes	Caches
DistributedCache	MACHINE-SAFE	3	3 12	4
DistributedCache-A	MACHINE-SAFE	3	30 14	5
MyDistributedCache	MACHINE-SAFE	2	.0 9	4
OnLineStoreCache	MACHINE-SAFE		5 2	2
Cumulative	Most Gets	6,803		Largest Ca
Product			foo	
UserSession 🧾			foo2	
foo 📃			near-foo	
A-Cache01			near-foo2	
Product Product UserSession foo	Most Gets	6,803	foo foo foo2 foo2 near-foo foo2	Largest

UNIX

- 1. Copy the .war files, located in the rtvapm/ocmon/webapps directory, and deploy them to your Application Server.
- 2. Start your Application Server.
- **3.** "Initialize a Command Prompt or Terminal Window" and change directory (cd) to the ocmon/projects/mysample directory.
- 4. In the **projects/mysample** directory, start the Monitor applications by typing: start_rtv.sh default

NOTE: The **start_rtv.sh** default command starts all the Monitor applications at once. To start each Monitor application individually use the following commands:

start_rtv.sh default database Starts the HSQLDB database.

start_rtv.sh default dataserver Starts the Data Server.

start_rtv.sh default historian Starts the Historian.

start_rtv.sh default displayserver Starts the Display Server.

A log file is created for each application and stored in the logs subdirectory.

5. Open a Web browser and go to http://host:port/ocmon. Where host is the IP or host name where your Application Server is running, and port is the port used by your Application Server. The login display opens in the Web browser.

6. Login. The default user name and password are:

User Name: demo

Password: demo

The Cluster Overview display opens.

n/a:DemoCluster		Cluster - O	verview	30-Jun-201
Total Nodes: 60	Storage: 24	Client: 36	Caches: 15	Coherenc
Senior Node: 20		Cluster Memory L	Jsage Totals	Alerts
Client Nodes:	Max MB: 10,543	Used MB: 361	Used %: 3.4	🧠 N
Storage Nodes:	Max MB: 6,061	Used MB: 1,838	Used %: 30.3	3 (N
		Cache Se	rvices	
Service Name	StatusHA	Total Nodes	Storage Nodes	Caches
DistributedCache	MACHINE-SAFE	33	12	4
DistributedCache-A	MACHINE-SAFE	. 30	14	5
MyDistributedCache	MACHINE-SAFE	. 20	9	4
OnLineStoreCache	MACHINE-SAFE	5	2	2
Cumulative	Most Gets	6,803		Largest Cac
Product			foo	
UserSession			foo2	
foo 📃			near-foo	
A-Cache01			near-foo2	

Stop the Monitor

This section describes how to stop Thin Client Browser deployment components, and includes steps for:

- "Windows," next
- "UNIX" on page 28

Windows

 "Initialize a Command Prompt or Terminal Window" and change directory (cd) to the ocmon\projects\mysample directory. In the projects\mysample directory, stop the Monitor applications by typing: stop_rtv default.

NOTE: The **stop_rtv default** command stops all the Monitor applications at once. To stop each Monitor application individually use the following commands:

stop_rtv default dataserver Stops the Data Server.

stop_rtv default historian Stops the Historian.

stop_rtv default displayserver Stops the Display Server.

If the database is running in a command window, close it by clicking in the command window to get focus, pressing **CTRL-C** until the application terminates and closing the command prompt window.

UNIX

- 1. "Initialize a Command Prompt or Terminal Window" and change directory (cd) to the ocmon/projects/mysample directory.
- 2. In the **projects/mysample** directory, stop the Monitor applications by typing: **stop_rtv.sh default**.

NOTE: The **stop_rtv default** command stops all the Monitor applications at once. To stop each Monitor application individually use the following commands:

stop_rtv.sh default dataserver Starts the Data Server.

stop_rtv.sh default historian Starts the Historian.

stop_rtv.sh default displayserver Starts the Display Server.
CHAPTER 4 Using the Monitor

This section describes RTView Oracle Coherence Monitor features, graphs and functionality as well as Monitor displays. The Coherence displays are located in the **COMPONENTS** tab under **Middleware** in the navigation tree. This section includes:

- "Overview" on page 29: Describes Monitor features and functionality.
- "Cluster Selector" on page 42: See all your Coherence clusters and Data Servers and choose which cluster to display data for.
- "Cluster Views" on page 44: Use these displays to assess Coherence cluster-level performance and utilization.
- "Proxy Services" on page 64: Use these displays to assess proxy service performance metrics.
- "Cache Services" on page 79: Use these displays to assess performance and utilization of all caches in the cluster.
- "Federated Clusters" on page 91: Use these displays to assess performance and utilization of Federated clusters.
- "All Caches" on page 102: Use these displays to investigate performance, utilization and activity metrics of a single cache.
- "Single Cache" on page 108: Use these displays to assess node-level performance and utilization in the cluster.
- "All Nodes" on page 124: Use these displays to investigate performance and utilization metrics of a single node.
- "Single Node" on page 133: Use these displays to investigate performance and utilization metrics of a single node.
- "Time Range Analysis" on page 148: Use these displays to manage your Oracle Coherence metrics, nodes and caches.
- "OC Administration": Use these displays to manage your Oracle Coherence metrics, nodes and caches.
- "Alert Views" on page 162: The display in this View presents the status of all alerts across all BW Servers, and allows you to track, manage and assign alerts.
- "Administration" on page 166: The displays in this View enable you to set global alerts and override alerts. You can also view internal data gathered and stored by RTView (used for troubleshooting with SL Technical Support).

Overview

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

 "Monitor Main Display" on page 30: Describes the Monitor display that opens by default as well as the navigation tree.

- "Heatmaps" on page 32: Describes how to read heatmaps and heatmap functionality.
- "Tables" on page 33: Describes how to read tables and table functionality.
- "Trend Graphs" on page 38: Describes how to read trend graphs and trend graph functionality.
- "Title Bar" on page 40: Describes the top layer of the title bar shared by Monitor displays.
- "Context Menu" on page 41: Describes right-click popup menu in the Monitor.
- "Multiple Windows" on page 42: Describes opening multiple windows in the Monitor.

Monitor Main Display

The **Cluster Overview** is the main display of the Monitor. This display enables you to quickly assess the configuration, activity and health of all of your Coherence clusters.

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

The main is organized by the following areas:



- Coherence Cluster Configuration: Get the cluster name and total counts of members (JVMs) in each cluster. Counts include storage nodes, client nodes (non-storage nodes), the total number of caches and the version of Coherence used in the cluster.
- Memory: Get memory information, including heap size and used memory totals for all storage and client (non-storage) nodes. Also see the total percent memory usage for storage and client nodes. A blue-colored recent memory usage trend chart is displayed for storage nodes and another for client nodes.

Service Configuration & HA Status: Check the high-availability (HA) status for all Coherence protocol-related cache services used by applications in the cluster. The StatusHA column indicates whether primary and backup objects are distributed for surviving machine failure or storage node failure. The most secure status is MACHINE-SAFE which indicates that an entire host could fail and all data could be recovered. NODE-SAFE indicates that a storage node could fail and data could be recovered, but data could be lost with a host failure. ENDANGERED indicates that the loss of a single storage node could result in data loss in the cluster. Note that Coherence does not track whether enough free memory is available for surviving machine or storage node failure without data loss.

Also, see the number of caches in each cache service, the number of storage nodes participating in each cache service, the number of objects in each cache service, the senior member for the cluster and the senior members for each service. Click a cache service to view details in the "Single Service Summary" display.

- Most Gets: Check on your four busiest caches (in terms of the number of gets reported in the last measurement). Mouse-over the bar charts to see the most recent metric for each cache. Select the Cumulative box to display the total number of gets since the cluster started (or since statistics were last reset). Statistics from the most active cache are displayed in the upper right field.
- Largest Cache: Check on your four largest caches (in terms of memory usage). The largest cache's current size, in units, is displayed in the upper right field. Mouse-over the bar charts to see the most recent metric for each cache. Units are user defined in your cache configuration file. Typically units are set to either the number of objects or the number of bytes consumed by objects.
- Node Uptimes: Assess cluster stability. View how long nodes in the cluster have been members of the cluster. The Monitor looks at the start time of every node in the cluster, determines how much time the node has been part of the cluster and categorizes the nodes into seconds, hours, days, weeks or months. Typically, if a node leaves the cluster as the result of a fault condition, the node subsequently rejoins, appears as a "younger" node, and is placed in the seconds, minutes or hours category. This metric can be an important indicator of cluster instability. If a node leaves the cluster and fails to rejoin, the node is subsequently shown in the Departed Nodes box.
- Memory Utilization: Determine whether cluster memory usage has been increasing over the last hour and by how much. These bar charts are capacity indicators as well as garbage collection indicators. The memory usage is shown for the most recent data collection, and as an average for the last hour. The red bar chart is for the node in the cluster with the highest memory use, and the yellow bar shows memory use averaged across the whole cluster.
- Communication Success Rate: Determine whether packet loss is occurring. The bar graph uses pairs of bars in which one bar represents the publisher success rate and the other represents the receiver success rate. The bar charts are the most important indicator for any issue affecting cluster health or performance. The chart shows the TCMP publisher and receiver success rates for the last 20 minutes. The success rate is typically 99% or greater in healthy clusters. Publisher/Receiver failures (the inverse of success) indicate that packets sent between nodes are not being acknowledged within the timeout period (which is typically 250 ms). This is typically the result of a node being unavailable due to garbage collection. However, there are many other possible causes as well (such as a network issue, a defective NIC card, a garbage collection issue, disk swapping, or a shortage of CPU on a single machine). Investigate further by clicking the bar chart to view details in the Cluster "Memory/Network Health" display.

Heatmaps

Heatmaps organize your resources into rectangles and use color to highlight the most critical values in each. Heatmaps enable you to view various metrics in the same heatmap using dropdown 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 to filter data by. The filtering options vary among heatmaps.

For example, each rectangle in the **Proxy / Extend Overview** heatmap represents a node, where color is representative of the selected **Metric**.



The **Metric** drop-down menu in this heatmap contains options to show **Alert Severity**, **Alert Count**, as well as other metrics. Menu options vary according to the data populating the heatmap. **Alert Severity** is selected and its corresponding color gradient **example** bar is shown. **Alert Severity** is the maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient **example** bar, where **2** is the highest **Alert Severity**:

Red indicates that one or more services associated with that node currently has an alert in an alarm state.

• Yellow indicates that one or more services associated with that node currently have an alert in a warning state.

• Green indicates that no services associated with that node have alerts in a warning or alarm state.

In most heatmaps, you can also drill-down to a *Summary* display containing detailed data for the resource. You can also open a new window **I** and then drill-down. The drill-down opens a display that contains relevant and more detailed data.

Mouse-over

The mouse-over functionality provides additional detailed data in an over imposed pop-up window when you mouse-over a heatmap. The following figure illustrates mouse-over functionality in a heatmap object.



Log Scale

Typically, heat maps provide the Log Scale option, which 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.

Tables

Monitor tables contain the same data that is shown in the heatmap in the same View. Tables provide you a text and numeric view of the data shown in that heatmap, and additional data not included the heatmap.

Domain	SLBW6	AppSpace: SimSpace		Count	: 2	Running:	Running Only
Application	Name Filter:			Clear	RegEx		
Domain	AppSpace	Name	Alert Level	Alert Count	State	AppNodes	Version
SLBW6	SimSpace	com.SL.BWApp.Two	(*)	0	Running	6	1.0 E
SLBW6	SimSpace	com.SL.BWApp.One	6	3	Running	6	1.0 E

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

Additional features are:

- "Multiple Column Sorting," next
- "Column Visibility" on page 34
- "Column Filtering" on page 34
- "Column Locking" on page 36
- "Column Reordering" on page 36
- "Saving Settings" on page 37
- "Row Paging" on page 37
- "Row Color Code" on page 38
- "Row Keyboard Selection" on page 38

Multiple Column Sorting

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

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

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

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

Column Visibility

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

E Sort Ascending		
F Sort Descending		
III Columns	+	🗹 Timestamp
▼ Filter	۲	🗹 Name
Settings	۲	Region
		Call Rate
		Active Calls

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

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

Column Filtering

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

The background of a column's menu icon changes to white to indicate that a filter is defined on that column. This is intended to remind you which columns are filtered. You can configure a filter on any column by clicking on the column's menu icon and choosing **Filter** from the menu. This opens the **Column Filter** dialog:

1			
1	Sort Ascending		
₹	Sort Descending		
ш	Columns	۲	
т	Filter	Þ	Show items with value that:
Set	tings	►	Contains 🔻
			abo
			And v
			Does not contain 🛛 🔻
			хуz
			Filter Clear

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

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

The numeric and boolean filter dialogs are shown below.



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



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

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

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

Column Locking

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

	Columns	•
T	Filter	•
	Lock	
<u>_</u>	Unlock	
90	Hinac	

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

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

Column Reordering

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

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

Saving Settings

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

▼ Filter	*	
Lock		
🚡 Unlock		
Settings	+	Save All
		Clear All

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

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

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

Revert Table Settings

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

Row Paging

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

211	TUNOW	alitivicw.aqi.aqiub	Intervision of the root my-secret-pw jube.mysql.misz	
217	emreference	sl.rtview.sub	<pre>\$rtvConfigDataServer:CONFIG_SERVER</pre>	
229	emreference	sl.rtview.properties.queryTimeOut	10	
216	emreference	sl.rtview.sql.sqldb	ALERTDEFSnone	Ŧ
•		II III	4	
	Page 1 of 2	H	1 - 200 of 235 items	

Row Color Code

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

JVM	localhostGLASSFISH_SERVER_8	(1	10	
JV14	IocalhostMYDEMO_DATASERVER		1	8	
JVM	IocalhostMYDEMO_DISPLAYSERVER		1		
JVM	sidemos.com/213415_RTVDB		1	10	
JVM	localhost8WM-DB-1		1	5	
WAS	SLHOST12Node01Ce8,SLHOST12Node01;server1		1	5	
2VM	localhostRTVMGR_DATABASE		 1	6	
JVM	localhostRTVMGR_DATASERVER		0	0	
JVM	IocalhostWLM_DATABASE		0	0	
EMS	tc#//SLHOST10.7021	<	0	0	
EMS	tcp://3LH06T10.7020		0	0	
WLS .	TestDomain:ManagedServer2		 0	0	

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

Row Color Code:

Tables with colored rows indicate the following:

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

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

Row Keyboard Selection

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

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

Trend Graphs

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

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



Time Range

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



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

Mouse-over

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



Log Scale

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

Title Bar

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

	e same top laye		
n/a:DemoCluste	er 🔻	Cluster - Overview	02-Jun-2016 11:38 💠 Data OK + 🖓
The following table	e describes the	functionality in the displa	ay title bar.
<	Opens the previ	ously open display.	
•	Opens the displa	ay that is up one level.	
Table	Navigates to a d The target displa	display that is most common ay differs among displays.	ly accessed from the current display
Menu 💌	Navigates to dis The drop-down	plays that are most common menu options differ among o	nly accessed from the current display displays.
	Opens the Alerts	s Table display in a new wind	dow.
23-Mar-2017 12:04	The current date RTView stopped indicator is gree and valid data.	e and time. If the time is inc running. When the date and n, this is a strong indication	orrect, this might indicate that d time is correct and the Data OK that the platform is receiving curren
财 Data OK	The data connect example, if the land not receive data source is connect indicator is gree and valid data.	ction state. Red indicates the Data Server is not receiving from the Data Server, this v cted. When the date and tim n, this is a strong indication	e data source is disconnected (for data, or if the Display Server does will be red). Green indicates the data is correct and the Data OK that the platform is receiving curren
+	Opens an instan independently, a EM, and compar Windows.	nce of the same display in a n allowing you to switch views, re server performance data.	new window. Each window operates navigate to other displays in RTViev For illustration, see Multiple
0	Opens the online	e help page for the current o	display.
6,047	The number of i	tems (for example, CIs or A	reas) in the display.
Area Count: 9			

Context Menu

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

	.	
Refresh		
Back		
Next		
Execute Command		
Drill Down		
Export Table to Excel		
Export Table to HTML		
Export PDF		
Status		
Log Off		

Export Report

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

To generate a report for a display:

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

Expo Report	ort Type O ^{pisplay}
Orie	ntation
Portrait	t 🕒 andscape
Ma	rgins
Left 1.0	Right 1.0
Тор 1.0	Bottom 1.0
ОК	Cancel

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

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

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

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



Multiple Windows

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



Cluster Selector

Cluster Selector

This display shows details about your Coherence clusters and Monitor Data Servers.

Use this display to see all the Coherence clusters you can monitor, as well as their status. Choose a cluster to view performance details for the cluster in the "Cluster - Overview" display. Each row in the table is a different Coherence cluster. The columns contain information pertaining to each cluster. When you select a cluster you are also selecting the Data Server corresponding with that cluster. After you make your selection, all displays subsequently show data for that cluster/Data Server (except for alert displays which consolidate alerts from all Data Servers). For example, the "Node Summary" display will then show data for the selected cluster/Data Server.

For details about Oracle Coherence data, refer to vendor documentation at www.oracle.com.

Connection	Alert	Alert	ClusterSize	Caches	Objects	Data Server
DemoCluster		0	60	15	1,856,470	defau

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

Connection	The name of the user defined connection that is used to connect to the monitored Coherence cluster.
Alert Severity	 The maximum level of alerts on the cluster. Red indicates that one or more exceeded their ALARM LEVEL threshold. Yellow indicates that one or more exceeded their WARNING LEVEL threshold. Green indicates that none have exceeded their alert thresholds.
Alert Count	The number of current alerts for the cluster.
Cluster Size	The total number of nodes for the cluster.
Caches	The total number of caches for the cluster.

Objects The total number of objects stored in the cluster.

Data Sever The name of the Data Server (connection) that is used to monitor the cluster.

Cluster Views

Cluster Views displays present high-level performance metrics for the cluster. Use the Cluster Views displays to quickly assess Coherence cluster-level performance metrics.

- "Cluster Overview" on page 44: Quickly assess general cluster stability, cluster size (number of nodes, clients and caches), service and cache capacity utilization/distribution and HA status.
- "Caches / Nodes / Alerts" on page 48: View cache and node utilization hot spots and currently active alerts.
- "Memory/Network Health" on page 51: Assess cluster memory utilization and packet transmission success/failure trends, and see weakest nodes.
- "Stability Metrics" on page 53: Troubleshoot nodes joining and leaving the cluster, view HA status for cache services.
- "All Services History" on page 55: Assess capacity utilization, over time, by all services in a cluster.
- "All Caches History" on page 58: Assess capacity utilization and distribution for all caches in a cluster, and quickly identify potential bottlenecks.
- "All Nodes History" on page 62: Assess capacity utilization, over time, for all nodes in a cluster.

Cluster - Overview

Use this display to quickly assess the cluster size (number of nodes, clients and caches) and stability, service and cache capacity utilization and HA status. This display is the initial view in the Monitor.

Choose a cluster from the drop down menu. Check the Communication Success% bar charts for cluster packet loss. If the pairs of bar graphs are uneven, this indicates that packet loss is occurring. The cause for the packet loss could be a network issue, a single defective NIC card, a garbage collection issue, disk swapping or a shortage of CPU on a single machine. Investigate further by clicking the bar chart to view details in the Cluster - "Memory/Network Health" display.



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

Fields and Data:

Coherence Cluster Configuration

Total
NodesTotal number of nodes being monitored, including storage enabled nodes, client
nodes, and management (JMX) nodes.StorageTotal number of nodes in the cluster which have storage enabled for any cache.
This value is equal to the total nodes when replicated caches are being used. The
number is less when only distributed cache types are utilized.ClientsTotal number of nodes in the cluster which do not have storage enabled for any
cache. These are usually process nodes, proxy nodes, extend nodes, or MBean
server nodes.CachesTotal number of caches in the cluster.

Version	Version of Oracle Coherence running.			
Cluster Mem	lemory Usage Totals			
Senior Node	Node ID of th	ne senior node of the cluster.		
Client Nodes	Monitor clien	t node memory utilization for the cluster.		
	Max MB	Total memory allocated.		
	Used MB	Total memory used.		
	%	Percent of allocated memory being used.		
Storage Nodes	Monitor storage node memory utilization for the cluster.			
	Max MB	Total memory allocated.		
	Used MB	Total memory used.		
	%	Percent of allocated memory being used		
Alert Severity	 The maximum level of alerts for all nodes in the cluster. Click to drill down to the Alert Detail Table. Red indicates that one or more exceeded their ALARM LEVEL threshold. Yellow indicates that one or more exceeded their WARNING LEVEL threshold. Green indicates that none have exceeded their alert thresholds. 			
	Memory	Represents the current most critical state of alerts related to heap and memory alerts for all nodes in the cluster. For example, the AvailableMemoryLowNode alert.		
	Network	Represents the current most critical state of alerts related to network and communication protocols for all nodes in the cluster. For example, the BadCommunicationCluster alert.		
	Stability	Represents the current most critical state of alerts related to cluster stability for all nodes in the cluster. For example, the DepartedNodePercentage alert.		
	Tasks	Represents the current most critical state of alerts related to queries, entry processors and invocations for all nodes in the cluster. For example, the HighTaskBacklogNode alert.		
	Data Quality	Represents the current most critical state of alerts related to the quality of data in the Data Server for all nodes in the cluster. For example, the JmxProcessingTime alert.		
	Other	Represents the current most critical state of alerts related to all alerts not represented in the other five status indicators for all nodes in the cluster. For example, the CapacityLimiitAllCaches alert.		
	Memory	Represents the current most critical state of alerts related to heap and memory alerts for all nodes in the cluster. For example, the AvailableMemoryLowNode alert.		

Service Configuration & HA Status

Cache Services Assess size, distribution and status of Coherence protocol-related cache services used by applications in the cluster. Determine whether cache services are distributed properly across the cluster. The list includes distributed, replicated and mirrored caches. Note that Management and Invocation services are intentionally not listed.

Service The name of the service in the cluster. These are defined in each server cache configuration XML file.

StatusHA	The high availability status for each of the services.MACHINE- SAFEIf a machine for the service goes offline the data stored on the machine remains available in the cluster (no data loss).		
	NODE- SAFE	If a node for the service goes offline (or is taken offline using kill-9) data stored on the node remains available in the cluster (no data loss).	
	ENDANGE RED	If a node for the service goes offline the data stored on the node is potentially unavailable in the cluster (potential data loss).	
Total Nodes	The number of nodes in the cluster that are running a thread for the service.		
Storage Nodes	The number of nodes for the service where storage is enabled.		
Caches	The number of caches for the service.		
Objects	The number of objects in all caches for the service.		
Senior	The node ID of the most senior node in the cluster for the service.		

Caches - Busiest & Largest

- Track services performing the greatest number of gets in the cluster. The total is Most Gets the number of gets by nodes in the cluster since the last sample was retrieved. Click to drill-down to the All Caches - "Current Activity Chart" display.
- Cumulative Select the checkbox to show only the cumulative total for all nodes for the service since they started in the Most Gets bar chart.
- Track caches that consume the greatest amount of capacity. Click to drill-down to the All Caches "Current Size Chart" display. Largest Cache

Cluster Stability

Monitor cluster stability and how often nodes are restarted (for example, every Node month, every day, every hour, and so forth). If the number of nodes running for Uptimes seconds of time increases (and your nodes are restarted weekly), consider investigating. Click in the Node Uptimes region to view details on the "Stability Metrics" display.

> Solid colors in the graph indicate the amount of time since the nodes were started. Longer uptimes generally represent a more stable cluster. Departed Nodes specifies the number of nodes that have departed and not returned since monitoring of the cluster was started. If a node departs and returns with the same name, the count is decremented.

Memory Monitor memory utilization for all nodes in the cluster. Utilization %

> The average memory utilization for all nodes in the cluster. Average

The most amount of memory consumed by a single node in the Worst cluster. A slow node that provides data to other nodes can cause Node latency issues for the entire cluster. If a node is consuming too much memory, investigate by clicking the bar chart to view details in the Cluster - "Memory/Network Health" display.

Communica tion Success%	Monitor cluster packet lossan excellent indicator of systemic issues in the cluster. If the pairs of bar graphs are uneven, this indicates that packet loss is occurring and analysis is needed. Investigate further by clicking the bar chart to view details in the Cluster - "Memory/Network Health" display.		
	Ine bar charts show the percent (%) successful UDP packet transfers in the cluster for the last twenty minutes. Each pair of bars show the Publish and Receive success rates for all nodes in the cluster. Compare each pair of Publish and Receive bars. The bars should have similar rates. If they do not have similar rates this indicates packet loss in the cluster. For example, if the Publish success rate is much lower than the Receive success rate, packets are being resent and the receiver is not getting them.		
	Compare and track the pairs of bars across twenty minutes. The bars should track evenly. If the bars do not track evenly this also is a sign of packet loss in the cluster.		
	The cause for card, a garb machine.	The cause for the packet loss could be a network issue, a single defective NIC card, a garbage collection issue, disk swapping or a shortage of CPU on a single machine.	
	Publish	The Publish success rate is the percent (%) of packets in the cluster successfully sent by nodes, without having to be resent. A 100% success rate occurs when a packet is sent and does not have to be re-sent. When a packet must be resent the success rate is reduced.	
	Receive	The Receive success rate is the percent (%) of packets in the cluster successfully received by nodes, without being received twice. A 100% success rate occurs when a packet is received once. When a packet is received twice the success rate is reduced.	

Caches / Nodes / Alerts

Use this display to view cache and node utilization hot spots and currently active alerts. Observe how much capacity is taken from memory and how much is taken from consumption. Identify caches and nodes that are slow due to a shortage of capacity or memory. Verify nodes are configured properly (using the mouseover tool-tip). View time-ordered list of current alerts in the cluster.

← n/a:DemoCluster		Cache	es / Nodes / Alerts	3	0-Jun-2016 16:27 🧉	🔊 Data OK 🛛 💠 🕜
Total Nodes: 60	Storage: 24	Client: 36	Caches: 15		Versi	on: 3.5.3/465
All Ca	ches - Size and Activity			All Nodes - Memo	ory Usage	
			Process N	Nodes	Sto	orage Nodes
					HostName:	SLHOST4
					Location: Sti Id: 13 Max Mem: 2 Used Mem:	oreNode05n.SLHOST 254 Mb 77.6 %
						£
	Cache Size Info			Node Memory / C	Comm Info	
Cache	Objects Units	Expired •	Location		Mem %	Comm %
foo	527,845 143,573	840 🔲 🔤 🛛	StoreNode08.SLHOST4		38.6	.0
foo2	440,012 126,723,	456 🔲 🗍	StoreNode08.SLHOST2		84.3	.0
near-foo	416,937 120,077,	856 🔲 🕴	StoreNode07.SLHOST4		2.4	.0
near-foo2	430,277 117,035,	344	StoreNode07.SLHOST2		1.6	.0
Product	34,827 15,429,9	536	StoreNode06.SLHOST4		2.4	.0
A-Cache04	2,608 751,1	104	StoreNode06.SLHOS12		3.1	
A-Cache03	2,751 748,7	272	StoreNode05n.SLHOS14		//.6	.0
			Current Alerts			
Time	Alert Name	Alert In	idex		Alert Text	
•	III					4

Title Bar (possible features are):	on Data OK Data connection state. Red indicates the Data
🗲 个 Open the previous and upper display.	Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the
Open an instance of this display in a new window.	data source is connected.
Open the online help page for this display. Menu Table open commonly accessed displays.	might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid
6,047 The number of items currently in the display.	Open the Alert Views - RTView Alerts Table display.

Fields and Data:

Total Nodes	Total number of nodes being monitored, including storage enabled nodes, client nodes, and management (JMX) nodes.
Storage	Total number of nodes in the cluster which have storage enabled for any cache. This value is equal to the total nodes when replicated caches are being used. The number is less when only distributed cache types are utilized.
Clients	Total number of nodes in the cluster which do not have storage enabled for any cache. These are usually process nodes, proxy nodes, extend nodes, or MBean server nodes.
Caches	Total number of caches in the cluster.
Version	Version of Oracle Coherence running.
Capacity & Mer	nory Usage
All Caches - Size and Activity	Use the heatmap to identify a cache with high capacity or memory usage, indicated by a dark rectangle. Observe how much capacity is taken from memory and how much is taken from consumption. View cache metrics using the mouseover tool-tip. Investigate cache utilization trends over time in the

"All Caches History" display. Člick on a rectangle to drill-down to the All Caches - "All Caches Heatmap". The heatmap is grouped by service. Each rectangle represents a cache within the service. The size of each rectangle represents the size of a cache in units. The color of each rectangle represents the number of gets on the cache. The color is linearly scaled, where white is the minimum gets seen and dark green

- **Cache Size** Info The table lists each cache in the cluster and enables you to sort the by most/least amount of objects or units. Click a row to view details in the "Single Cache Summary" display.
- **Cache** The name of the cache.

is the maximum gets seen.

- **Objects** The number of objects currently in the cache.
- **Units** The number of units currently used by the cache.

All Nodes-Memory Usage Use the heatmap to identify a node with high memory usage, indicated by a dark rectangle. Verify nodes are configured properly using the mouseover tooltip. Click on a rectangle to drill-down to the "All Nodes by Type/Host/ Memory".

> The heatmap is divided into two sections: Process Nodes and Storage Nodes. Each rectangle represents a node in the cluster. The size of the rectangle represents the value of the maximum node memory. The color of the rectangle represents the value of the memory used. The color is linearly scaled, where white is 0% memory used and dark green is 80% memory used.

> Node The table lists each node in the cluster and enables you to sort the by most/least amount of objects or units. Click a row to view details in the "Node Summary" display.

Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site .
Mem%	The percent memory utilization for the node.
Comm%	The percent memory utilization used for packet transfer by the node.

All Active Alerts (in selected cluster)

Current Alerts The table lists all alerts for all sources (nodes and caches) in the selected cluster that have exceeded an alert threshold. Sort the data by column using the button. By default, critical and warning alerts are shown. Select an alert in the list to open the **Alert Detail Table** dialog and acknowledge an alert or add comments. Where:

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

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

Green indicates that no resources have exceeded their alert thresholds.

For details about alerts, see Appendix, Alert Definitions.

Alert Name The alert type. Alert Types contain alert threshold definitions. A single alert type applies to all nodes or caches in the cluster. For example, the OcAvailableMemoryLowNodeSpike alert type applies to multiple nodes, and the OcCapacityLimitCache alert type applies to multiple caches. (The Alert Index identifies the source node for the alert.)

For details about alerts, see Appendix, Alert Definitions.

- Alert Index The Oracle Coherence source (node or cache) from which the alert originated. As with nodes, a cluster can have multiple caches. A single alert type, such as OcCapacityLimitCache, applies to all caches in the cluster. The Alert Index identifies the cache from which the alert originated.
- **Alert Text** Descriptive information about the alert.
- **Cleared** The checkbox is selected if this alert has cleared. An alert is considered cleared when the source for the alert (node or cache) returns to below the alert threshold. To include acknowledged alerts in the table, select Show Cleared.
- Acknowled ged The checkbox is selected if this alert has been acknowledged. Acknowledged alerts have been manually acknowledged by an administrator. Acknowledged alerts are automatically removed from the Current Alerts table. To include acknowledged alerts in the table, select Show Acknowledged.
- **ID** Unique ID for the alert.
- **Comments** Comments about the alert previously entered by an administrator.
- Cleared An alert is in a cleared state when the source for the alert (node or cache) returns to below the alert threshold. Or, with the OcDepartedNode alert type, when the node rejoins the cluster the alert is cleared.
- Cleared The time the alert was cleared. Time

Alert Index The Oracle Coherence source (node or cache) from which the alert originated.

Cluster The name of the cluster in which the alert source (node or cache) is a member.

Memory/Network Health

Use this display to assess cluster memory utilization and packet transmission success/failure trends, and to see the weakest nodes.

fn/a:DemoCluster	Cluster - Memory/Network Health	30-Jun-2016 16:30 🍁 💮
Total Nodes: 60	Weakest Node: StoreNode03.SLHOST1	Pct Weak: 35%
Storage Nodes: 24	Worst Network: OLS-StoreNode-1.SLHOST4	Send Queue Size: 53
Rx Success: 99.98%	Worst Memory: ProcessNode08.SLHOST4	Mem Used: 86.61%
	Average over all Process / Storage Nodes	Time Range: All Data 💌
Storade Nodes 0 0 0 0 0 0 0 0 0 0 0 0 0	D:00 00:00:00 00:000 00:	Proc Pub Fail % Proc Rcv Fail % Proc Rcv Fail % Proc Rcv Fail % Proc Rcv Fail % Proc Mem % Avg. CPU % Stor Pub Fail % Stor Rcv Fail % Stor Rcv Fail % Stor Mem % Avg. CPU % Avg. CPU % Stor Mem % Stor Mem % Avg. CPU % Stor Mem % S

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

Fields and Data:

Total Nodes	The total number of nodes in the cluster. This includes storage enabled nodes, client nodes, and management (JMX) nodes.
Storage Nodes	The total number of nodes in the cluster which have storage enabled for any cache. This value is equal to the total nodes when replicated caches are being used. The number is less when only distributed cache types are utilized.
Tx Success	The publisher success rate, in percent. The Publish success rate is the percent (%) of packets in the cluster successfully sent by nodes, without having to be resent. A 100% success rate occurs when a packet is sent and does not have to be re-sent. When a packet must be resent the success rate is reduced.
Rx Success	The receiver success rate, in percent. The Receive success rate is the percent (%) of packets in the cluster successfully received by nodes, without being received twice. A 100% success rate occurs when a packet is received once. When a packet is received twice the success rate is reduced.

Weakest Node	The node vo often points most often a The format c	oted by Coherence as the weakest in the cluster. The Weakest Node to a server/node that is causing performance issues. The node value appears in the "weakest node" attribute of all the JMX "node" objects. of this string is <node address="" ip="">:< Node Port >/<nodeid< b="">>.</nodeid<></node>	
	Weak	The percent of weakest.	of the Coherence nodes that "elected" the node as the
Worst Network	The node that	at has the long	gest network queue in the cluster.
	Send Queue	The number of packets sent receive an ac automatically	of packets currently scheduled for delivery, including and still awaiting acknowledgment. Packets that do not knowledgment within the ResendDelay interval are resent.
Worst Memory	The node that	at has the lowest available memory of any node in the cluster.	
	Mem Used	The percent of	of memory consumed on the Worst Memory node.
Average over all Process / Storage Nodes	Trend Grap The trend gr nodes.	hs raphs show aggregated performance metrics for storage and process	
	Time Range	Select a time Minutes to L range, click C	range from the drop down menu varying from 2 .ast 7 Days , or display All Data . To specify a time calendar [].
		By default, the the time rang time from the using the follo Aug 21, 201 Use the navig time period. I from the Tim Click Restore current time.	Select or Enter Date and Time: Image: Restore to Now Image: OK Apply Cancel Image: OK OK Apply Cancel Image: OK Apply Cancel The time range end point is the current time. To change the calendar or enter the date and time in the text field the calendar or enter the date and time in the text field powing format: MMM dd, YYYY HH:MM. For example, the calendar or enter the time period is determined by your selection the NOTE: The time period is determined by your selection the Range drop-down menu.
	Process Nodes	Publish Failures and Received Failures	Indicates the trending of process node publisher and receiver failure rates. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole. The Weakest Node information often points to the server/nodes that are the cause of these issues.
		Memory Utilization%	Indicates the trending of process node memory utilization. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole.

Storage Nodes	Publish Failures and Received Failures	Indicates the trending of storage node publisher and receiver failure rates. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole. The Weakest Node information often points to the server/nodes that are the cause of these issues.
	Memory Utilization%	Indicates the trending of storage node memory utilization. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole.

Stability Metrics

Use this display to troubleshoot nodes joining and leaving the cluster, and view HA status for cache services. This display presents information about node up times and the stability of the cluster.





Fields and Data:

Cluster Name	Select a cluster from the drop-down menu.
Data Grid Total Nodes	The total number of nodes being monitored. This includes storage enabled nodes, client nodes, and management (JMX) nodes.

Node

Startup History

Location

Storage
NodesThe total number of nodes in the cluster which have storage enabled for any
cache. This value is equal to the total nodes when replicated caches are being
used. The number is less when only distributed cache types are utilized.

Use this table to identify nodes that have departed and returned to the cluster recently. This table contains a list of nodes in the cluster, sorted by start time (the most recently created node is listed first).

A unique identifier for each node. It is defined as:

member_name.machine.rack.site. The date and time that the node joined the cluster. Start Time Indicates whether storage is enabled (0 or 1). StorageFlag The short member id that uniquely identifies this member. ١d Avail MB The amount of available memory for this node, in megabytes. The maximum amount of memory for this node, in megabytes. Max MB The cumulative number of packets sent by this node since the Pkts Sent node statistics were last reset. Delta The number of packets sent by this node since the last update. Pkts Rcvd The cumulative number of packets received by this node since the node statistics were last reset. The number of packets received by this node since the last Delta update. The cumulative number of duplicate packets received by this node Pkts Rptd since the node statistics were last reset. Delta The number of duplicate packets received by this node since the last update. The cumulative number of packets resent by this node since the Pkts Resent node statistics were last reset. The number of packets resent by this node since the last update. Delta The publisher success rate for this node since the node statistics Pub Succ were last reset. Publisher success rate is a ratio of the number of Rate packets successfully delivered in a first attempt to the total number of sent packets. A failure count is incremented when there is no ACK received within a timeout period. It could be caused by either very high network latency or a high packet drop rate. The receiver success rate for this node since the node statistics **Rec Succ** were last reset. Receiver success rate is a ratio of the number of Rate packets successfully acknowledged in a first attempt to the total number of received packets. A failure count is incremented when a re-delivery of previously received packet is detected. It could be caused by either very high inbound network latency or lost ACK packets. The member name for this node. Member The machine name for this node. Machine The rack name for this node. Rack Site The site name for this node. The process name for this node. Process The unicast address. This is the IP address of the node's Uni Addr DatagramSocket for point-to-point communication.

- **Uni Port** The unicast port. This is the port of the node's DatagramSocket for point-to-point communication.
- **RoleName** The role name for this node.
- ProductEditi
onThe product edition this node is running. Possible values are:
Standard Edition (SE), Enterprise Edition (EE), Grid Edition (GE).

Membership Trends Track the total number of nodes and the total number of storage nodes in the cluster for the duration of the user session. These lines are normally unchanging or "flat". If there are fluctuations in this graph, check the debugging guide for appropriate actions.

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

Select or Enter Date and Time:
Restore to Now
Ok Apply Cancel

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

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

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

Departed
NodesTrack departed nodes by IP address, port number and time last seen.LocationA unique identifier for each node. It is defined as:
member_name.machine.rack.site.HostNameThe name of the host on which the node resides.IPThe node IP address.PortThe unicast port the node used while in the cluster. This is the port
of the node's DatagramSocket for point-to-point communication.Last SeenThe date and time that the node left the cluster.

All Services History

Use this display to assess utilization of cache capacity, over time, by all services in a cluster. Analyze load distribution across services and caches, check for bottlenecks and quickly identify services that need more threads. Answer questions such as:

- Is their enough cache capacity available for the service?
- Is their enough storage capacity available for the service?

Use the mouseover tool-tip to see how many caches the service runs on, and data for the selected metric.

n/a:DemoCluster One All Service:	s - History Heatmap	30-Jun-2016 16:34 🚸 Data OK 💠 💮
Service Metric: CPU%	-	Range: All Data 💽
Enable MouseOver History Heatmap of Selected Metric O seore Seore	Over Time by Service - Total	Cover All Nodes Ver All Nodes Time Stamp: 06/26/16 16:00:00 Service: OnLineStoreCache Cache Count: 2 CPU%: 5.38 Active Threads: 0 Messages (delta): 25,857 Request Avg Duration: 0.64 Request Avg Duration: 0.64 Request Avg Duration: 0.64 Request Pending Count: 0 Task Backlog: 1
 Title Bar (possible features are): Open the previous and upper display. Open an instance of this display in a new window. Open the online help page for this display. Menu , Table open commonly accessed displays. 	Data OK Data conne Server is not receiving d receiving data from the data source is connected 23-Mar-2017 12:04 Curren might indicate the Monit and green Data OK icon current and valid.	ection state. Red indicates the Data lata or the Display Server is not Data Server. Green indicates the d. Int date and time. Incorrect time for stopped running. Correct time in is a strong indication that data is

6,047

Filter By: The display might include these filtering options:

The number of items currently in the display.

Service Metric:	 Choose a service metric for which to display data in the heatmap. Use the movement of the view metrics. Identify a service with high utilization. Perform analysis by clicking One to view the "Single Service History" display. CPU% Percent of CPU utilization in the specified time range. 		
	Requests The number of client requests issued to the cluster in the time range. This metric is a good indicator of end-user u the service.		
	Messages	The number of messages for the given node in the specified time range.	
	ActiveThr eads	The number of threads in the service thread pool, not currently idle.	

Open the Alert Views - RTView Alerts Table display.

The size of the backlog queue that holds tasks scheduled to be TaskBackI executed by one of the service threads. Use this metric for og determining capacity utilization for threads running on a service. For example, if the service has a high **TaskBacklog** rate and a low amount of CPU available, consider increasing the number of threads for the service to improve performance.

The number of pending requests issued by the service. RequestP endingCo unt

The average duration (in milliseconds) of an individual request issued RequestA by the service since the last time the statistics were reset. verageDur ation

Time Range

Enable

er

of

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



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

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

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

Select this option to make service details visible upon mouseover. MouseOv

Use the heatmap to view utilization trends for all services, over time, and quickly History identify heavy usage, indicated by a dark color (by default, dark green). Look for a Heatmap consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing service details in the "Single Service Summary" display. Selected Metric by Service

Two heatmaps, one for Process Nodes and another for Storage Nodes, show utilization trends for the selected metric, for all services running in the cluster. Each row represents a service. Cells in a row are sized uniformly. Each column represents a time period (typically in 10 second intervals). The color of the row cells represent the relative value of the selected service Metric, where a darker shade is a larger value.

Use the mouseover tool-tip to see how many caches the service runs on, and data for the selected metric.

Services	Each row represents a service. The color of the cells represents the
on	relative value of the selected Service Metric, where a darker shade is
Process	a larger value. The size of the cells are uniform as they each
Nodes	represent one process node. Use the mouseover tool-tip to see how
	many caches the service runs on, and data for the selected metric.

Services on Storage Nodes	Each row represents a service. The color of the cells represents the relative value of the selected Service Metric, where a darker shade is a larger value. The size of the cells are uniform as they each represent one storage node. Use the mouseover tool-tip to see how many caches the service runs on, and data for the selected metric.
- · · ·	

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

All Caches History

Use this display to assess capacity utilization, over time, for all caches in a cluster. Analyze load distribution, check for bottlenecks and quickly identify caches with high usage. Answer questions such as:

- Is the cluster using what I expect?
- Is the cluster using it in a uniform scale?

Use the mouseover tool-tip to see the name of the cache and data for the selected metric.

← n/a:DemoCluster	All Caches - History Heatmap	30-Jun-2016 16:37 🚸 Data OK 🛛 🕂 💮
Service: DistributedCache	Metric: Total Gets (delta)	Range: All Data 🗾
Service: DistributedCache	o of Selected Metric Over Time for Cach	Range: All Data nes by Service ✓ Log Scale Time Stamp: 6/28/16 00:00:00 Service: DistributedCache Cache: near-foo Cache: Store Total Gets (delta): 24,968 Cache Hisses (delta): 0 Store Read Ms (delta): 0 Store Read Ms (delta): 0 Store Writes (delta): 0
06/22/16 10:00:00		06/30/16 16:37:45

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

Filter By:			
Cluster:	Select a cluster for which to display data in the heatmap.		
Service:	Select a service for which to display data in the heatmap.		
Metric:	Select a metric for which to display data in the heatmap.		
	Total Gets	The total number of requests for data from this cache.	
	Total Puts	The total number of data stores into this cache.	
	Cache Hits	The total number of successful gets for this cache.	
	Cache Misses The total number of failed gets for this cache. This metric indicates whether cache utilization is effective. For example, ho often requests are made for data that does not exist in the cache. If a cache has a high rate of misses, consider performin a lower level analysis by viewing the cache in the "Single Cach Summary" display. Check the metrics for Size, Evictions and Misses to determine whether more capacity is needed.		
	Cache Size	The total number of objects in the cache.	
	StoreFailure s (Delta)	The total number of store failures on this cache since the last data sample.	
	StoreReads (Delta)	The total number of load operations on this cache since the last data sample.	
	StoreReadMi Ilis (Delta)	The cumulative amount of time (in milliseconds) of load operations for this cache since the last data sample.	
	StoreWrites (Delta)	The total number of store and erase operations for this cache since the last data sample.	
	StoreWrites Millis (Delta)	The cumulative amount of time (in milliseconds) of store and erase operations on this cache since the last data sample.	
	Total Gets	The total number of requests for data from this cache.	
Range:	Select a time ra	inge from the drop down menu varying from 2 Minutes to Last 7	



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

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

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

AppName: Choose an AppName to show data for in the display.

7

Fields and Data:

AppSlice Information	Last Update:	The date and time the data was last updated.
	Completed:	The total number of completed processes summed across all processes in one AppSlice of the application.
	Suspended:	The total number of suspended processes
	Failed:	The total number of failed processes
	Created Rate:	The number of application processes created per second.
	Failed Rate:	The number of failed application processes per second.
	Avg Exec:	The average number of seconds for processes to execute.
	Avg Elap:	The average amount of elapsed time, in seconds.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last	

Days, or display All Data. To specify a time range, click Calendar .

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

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

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

Enable Select this option to make cache details visible upon mouseover. MouseOver

History Heatmap of Selected Metric Use the heatmap to view utilization trends for all caches, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing cache details in the "Single Cache Summary" display.

Also look for a dark vertical line, which indicates that all the caches, nodes or services are being used simultaneously. Typically this indicates further analysis is needed.

The heatmap shows cache utilization trends for the selected service and metric, for all caches running in the cluster. Each row represents a cache. Cells in a row are sized uniformly and represent one process node. Each column represents a time period (typically in 10 second intervals). The heatmap is grouped vertically by service. The color of the row cells represent the relative value of the selected service Metric, where a darker shade is a larger value.

Use the mouseover tool-tip to see the name of the cache and data for the selected metric.

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



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

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

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

All Nodes History

Use this display to assess capacity utilization, over time, for all nodes in a cluster. Analyze load distribution, check for bottlenecks and quickly identify nodes with high usage. Use the mouseover tool-tip to see the node hostname and data for the selected metric.



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

Filter By:

Select a cluster for which to display data in the heatmap.		
Click to open the "All Nodes History" display which shows GC Duty Cycle for all the nodes in a cluster.		
Select a metric for which to display data in the heatmap.		
Mem Used%	The percent (%) of memory used by the node.	
Packets Sent Fail%	The percent (%) of packets that had to be resent by this node.	
Packets Rcvd Fail%	The percent (%) of packets that failed to be received by this node.	
	Select a cluster Click to open th the nodes in a o Select a metric Mem Used% Packets Sent Fail% Packets Rcvd Fail%	

DeltaThe number of packets sent by this node since the last dataPackets Sentsample.

DeltaThe number of packets received by this node since the last dataPacketssample.RcvdSample.

Delta Nacks Sent The number of TCMP packets sent by this node since the last data sample. Use this data to troubleshoot communication errors.

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



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

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

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

Enable MouseOver

Select this option to make cache details visible upon mouseover.

History Heatmap of Selected Metric Use the heatmap to view utilization trends for all nodes, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing node details in the "Node Summary" display.

Two heatmaps, one for Process Nodes and another for Storage Nodes, show utilization trends for the selected metric, for all nodes running in the cluster. Each row represents a node. Cells in a row are sized uniformly. Each column represents a time period (typically in 10 second intervals). The color of the row cells represent the relative value of the selected service Metric, where a darker shade is a larger value.

Use the mouseover tool-tip to see the node hostname and data for the selected metric.

Process Each row represents a node. The color of the cells represents the relative value of the selected Service Metric, where a darker shade is a larger value. The size of the cells are uniform. Use the mouseover tool-tip to see the node hostname and data for the selected metric.

	Storage Nodes	Each row represents a node. The color of the cells represents the relative value of the selected Service Metric, where a darker shade is a larger value. The size of the cells are uniform. Use the mouseover tool-tip to see the node hostname and data for the selected metric.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations fo data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.	

Proxy Services

Proxy Services displays present detailed proxy server performance metrics for the cluster. Use the Proxy Services displays to quickly identify overloaded proxy services and locate the extend client connection causing the issue.

Proxy Services performance metrics include: CPU%, Requests, Request Average Duration, Request Pending Count, Task Backlog and Active Threads.

- "Proxy / Extend Overview": Heatmap shows the extend connections and a trend graph shows the total connections and total bytes transferred across all proxies for the selected host or hosts.
- "Proxy / Extend Connections" on page 66: Table shows proxy services data with trend graphs/tables of extend connection detail for a specified location.
- "Proxy / Extend Detail" on page 72: Table shows data for proxy services and extend client connection data, including remote endpoint, time stamp, connect time and outgoing byte backlog.
- "Proxy Nodes History" on page 76: Heatmap shows performance utilization, over time, for all proxy service nodes in the selected cluster.
- "Extend Connections History" on page 78: Heatmap shows performance utilization, over time, for all extend connections in the selected cluster.
Proxy / Extend Overview

Heatmap shows performance utilization and a trend graph shows the total connections and total bytes transferred for all proxy services for the selected host or hosts.





Cluster:	Select a cluster for which to display data in the heatmap.
----------	--

Hosts Click to open display that shows GC Duty Cycle for all the nodes in a cluster.

Metric: Select a metric for which to display data in the heatmap.

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

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

Current
Bytes SentTotal number of bytes sent by the selected proxy in the time
range specified.

CurrentTotal number of bytes received by the selected proxy in the timeBytesrange specified.ReceivedReceived

Proxy CPU% The average percent CPU utilization for the selected proxy.

Bytes The number of pending bytes in the Extend outgoing queue. Backlog

Proxy Bytes The number of pending bytes in the Proxy outgoing queue. **Backlog**

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

Select or Enter Date and Time:
Restore to Now
Ok Apply Cancel

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

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

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

Total Connections

The number of extend clients connected to the selected proxy.

Proxy / Extend Connections

Table shows proxy services data, including connections, CPU usage and bytes sent and received, and a trend graph displays messages and bytes sent and received for the selected remote endpoint.

The table data is the result of joins of metric from the following Coherence MBeans: Service and ConnectionManager. For details on attributes of these MBeans go to: http://download.oracle.com/otn_hosted_doc/coherence/350/com/tangosol/net/management/Registry.html.



displays. and green Data OK icon is a strong indication that data is current and valid.
_

Filter By:

The display might include these filtering options:

Cluster:	Select a cluster from the drop-down menu.				
Proxy Services	Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site.			
	HostIP	The IP address of the host where the proxy service resides.			
	Alert Level	 The maximum level of alerts in the row: Red indicates that one or more exceeded their ALARM LEVEL threshold. Yellow indicates that one or more exceeded their WARNING LEVEL threshold. Green indicates that none have exceeded their alert thresholds. 			

Alert Count	The number of alerts in the row.
Running	Indicates that the proxy service is running when selected.
Connections	The number of extend clients connected to the selected host or hosts.
CPU%	The average percent CPU utilization for each proxy service in the cluster.
Bytes Sent	The number of bytes sent by the proxy service since the proxy service joined the cluster.
Delta	The number of bytes sent by the proxy service since the last data sample.
Backlog	The size (in kilobytes) of the backlog queue.
Bytes Rcvd	The number of bytes received by the proxy service since the proxy service joined the cluster.
Delta	The number of bytes received by the proxy service since the last data sample.
MsgsSent	The number of messages sent by the proxy service since the proxy service joined the cluster.
Delta	The number of messages sent by the proxy service since the last data sample.
Backlog	The size of the backlog queue that holds messages scheduled to be sent by one of the proxy service pool threads.
Msgs Rcvd	The number of messages received by the proxy service since the proxy service joined the cluster.
Delta	The number of messages received by the proxy service since the last data sample.
Tasks	The number of tasks performed by the proxy service since the last time the statistics were reset.
RequestAve rageDuratio n	The average duration (in milliseconds) of an individual synchronous request issued by the proxy service since the last time the statistics were reset.
RequestMax Duration	Maximum duration (in milliseconds) of an individual proxy service request since the last time the statistics were reset.
RequestTot alCount	The number of requests issued and received by the proxy service.
TaskAverag eDuration	The average duration (in milliseconds) of an individual task execution.
TaskBacklog	The size of the backlog queue that holds tasks scheduled to be executed by one of the proxy service pool threads.
TaskCount	The number of tasks performed by the proxy service since the last time the statistics were reset.
TaskHungCo unt	The total number of currently executing hung tasks.
TaskHungD uration	The longest currently executing hung task duration in milliseconds.
TaskHungTa skId	The id of the of the longest currently executing hung task.

TaskHungTh resholdMilli s	The duration (in milliseconds) that a proxy service task can execute before it is considered hung. Note that a posted task that has not yet started is never considered as hung.
TaskMaxBac klog	The maximum size of the proxy service backlog queue since the last time the statistics were reset.
TaskTimeou tCount	The total number of timed-out proxy service tasks since the last time the statistics were reset.
RequestPen dingCount	The number of pending proxy service requests.
RequestPen dingDuratio n	The average duration (in milliseconds) that an individual proxy service request waits before being executed.
RequestTim eoutCount	The total number of timed-out proxy service requests since the last time the statistics were reset.
RequestTim eoutMillis	The duration (in milliseconds) for a proxy service request to reach the specified timeout threshold.
TaskTimeou tMillis	The default timeout value (in milliseconds) for tasks that can be timed-out but do not explicitly specify the task execution timeout value.
IncomingBu fferPoolSize	The number of buffers in the incoming pool.
ThreadAban donedCount	The number of abandoned threads from the proxy service thread pool. A thread is abandoned and replaced with a new thread if it executes a task for a period of time longer than execution timeout and all attempts to interrupt it fail.
ThreadCoun t	The number of threads in the proxy service thread pool.
ThreadI dleC ount	The number of currently idle threads in the proxy service thread pool.
AverageActi veThreadCo unt	The average number of proxy service active threads, not currently idle, since the last time the statistics were reset.
ThreadAver ageActiveCo unt	The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.
AverageTas kDuration	The average duration (in milliseconds) to perform a proxy service task since the last time the statistics were reset.
MaximumBa cklog	The maximum size of the backlog queue since the last time the statistics were reset.
Throughput	The amount of data (in kilobytes) that is transferred by the proxy service.
Throughput Inbound	The amount of data (in kilobytes) that is transferred from clients to the proxy service.
Throughput Outbound	The amount of data (in kilobytes) that is transferred from the proxy service to clients.
IncomingBu fferPoolCap	The size (in kilobytes) of the proxy service incoming buffer pool.

	OutgoingBu fferPoolCap acity	The size (in kilobytes) of the proxy service outgoing buffer pool.			
	OutgoingBu fferPoolSize	The number of buffers in the proxy service outgoing pool.			
	nodeld	The unique identifier for the proxy service.			
	RefreshTim e	The timestamp when this model was last retrieved from a corresponding node. For local servers it is the local time.			
	HostName	The name of the host where the proxy service resides.			
	MemberNa me	A specified, unique name of the host where the proxy service resides.			
	SeniorMemb erId	The proxy service senior member id. If the proxy service is not running, it is -1.			
	Rate	The number of errors accumulated per second.			
Execution	Min	The shortest execution time of any process instance, in milliseconds.			
	Мах	The longest execution time of any process instance, in milliseconds.			
	Average	The average execution time for all completed process instances, in milliseconds.			
	Current	The amount of time accumulated this update cycle.			
	Rate	The amount of time accumulated per second.			
Elapsed	Min	The shortest elapsed time of any process instance, in milliseconds.			
	Max	The longest elapsed time of any process instance, in milliseconds.			
	Average	The average elapsed time for all completed process instances, in milliseconds.			
	Current	The amount of elapsed time accumulated this update cycle.			
	Rate	The amount of elapsed time accumulated per second.			
Selected Proxy	This field is p	opulated by the selection made in the Proxy Services table.			
Selected Endpoint	This field is populated by the selection made in the Remote Endpoint table.				

Trend Graphs

Select a host from the Proxy Services table and a connection from the Remote Endpoint table. This table is populated by the selection made in the Proxy Services table. Alert Level shows the maximum level of alerts in row:

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

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

Green indicates that none have exceeded their alert thresholds.

Msgs Received: Traces the number of messages received by the selected proxy service from the remote endpoint.

Msgs Sent: Traces the number of bytes received by the selected proxy service from the remote endpoint.

Bytes Received: Traces the rate at which the application is accumulating process execution time, in milliseconds per second.

Bytes Sent: Traces the number of executed activities per second.

All Activities Exec Time/sec: Traces the number of bytes sent by the selected proxy service to the remote endpoint.

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



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

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

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

Proxy / Extend Detail

Table shows data for proxy services and extend client connection data, including remote endpoint, time stamp, connect time and outgoing byte backlog.

🗲 n/a:DemoCluster	Р	roxy / E	xtend De	tail			30-Jun-2010	6 16:49 📫	Data OK	+ 📀
Proxy Services										
Location	Service Name		HostIP		Alert Level	Alert Count	Running	Conns	CPU %	Tota
ProxyNode01.SLHOST2	ProxyService	SLHOST?	2:9099		- ()		0 🔽	1	0.(0 2
ProxyNode01.SLHOST3	ProxyService	SLHOST	3:9099				0	1	0.0	0 2
ProxyNode01.SLHOST4	ProxyService	SLHOST/	4:9099				0 🔽	0	0.0	0 4
ProxyNode02.SLHOST2	ProxyService	SLHOST?	2:9099				0	0	0.0	0
ProxyNode02.SLHOST3	ProxyService	SLHOST	3:9099				0 🔽	0	0.0	0
ProxyNode02.SLHOST4	ProxyService	SLHOST/	4:9099				0	0	0.0	0
< III										•
	Extend Client Cor	nnections	s for Proxy	vNode	e02.SLH	OST3;*				
	Remote	Alert	Alert			(Connect Tim	e Byt	es N	lessage
Location	Endpoint	Level	Count	Tr	me Stam	(P	(mins)	Back	dog	Backlog
ProxyNode02.SLHOST3	192.168.200.103:1131		0	06/1	7/16 01:0	03:29	2,098	3.6	0	
ProxyNode02.SLHOST3	192.168.200.103:1136		0	06/1	7/16 01:0	03:34	2,098	3.5	0	
ProxyNode02.SLHOST3	192.168.200.103:1128		0	06/1	9/16 01:0	03:33	659).1	0	
ProxyNode02.SLHOST3	192.168.200.103:1121		0	06/1	9/16 01:0	03:27	1,437	.2	0	
ProxyNode02.SLHOST3	192.168.200.103:1124	6	0	06/2	2/16 01:0	03:32	2,818	3.6	0	
ProxyNode02.SLHOST3	192.168.200.103:1132		0	06/2	2/16 01:0	03:37	2,098	8.7	0	
ProxyNode02.SLHOST3	192.168.200.103:1136	(a)	0	06/2	.4/16 01:0	03:40	1,018	3.4	0	
ProxyNode02.SLHOST3	192.168.200.103:1140		0	06/2	4/16 01:0	03:45	1,018	3.3	0	
ProxyNode02.SLHOST3	192.168.200.103:1121	- ()	0	06/2	6/16 01:0	03:29	1,437	.2	0	
ProxyNode02.SLHOST3	192.168.200.103:1127		0	06/2	6/16 01:0	03:34	1,437	.1	0	
▲ III										•

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

Cluster

Select a cluster from the drop-down menu.

Proxy Services

Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site.
HostIP	The IP address of the host where the proxy service resides.
Running	Indicates that the proxy service is running when selected.
Connections	The number of extend clients connected to the selected host or hosts.
CPU%	The average percent CPU utilization for each proxy service in the cluster.

Bytes Sent	The number of bytes sent by the proxy service since the proxy service joined the cluster.
Delta	The number of bytes sent by the proxy service since the last data sample.
Backlog	The size (in kilobytes) of the backlog queue.
Bytes Rcvd	The number of bytes received by the proxy service since the proxy service joined the cluster.
Delta	The number of bytes received by the proxy service since the last data sample.
MsgsSent	The number of messages sent by the proxy service since the proxy service joined the cluster.
Delta	The number of messages sent by the proxy service since the last data sample.
Backlog	The size of the backlog queue that holds messages scheduled to be sent by one of the proxy service pool threads.
Msgs Rcvd	The number of messages received by the proxy service since the proxy service joined the cluster.
Delta	The number of messages received by the proxy service since the last data sample.
Tasks	The number of tasks performed by the proxy service since the last time the statistics were reset.
RequestAverage Duration	The average duration (in milliseconds) of an individual synchronous request issued by the proxy service since the last time the statistics were reset.
RequestMaxDura tion	Maximum duration (in milliseconds) of an individual proxy service request since the last time the statistics were reset.
RequestPending Count	The number of pending proxy service requests.
RequestPending Duration	The average duration (in milliseconds) that an individual proxy service request waits before being executed.
RequestTimeout Count	The total number of timed-out proxy service requests since the last time the statistics were reset.
RequestTimeout Millis	The duration (in milliseconds) for a proxy service request to reach the specified timeout threshold.
RequestTotalCou nt	The number of requests issued and received by the proxy service.
TaskAverageDur ation	The average duration (in milliseconds) of an individual task execution.
TaskBacklog	The size of the backlog queue that holds tasks scheduled to be executed by one of the proxy service pool threads.
TaskCount	The number of tasks performed by the proxy service since the last time the statistics were reset.
TaskHungCount	The total number of currently executing hung tasks.
TaskHungDurati on	The longest currently executing hung task duration in milliseconds.

TaskHungTaskId	The id of the of the longest currently executing hung task.
TaskHungThresh oldMillis	The duration (in milliseconds) that a proxy service task can execute before it is considered hung. Note that a posted task that has not yet started is never considered as hung.
TaskMaxBacklog	The maximum size of the proxy service backlog queue since the last time the statistics were reset.
TaskTimeoutCou nt	The total number of timed-out proxy service tasks since the last time the statistics were reset.
TaskTimeoutMilli s	The default timeout value (in milliseconds) for tasks that can be timed-out but do not explicitly specify the task execution timeout value.
IncomingBufferP oolSize	The number of buffers in the incoming pool.
ThreadAbandone dCount	The number of abandoned threads from the proxy service thread pool. A thread is abandoned and replaced with a new thread if it executes a task for a period of time longer than execution timeout and all attempts to interrupt it fail.
ThreadCount	The number of threads in the proxy service thread pool.
Thread I dleCount	The number of currently idle threads in the proxy service thread pool.
AverageActiveTh readCount	The average number of proxy service active threads, not currently idle, since the last time the statistics were reset.
ThreadAverageA ctiveCount	The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.
AverageTaskDur ation	The average duration (in milliseconds) to perform a proxy service task since the last time the statistics were reset.
MaximumBacklo g	The maximum size of the backlog queue since the last time the statistics were reset.
Throughput	The amount of data (in kilobytes) that is transferred by the proxy service.
Throughput I nbo und	The amount of data (in kilobytes) that is transferred from clients to the proxy service.
ThroughputOutb ound	The amount of data (in kilobytes) that is transferred from the proxy service to clients.
IncomingBufferP oolCapacity	The size (in kilobytes) of the proxy service incoming buffer pool.
OutgoingBufferP oolCapacity	The size (in kilobytes) of the proxy service outgoing buffer pool.
OutgoingBufferP oolSize	The number of buffers in the proxy service outgoing pool.
nodeld	The unique identifier for the proxy service.
RefreshTime	The timestamp when this model was last retrieved from a corresponding node. For local servers it is the local time.
HostName	The name of the host where the proxy service resides.

	MemberName	A specified, unique name of the host where the proxy service resides.
	SeniorMember I d	The proxy service senior member id. If the proxy service is not running, it is -1.
Extend Client Connections	Select a row from th	e Proxy Services table to populate client data in the table.
	Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site.
	RemoteEndpoint	The IP address of the client.
	Timestamp	The date and time (in cluster time) that this client joined the proxy service.
	Connect Time (mins)	The duration (in minutes) the client has been connected to the proxy service.
	OutgoingByteBa cklog	The size of the backlog queue (in bytes) that holds outgoing bytes scheduled to be executed by one of the proxy service pool threads for the client.
	OutgoingMessag eBacklog	The number of messages in the backlog queue that holds outgoing messages scheduled to be sent to the client by one of the proxy service pool threads.
	TotalBytesReceiv ed	The number of bytes received from the client by the proxy service since the client connected to the proxy service.
	Delta	The number of bytes received from the client by the proxy service since the last data sample.
	TotalBytesSent	The number of bytes sent to the client by the proxy service since the client connected to the proxy service.
	Delta	The number of bytes sent to the client by the proxy service since the last data sample.
	TotalMessagesR eceived	The number of messages received from the client by the proxy service since the client connected to the proxy service.
	Delta	The number of messages received from the client by the proxy service since the last data sample.
	TotalMessagesSe nt	The number of messages sent to the client by the proxy service since the client connected to the proxy service.
	Delta	The number of messages sent to the client by the proxy service since the last data sample.
	UUID	The unique identifier for the extend client application.
	nodeld	The unique identifier for the proxy service the extend client is connected to.

Proxy Nodes History

Heatmap shows performance utilization, over time, for all proxy service nodes in the selected cluster. Use this display to assess performance, over time, for all proxy service nodes in a cluster. Analyze load distribution, check for bottlenecks and quickly identify proxy service nodes with high usage.

fn/a:DemoCluster	Proxy Nodes - History Heatmap	11-Jul-2016 10:43 🚸 Data OK 💠 💮
	Metric: Refresh Milliseconds	Range: 5 Mins 🔽
Enable MouseOver	History Heatmap of Selected Metric Over Time for Proxy	Nodes Log Scale
Proxy Nodes	Time Stamp: Jul 11, 2016 10:39: Location: ProxyNode02.SLHOST3 Requests Pending: 0 Connections: 2 Refresh Milliseconds (delta): 11,3 Delta Bytes Sent: 4,753 Delta Bytes Rovd: 5,992 Bytes Sent per Second: 425.44 Bytes Rovd per Second: 536.34	57 AM 191
07/11/16 10:37:35		07/11/16 10:43:30

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

Select a cluster from the drop-down menu.		
Select a metric from the drop-down menu.		
Request Pending	The number of pending requests issued by the node.	
Connections	Total number of connection for the node.	
Refresh Milliseconds	The amount of time, in milliseconds, since the last data sample.	
Delta Bytes Sent	Total number of bytes sent by the node since the last data sample.	

Cluster

Metric

Delta Bytes Rcvd Total number of bytes received by the node since the last data sample.

Bytes Sent Per Total bytes sent, per second, by the node. Second

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



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

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

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

Enable Select this option to make details visible upon mouseover. MouseOver

- Log Scale Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- **Proxy Nodes Heatmap Use the heatmap to view utilization trends for all Process and Storage nodes, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing node details in the "Node Summary**" display.

Use the mouse over tool-tip to see the node hostname and data for the selected metric.

Extend Connections History

Heatmap shows performance utilization, over time, for all extend connections in the selected cluster.

← n/a:DemoCluster <	Proxy Nodes - History Heatmap	11-Jul-2016 10:43 🚸 Data OK + 👔
Metric:	Refresh Milliseconds	Range: 5 Mins 🔽
Sapon History Heatman Sapon 07/11/16 10:37:35	Time Stamp: Jul 11, 2016 10:39: Location: ProxyNode02.SLHOST3 Requests Pending: 0 Connections: 2 Refresh Milliseconds (delta): 11,31 Delta Bytes Sent: 4,753 Delta Bytes Sent: 4,753 Delta Bytes Sent: 4,753 Delta Bytes Revd: 5,992 Bytes Sent per Second: 425.44 Bytes Revd per Second: 536.34	Nodes Log Scale
 Title Bar (possible features are): Open the previous and upper displa Open an instance of this display in a new Open the online help page for this display Menu , Table open commonly accidisplays. 6,047 The number of items currently in the 	Data OK Data conne Server is not receiving d vindow. data source is connected 23-Mar-2017 12:04 Curren might indicate the Monit and green Data OK icon current and valid. e display. Open the Alert View	ction state. Red indicates the Data ata or the Display Server is not Data Server. Green indicates the l. It date and time. Incorrect time or stopped running. Correct time is a strong indication that data is ws - RTView Alerts Table display.

Cluster	Select a cluster from the drop-down menu.	
Metric	Select a metric from the drop-down menu.	
	Delta Bytes Sent	Total number of bytes sent by the node since the last data sample.
	Delta Bytes Rcvd	Total number of bytes received by the node since the last data sample.
	Delta Messages Sent	Total number of messages sent by the node since the last data sample.
	Delta Messages Rcvd	Total number of messages received by the node since the last data sample.

Bytes Sent per Total bytes sent, per second, by the node. Second

Bytes per Second Total bytes received, per second, by the node.

Msgs Sent per Total messages sent, per second, by the node. Second

Msgs Rcvd per Total messages received, per second, by the node. Second

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



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

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

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

Enable Select this option to make details visible upon mouseover.

MouseOver

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

Proxy Nodes Heatmap Use the heatmap to view utilization trends for all Process and Storage nodes, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing node details in the Single Node -Summary display.

Use the mouse over tool-tip to see the node hostname and data for the selected metric.

Cache Services

Cache Services displays present detailed service performance metrics for the cluster. Use the Cache Services displays to quickly identify overloaded services and locate the client connection causing the issue.

These displays show metrics for all cache services, including: CPU%, Requests, Request Average Duration, Request Pending Count, Task Backlog and Active Threads.

- "Single Service Summary" on page 80: Trend graphs show performance metrics for a single service aggregated across all nodes.
- "Service Metrics Overview" on page 84: Heatmap shows overview of the current behavior of the cluster, displaying metrics across nodes in the cluster for a selected service or for all services. Enables you to determine if the behavior of the cluster is balanced across all nodes or identify if some nodes are hot spots.
- "Service Metric Heatmap" on page 85: Heatmap shows current value of a selected metric, selected by service, across the cluster. Enables you to determine if the behavior of the cluster, for the selected metric, is balanced or identify if some nodes are hot spots.
- "Single Service History" on page 87: Use this display to perform low-level analysis of service capacity utilization, over time, per node. Heatmap enables you to view the impact of events across the cluster as well as the relative historical performance of nodes across the cluster.
- "Cache Service Detail" on page 89: Table view of attributes of a selected service for a selected host for nodes. Attribute values can be ordered to identify the nodes with the highest and lowest values of interest.

Single Service Summary

This display shows performance metrics for a single service aggregated across all nodes.



Title Bar (possible features are):	🕼 Data OK Data connection state. Red indicates the Data
Open the previous and upper display.	Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data server is server to determine the server.
Open an instance of this display in a new window.	data source is connected.
 Open the online help page for this display. Menu , Table open commonly accessed displays. 	might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
6,047 The number of items currently in the display.	Open the Alert Views - RTView Alerts Table display.

Cluster	Select a cluster to display.	
Service	Select a service to display.	
Storage Nodes	Select to display storage node data in the trend graphs of this display.	
Process Nodes	Select to display process node data in the trend graphs of this display.	
Caches	The number of caches managed by the service.	
Туре	The type of cache.	
Storage Nodes	The number of storage nodes in the cache.	
Process Nodes	The number of process nodes in the cache.	
Status	The high availability status of the service:	
	• ENDANGERED: There is potential data loss in the cluster if a node goes offline.	
	NODE-SAFE: There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.	
	• MACHINE-SAFE: There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.	
	• RACK-SAFE: There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.	
	SITE-SAFE: There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.	

Requests	Requests executed by the service.
	Total The number of requests executed.
	Rate / Delta Use the Use Rates checkbox to toggle between two value types: Rate and Delta (as labeled in the display upon selection).
	When the Use Rates (checkbox) is NOT selected the Delta values are shown here and in the trend graphs. Delta is the difference in the value since the last sample. When the Use Rates (checkbox) is selected the Rate values are shown here and in the trend graphs. Rate is the value per second. The Rate value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the Rate value does not vary if the sample period changes (whereas the Delta value does vary). The Rate value enables you to directly compare rates on systems with different sample periods. Pending The number of pending requests.
Messages	Messages executed by the service.
J. J	Total The number of messages executed.
	Rate / Delta Use the Use Rates checkbox to toggle between two value types: Rate and Delta (as labeled in the display upon selection).
	When the Use Rates (checkbox) is NOT selected the Delta values are shown here and in the trend graphs. Delta is the difference in the value since the last sample. When the Use Rates (checkbox) is selected the Rate values are shown here and in the trend graphs. Rate is the value per second. The Rate value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the Rate value does not vary if the sample period changes (whereas the Delta value does vary). The Rate value enables you to directly compare rates on systems with different sample periods.
	Req Avg Duration The average amount of time to process messages.
Tasks	Tasks performed by the service.
	Count The number of tasks performed.
	Backlog The number of tasks scheduled to be executed by one of the service threads.
	Queue The Write Back Queue total across all caches on the service.
Threads	Threads on the service. Count The number of threads on the service. Active The number of threads in the service not currently idle. Avg CPU% The average amount of CPU usage (%) for the service.

Storage / Process Node Totals The trend graphs show aggregated performance metrics for storage or process nodes. Choose **Storage Nodes** or **Process Nodes** at the top of this display.

Use Rates Select to show **Rate** values in the **Requests and Messages** fields and trend graphs.

Rate is the value per second. The **Rate** value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the **Rate** value does not vary if the sample period changes (whereas the **Delta** value does vary). The **Rate** value enables you to directly compare rates on systems with different sample periods. Deselect **Use Rates** to show the **Delta** values in the **Activity - Current (Delta)** fields and trend graphs. **Delta** is the difference in the value since the last sample.

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

Base at Zero Use zero for the Y axis minimum for all graphs.

Time RangeSelect a time range from the drop down menu varying from 2 Minutes to Last7 Days, or display All Data. To specify a time range, click Calendar ____.

1 5	0
Select or Enter Date and Tim	e:
Restore to Now	
Ok Apply Ca	ancel

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

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

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

Service Metrics Overview

Heatmap of Process (non-storage enabled) Nodes and Storage (enabled) Nodes. Size = One Node. Color = Relative Value of Selected Metric.





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

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

Open the Alert Views - RTView Alerts Table display.

Cluster Select a cluster to display.

Service Select a service to display.

Node Labels Select to display node labels.

Process Nodes Storage Nodes Color of the cells represents the relative value of the selected Metric; a darker shade is a larger value. The size of all cells is identical as they each represent one process node.

CPU% Percent of CPU utilization on the given node.

Requests Number of requests issued by the service in the measured period. **Messages** The number of messages for the given node in the measured interval.

Request Average Duration Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

Request Pending Count Number of pending requests issued by the service. **Task Backlog** Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads.

Active Threads Number of threads in the service thread pool, not currently idle.

Service Metric Heatmap

Heatmap of Process (non-storage enabled) Nodes and Storage (enabled) Nodes. Size = Number of Caches in Selected Service, Color = Relative Value of Selected Metric.



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

Cluster	Select a cluster to display.
Service	Select a service to display.
Node Labels	Select to display node labels.
Metric	 CPU% Percent of CPU utilization on the given node. Requests Number of requests issued by the service in the measured period. Request Average Duration Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset. Request Pending Count Number of pending requests issued by the service. Task Backlog Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads. Active Threads Number of threads in the service thread pool, not currently idle.
Node Labels	Select to view node locations. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
Process Nodes: Heatmap of Service Metric	Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value. Size of the cells is based the number of caches in the selected Service for that process node.
Storage Nodes: Heatmap of Service Metric	Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value. Size of the cells is based the number of caches in the selected Service for that process node.

Single Service History

Use this display to perform low-level analysis, node-by-node, of service capacity utilization. Heatmap of Process (non storage enabled) Nodes and Storage (enabled) Nodes. Color = Relative Value of Selected Metric.



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

Cluster	Select a cluster to display.
AII	Click to view the "All Services History" display.
Service	Select a service to display.

Metric CPU% CPU Utilization (as a percent) on the given node.

Requests Number of requests issued by the service in the measured period.

Request Average Duration Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

Request Pending Count Number of pending requests issued by the service. **Task Backlog** Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads.

Active Threads Number of threads in the service thread pool, not currently idle.

[Select or Enter Date and Time:
	Restore to Now
	Ok Apply Cancel

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

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

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

Process Nodes: History Heatmap of Service Metric	Color of the cells represents the relative value of the selected Metric for a given
	process node; a darker shade is a larger value. The value of the Metric is displayed over the specified History for all process
	nodes in the selected Service.
Storage Nodes: History Heatmap of Service Metric	Color of the cells represents the relative value of the selected Metric for a given storage node; a darker shade is a larger value.
	The value of the Metric is displayed over the specified History for all storage nodes in the selected Service.

Cache Service Detail

This display provides a table view of attributes of a selected service for a selected host for nodes. Attribute values can be ordered to identify the nodes with the highest and lowest values of interest.

🔶 n/a:D	emoCluster 🔤	Summary	Cache Se	rvice Det	ail	05-Jul-2	016 12:01 🗳	🖢 Data OK 🛛 🛖	0
Service:	DistributedCache	-	Host: All Ho	sts	-		Class:	All	•
	Location	Serv	vice	Running	StatusHA	Storage	CPU %	Messages	
JmxNode0	1.SLHOST1	DistributedCac	he	~	MACHINE-SAFE		0.0	195,635	
ProcessNo	de01.SLHOST2	DistributedCac	he	V	MACHINE-SAFE		0.0	376,689	
ProcessNo	de01.SLHOST3	DistributedCac	he	V	MACHINE-SAFE		0.0	1,149,263	
ProcessNo	de01.SLHOST4	DistributedCac	he	V	MACHINE-SAFE		0.0	1,131,618	
ProcessNo	de04.SLHOST2	DistributedCac	he	~	MACHINE-SAFE		0.0	188,285	
ProcessNo	de04.SLHOST3	DistributedCac	he	V	MACHINE-SAFE		0.0	574,088	
ProcessNo	de04.SLHOST4	DistributedCac	he	Ľ	MACHINE-SAFE		0.0	564,263	
ProcessNo	de05.SLHOST3	DistributedCac	he	V	MACHINE-SAFE		0.0	1,150,706	
ProcessNo	de05.SLHOST4	DistributedCac	he	Ľ	MACHINE-SAFE		0.0	1,134,618	
ProcessNo	de05n.SLHOST2	DistributedCac	he	V	MACHINE-SAFE		0.0	754,247	=
ProcessNo	de05n.SLHOST3	DistributedCac	he	Ľ	MACHINE-SAFE		0.0	2,334,486	
ProcessNo	de05n.SLHOST4	DistributedCac	he	V	MACHINE-SAFE		0.0	2,248,394	
ProcessNo	de08.SLHOST3	DistributedCac	he	Ľ	MACHINE-SAFE		0.0	573,289	
ProcessNo	de08.SLHOST4	DistributedCac	he	V	MACHINE-SAFE		0.0	565,373	
ProxyNode	01.SLHOST2	DistributedCac	he	2	MACHINE-SAFE		0.0	375,960	
ProxyNode	01.SLHOST3	DistributedCac	he	V	MACHINE-SAFE		0.0	2,431	
ProxyNode	01.SLHOST4	DistributedCac	he	2	MACHINE-SAFE		0.0	1,788	
ProxyNode	02.SLHOST2	DistributedCac	he	V	MACHINE-SAFE		0.0	307	
ProxyNode	02.SLHOST3	DistributedCac	he	V	MACHINE-SAFE		0.0	584,044	
ProxyNode	02.SLHOST4	DistributedCac	he	V	MACHINE-SAFE		0.0	1,127,855	
RTViewOC	DataServer.SLHOST1	DistributedCac	he	Ľ	MACHINE-SAFE		0.0	92	
StoreNode	01.SLHOST1	DistributedCac	he	V	MACHINE-SAFE	V	0.3	632,135	
StoreNode	01.SLHOST2	DistributedCac	he	V	MACHINE-SAFE		0.0	1,223,758	
StoreNode	01.SLHOST4	DistributedCac	he	V	MACHINE-SAFE	~	0.3	3,676,489	
StoreNode	04.SLHOST1	DistributedCac	he	2	MACHINE-SAFE	2	0.1	640,945	
StoreNode	04.S_HOST4	DistributedCac	he	V	MACHINE-SAFE	V	0.5	3,647,694	Ŧ
I		<u>B' ' ' ' 10</u>					0.0		

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

Cluster	Select a cluster to display.
Summary	Click to view the "Single Service Summary" display.
Service	Select a service to display.
Host	Select a host.
Class	Select the type of node to display: All, Storage or Process nodes

Cache Service Detail by Node:

The columns in this table, with the exception of **Location**, come from Service and Node MBeans. **Location** is a unique identifier for each node and defined as: **member_name.machine.rack.site**. For details on attributes of these MBeans go to: http://download.oracle.com/otn_hosted_doc/ coherence/350/com/tangosol/net/management/Registry.html.

- Location A unique identifier for each node. It is defined as: member_name.machine.rack.site.
- Service The name of the service.

Running Indicates that the service is running when checked.

Metric The high availability status of the service:

• ENDANGERED: There is potential data loss in the cluster if a node goes offline.

ODE-SAFE: There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.

• MACHINE-SAFE: There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.

RACK-SAFE: There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.

• SITE-SAFE: There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.

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



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

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

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

Process Nodes: History Heatmap of Service Metric	Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value. The value of the Metric is displayed over the specified History for all process nodes in the selected Service.
Storage Nodes: History Heatmap of Service Metric	Color of the cells represents the relative value of the selected Metric for a given storage node; a darker shade is a larger value.
	The value of the Metric is displayed over the specified History for all storage nodes in the selected Service.

Federated Clusters

Federated Clusters displays present high-level and detailed cache performance metrics for the cluster. Performance statistics are derived from the cluster Destination and Origin MBeans. Destination information shows how efficiently each node in the local cluster participant is sending data to each destination cluster participant. Origin information shows how efficiently each node in the local cluster participant is receiving data from destination cluster participants.

Use these displays to quickly assess total utilization and throughput metrics for all caches in the cluster.

- "Federated Destination Detail" on page 92: Shows current information for all participating nodes for a selected cluster.
- "Federated Destination Summary" on page 95: Shows current information and trended historical rate information.
- "Federated Origin Detail" on page 97: Shows current information for all participating nodes for a selected cluster.
- "Federated Origin Summary" on page 100: Shows current information and trended historical rate information.

Federated Destination Detail

Table shows performance and utilization data, such as bandwidth usage and bytes sent, for Federated Destinations on the selected cluster. Use this display to do high level utilization analysis. Each row is a different Destination MBean. Click a row to see details in the "Federated Destination Summary" display. Sort data by the highest and lowest values of interest by clicking on the column heading.

Cluster1:Cluster1Jmx	Federat	ed Destination Detail	24-Mar-201	17 16:53 🗳 Data OK 📲	ŀ ()
Host: All Hosts					
	Federated	Destination Detail by Nod	e		
Location A	BytesSentSecs	ConnectRetryTimeoutMillis	Connection	CurrentBandwidth	Deltal
1.SLNB50	4,401	0	Cluster1Jmx	0.0	
2.SLNB50	4,615	0	Cluster1Jmx	0.0	
3.SLNB50	6,210	0	Cluster1Jmx	0.0	

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

Filter By:

Cluster: Select a cluster from the drop-down menu.

Host: Select a host from the drop-down menu.

Federated Destination Detail by Node

Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site.
BytesSentSecs	The number of bytes sent per second.
ConnectRetryTimeoutMillis	The configured connect retry timeout.

Connection	The name of the JMX connection used to access the cluster data.
CurrentBandwidth	The current amount of bandwidth being used, in megabits per second, for sending replicate message.
DeltaReplicateAllTotalTime	The difference in the total amount of time the replicateAll request took since the last data sample.
DeltaTIME_STAMP	The amount of time since the last data sample.
DeltaTotalBytesSent	The difference in the total number of bytes sent since the last data sample.
DeltaTotalEntriesSent	The difference in the total number of entries sent since the last data sample.
DeltaTotalErrorResponses	The difference in the total number of error responses since the last data sample.
DeltaTotalMsgSent	The difference in the total number of messages sent since the last data sample.
DeltaTotalMsgUnacked	The difference in the total number of unacknowledged messages since the last data sample.
DeltaTotalRecordsSent	The difference in the total number of records sent since the last data sample.
ErrorDescription	A description of the error. A value exists only if the sender is in an error state.
EstimatedReplicateAllRemai ningTime	The estimated remaining time, in milliseconds, to complete the replicateAll request.
Expired	When checked, this connection is expired due to inactivity.
Geolp	The Geo-IP metadata
HostName	The name of the host.
MaxBandwidth	The maximum amount of bandwidth per second, in megabits, for sending replicate message, where -1.0 means the maximum bandwidth is not specified.
Member	The member information of the destination node.
MemberName	The name of the member.
MsgApplyTimePercentileMill is	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the destination.
MsgNetworkRoundTripTime PercentileMillis	The 90-percentile value, in milliseconds, of the time taken by transmission of replication messages and the corresponding ack messages over the network.
MsgSentSecs	The number of messages sent per second.
Name	The sender name.
ParticipantType	The participant type. Valid types are cluster and interceptor.
RateReplicateAllTotalTime	The number of replicateAll requests per second.
RateTotalBytesSent	The total number of bytes sent per second.
RateTotalEntriesSent	The total number of entries sent per second.
RateTotalErrorResponses	The total number of error responses per second.
RateTotalMsgSent	The total number of messages sent per second.

RateTotalMsgUnacked	The total number of unacknowledged messages per second.
RateTotalRecordsSent	The total number of records sent per second.
RecordBacklogDelayTimePe rcentileMillis	The 90-percentile value , in milliseconds, of the time the journal records are in the cache waiting to be replicated.
ReplicateAllPercentComplet e	The percent of work completed for a replicateAll request.
ReplicateAllTotalTime	The total amount of time the replicateAll request took, in milliseconds.
SendTimeoutMillis	The configured send timeout.
State	The participant state, where: 0 is Ok 1 is Warning 2 is Error
Status	The participant status.
TIME_STAMP	The date and time of the data update.
TotalBytesSent	The total number of bytes sent.
TotalEntriesSent	The total number of cache entries sent.
TotalErrorResponses	The total number of responses with an error.
TotalMsgSent	The total number of replication messages sent. A replication message might contain multiple journal records
TotalMsgUnacked	The total number of unacknowledged replication messages.
TotalRecordsSent	The total number of journal records sent. A journal record might consist of multiple cache entries that are part of the same transaction.
name	The destination cluster name.
nodeid	The unique identifier for the node.
service	The Federated Service name.
subType	The Federated Service sub-type.
type	The Coherence MBean type (Federation, in this case).

Federated Destination Summary

Detailed performance and utilization data, such as bandwidth usage and bytes sent per second, for a Federated Destinations location. Use this display to do low level utilization analysis. Check the metrics for to determine whether more capacity is needed.

Cluster1:Cluster1	Jmx Detail Fed.	Dest Summary	24-Mar-2017 16:53 📫 D	ata OK 💠 😨
Host: All Hosts	Location: 1.SLNB50		~	ld: 1
Participant Type:	CLUSTER	Name:	Destination for SLFederated, Cluster	r2
State:	YIELDING	Msg Apply Time Percentile (ms)	5	Node
Bytes Sent Secs:	4250	Msg Network Round Trip	5	Summary
Connect Retry Timeout (ms):	0	Msgs Sent Secs:	28	
Current Bandwidth:	0.0	Record Backlog Delay	5	
Estimated Replicate All Remaining Time:	-1	Replicate All Percentile Complete	0	
Geo IP:		Replicate All Total Time:	0	
Max Bandwidth:	-1.0	Send Timeout (ms):	300000	
Status:	0	Error Description:		
Start Time: 3/24/17	4:53 PM	Ba	ase at Zero Time Range; 5 Mins	×
o			RateReplicateAll	TotalTime= 0.0
			RateTotalBytesSe	ent= 4,700.5
° -			RateTotalEntries	Sent= 31.0
0-			RateTotalErrorRe	sponses= 0.0
o_			RateTotalMsgSer	nt= 30.9
o _[;			RateTotalMsgUna	cked= 0.0
o;			RateTotalRecord	sSent= 31.0
o:			ReplicateAllPero	entComplete= 0.0
16:49:00 16:49:30 10 03/24 03/24	8:50:00 16:50:30 16:51:00 16:51:30 16:5 03/24 03/24 03/24 03/24 03/	2:00 16:52:30 16:53:00 1 24 03/24 03/24	16:53:30 16:54:00 03/24 03/24	

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

Filter By:

Cluster:	Select a cluster from the drop-down menu.
Host:	Select a host from the drop-down menu.
Location:	Select a location from the drop-down menu. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .

ld:	The unique identifier for the node.
Participant Type	The participant type. Valid types are cluster and interceptor .
State	The participant state, where: 0 is Ok 1 is Warning 2 is Error
Bytes Sent Secs	The number of bytes sent per second.
Connect Retry Timeout (ms)	The configured connect retry timeout.
Current Bandwidth	The current amount of bandwidth being used, in megabits per second, for sending replicate message.
Estimated Replicate All Remaining Time	The estimated remaining time, in milliseconds, to complete the replicateAll request.
Geo I P	The Geo-IP metadata
Max Bandwidth	The maximum amount of bandwidth per second, in megabits, for sending replicate message, where -1.0 means the maximum bandwidth is not specified.
Status	The participant status.
Name	The sender name.
Msg Apply Time Percentile (ms)	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the destination.
Msgs Sent Secs	The number of messages sent per second.
Record Backlog Delay Time Percentile (ms)	The 90-percentile value, in milliseconds, of the time the journal records are in the cache waiting to be replicated.
Replicate All Percentile Complete	The percent of work completed for a replicateAll request.
Replicate All Total Time	The total amount of time the replicateAll request took, in milliseconds.
Send Timeout (ms)	The configured send timeout.
Error Description	A description of the error. A value exists only if the sender is in an error state.

Trend Graph

Select a location from the drop-down menu to populate the trend graph. **Location** is a unique identifier for each node and defined as: **member_name.machine.rack.site**.

RateReplicateAllTotalTime: Traces the total number of replicateAll requests per second.

RateTotalBytesSent: Traces the total number of bytes sent per second.

RateTotalEntriesSent: Traces the total number of entries sent per second.

RateTotalErrorResponses: Traces the total number of error responses per second.

RateTotalMsgSent: Traces the total number of messages sent per second.

RateTotalMsgUnacked: Traces the total number of unacknowledged messages per second.

RateTotalRecordsSent: Traces the total number of records sent per second.

ReplicateAllPercentComplete: Traces the percent of completed replicateAll requests.

- **Start Time** The date and time the location was started. **Location** is a unique identifier for each node and defined as: **member_name.machine.rack.site**.
- **Base at Zero** Use zero for the Y axis minimum for all graphs.

Select or Enter Date and Time:
Restore to Now
Ok Apply Cancel

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

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

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

Federated Origin Detail

Table shows performance and utilization data, such as bandwidth usage and bytes sent, for Federated Origins on the selected cluster. Use this display to do high level utilization analysis. Each row is a different Origin MBean. Click a row to see details in the "Federated Origin Summary" display. Sort data by the highest and lowest values of interest by clicking on the column heading.

Cluster2:Cluster2Jmx	Fede	rated Origin Deta	uil 24-1	Mar-2017 16:54 🗳 Da	ata OK + 😨
Host: All Hosts					
	Federate	d Origin Detail by I	Node		
Location A	BytesReceived Secs	Connection	DeltaTIME_STAMP	DeltaTotalBytesRe	ceived Delta
1.SLNB50	5,910	Cluster2Jmx	10,540		52,791
2.SLNB50	9,305	Cluster2Jmx	10,540		99,502
					•

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

Filter By:

Cluster: Select a cluster from the drop-down menu.

Host: Select a host from the drop-down menu.

Federated Origin Detail by Node

Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site .
BytesReceivedSecs	The number of bytes received per second.

Connection	The name of the JMX connection used to access the cluster data.
DeltaTIME_STAMP	The amount of time since the last data sample.
DeltaTotalBytesReceived	The difference in the total number of bytes received since the last data sample.
DeltaTotalEntriesReceived	The difference in the total number of entries received since the last data sample.
DeltaTotalMsgReceived	The difference in the total number of messages received since the last data sample.
DeltaTotalMsgUnacked	The difference in the total number of unacknowledged messages since the last data sample.
DeltaTotalRecordsReceived	The difference in the total number of records received since the last data sample.
Expired	When checked, this connection is expired due to inactivity.
HostName	The name of the host.
Member	The member information of the destination node.
MemberName	The name of the member.
MsgApplyTimePercentileMill is	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the origin.
MsgReceivedSecs	The number of messages received per second.
RateReplicateAllTotalTime	The number of replicateAll requests per second.
RateTotalBytesReceived	The total number of bytes received per second.
RateTotalEntriesReceived	The total number of entries received per second.
RateTotalMsgReceived	The total number of messages received per second.
RateTotalMsgUnacked	The total number of unacknowledged messages per second.
RateTotalRecordsReceived	The total number of records received per second.
RecordBacklogDelayTimePe rcentileMillis	The 90-percentile value, in milliseconds, of the time the journal records are in the cache waiting to be replicated.
TIME_STAMP	The date and time of the data update.
TotalBytesReceived	The total number of bytes received.
TotalEntriesReceived	The total number of cache entries received.
TotalErrorResponses	The total number of responses with an error.
TotalMsgReceived	The total number of replication messages received. A replication message might contain multiple journal records
TotalMsgUnacked	The total number of unacknowledged unacknowledged messages.
TotalRecordsReceived	The total number of journal records received. A journal record might consist of multiple cache entries that are part of the same transaction.
name	The destination cluster name.
nodeid	The unique identifier for the node.
service	The Federated Service name.

subType	The Federated Service sub-type.
type	The Coherence MBean type (Federation, in this case).

Federated Origin Summary

Detailed performance and utilization data, such as bandwidth usage and received per second, for a Federated Origin location. Use this display to do low level utilization analysis. Check the metrics for to determine whether more capacity is needed.



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

The display might include these filtering options:

Cluster:	Select a cluster from the drop-down menu.
Host:	Select a host from the drop-down menu.
Location:	Select a location from the drop-down menu. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
Bytes Received Secs	The number of bytes received per second.
Msg Apply Time Percentile (ms)	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the origin.
Msgs Received Secs	The number of messages received per second.
Record Backlog Delay Time Percentile (ms)	The 90-percentile value, in milliseconds, of the time the journal records are in the cache waiting to be replicated.
Total Bytes Received	The total number of bytes received.
Total Entries Received	The total number of cache entries received.
Total Msg Received	The total number of replication messages received. A replication message might contain multiple journal records.
Total Msg Unacked	The total number of unacknowledged replication messages.
Total Records Received	The total number of journal records received. A journal record might consist of multiple cache entries that are part of the same transaction.

Trend Graph

Select a location from the drop-down menu to populate the trend graph. Location is a unique identifier for each node and defined as: member_name.machine.rack.site.

RateReplicateAllTotalTime: Traces the total number of replicateAll requests per second. RateTotalBytesReceived: Traces the total number of bytes received per second. RateTotalEntriesReceived: Traces the total number of entries received per second. RateTotalErrorResponses: Traces the total number of error responses per second. RateTotalMsgReceived: Traces the total number of messages received per second. RateTotalMsgUnacked: Traces the total number of unacknowledged messages per second. RateTotalRecordsReceived: Traces the total number of records received per second. RateTotalRecordsReceived: Traces the total number of records received per second. RateTotalRecordsReceived: Traces the total number of records received per second.

Start Time The date and time the.....

Base at Zero

Use zero for the Y axis minimum for all graphs.

Time Range

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

Select or Enter Date and Time:
Restore to Now
Ok Apply Cancel

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

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

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

All Caches

All Caches displays present high-level cache performance metrics for the cluster. Use the All Caches displays to quickly assess total utilization metrics for all caches in the cluster.

- "All Caches Heatmap" on page 103: Heatmap of caches by service where size represents Units and color represents Total Gets%.
- "Storage Nodes Cache Map" on page 104: Heatmap of memory usage on storage nodes by service where size represents Units and color represents Units Used%.
- "Current Size Chart" on page 106: Bar chart/table sorted by caches with largest size displays current size/capacity metrics.
- "Current Activity Chart" on page 107: Bar chart/table sorted by caches with greatest activity displays current activity metrics.

All Caches Heatmap

Heatmap of cache size and activity organized by service: Size = Number of Units or Objects, Color = Percent of Total Gets.

← n/a:	DemoCluster	✓ Chart	All Caches - Ac	tivity Heatmap	05-Jul-20)16 12:02	< Data OK	+ 0
Service:	DistributedCache	-		%		0	50	100
	Heatmap c	Use Objects Use Objects of Caches organiz	instead of Units ed by Service/Cac Distribute	✓ Log Scale (Activity) he where Size = Units a dCache	Cache I nd Color = `	Labels Total Ge	Cumu ts	ılative
								_
) Settocow	ervice: Dis ache: nea nits: 119,6 ets: 87 it %: 87.8 its: 41 isses: 46 uts: 68	stributedCache r-foo2 607,920 %	
					4			

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

Cluster	Select a cluster to display.
Chart	Toggle between heatmap view and chart view.
Service	Select a service to display.
%	Set the activity percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
Use Objects Instead of Units	Select to use Objects instead of Units for heatmap cell sizing and mouseover tool-tips.

Log Scale (Activity)	Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value. The value of the Metric is displayed over the specified History for all process nodes in the selected Service.
Storage Nodes: History Heatmap of Service Metric	Color of the cells represents the relative value of the selected Metric for a given storage node; a darker shade is a larger value. The value of the Metric is displayed over the specified History for all storage nodes in the selected Service.
Cache Labels	Select to display cache labels.
Cumulative	Select to show cumulative statistics for each cache.
Heatmap of Caches organized by Service/Cache	Activity heatmap where the activity metric is TotalGets. The levels of this heatmap are Service>Cache. The size of the cells is based on Units. The size of aggregate cells is based on the sum of the Units used by its component cells. The color of the cache cells is based on TotalGets.

Storage Nodes Cache Map

Heatmap of memory usage on all storage nodes organized by service: Size = Number of Units, Color = Percent of Units Used.



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

Cluster	Select a cluster to display.	
Host to Show	Select a host to display.	
Memory Used%	Set the memory used percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.	
Service	Select a service to display, or select All Services. NOTE: When you select a specific service, only data for nodes running that service is displayed. This enables you to view services that only run on a subset of nodes.	
Ignore High Units	Select to remove High Units from calculations. This results in all caches having 100% units used. The color of cache cells represents units used instead of percent Units used when this option is selected.	
Log Scale	Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.	
Cell Borders	Select to display heatmap cell borders.	
Node Labels	Select to display node labels.	
Heatmap of Nodes organized by Service/Cache	A heatmap of memory usage. The levels of this heatmap are Host>Node>Service>Cache. The size of the cells is based on Units. The size of aggregate cells is based on the sum of the Units used by its component cells. The color of cache cells is based on the percent of Units used unless Ignore High Units is selected.	

Current Size Chart

Toggle between bar chart and table views that present the latest values of total objects and total nits for each cache in the selected service.



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

Cluster	Select a cluster to display.
Table	Toggle between chart view and table view.
Service	Select a service to display.
Front	Select for front tier, deselect for back tier.
Cache Count	Number of caches in the selected server. This is not available in the Table view.

Current Size Chart	 Total Objects in each Cache summed over all Storage Nodes. This is the default view. Toggle between totals for Object Count and Units Used. Click the Table to view Current Size Table. Objects shows the total number of objects in this cache (Object Count). Units shows the highest number of units before evictions occur. Ignore High Units removes High Units bars from view.
Current Size Table	Totals for each Cache over all Storage Nodes. Click Chart to view Current Size chart. shortCacheName Abbreviated name of cache tier Front or back Objects Total number of objects in this cache Units Total number of units (typically bytes) in this cache LowUnits Low limit for cache evictions HighUnits Highest number of units before evictions occur Service Name of selected service(s). Name Full name of cache

Current Activity Chart

Toggle between bar chart and table views that present the latest values for activity metrics for each cache in the selected service.



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

Cluster	Select a cluster to display.
Table	Toggle between chart view and table view.
Service	Select a service to display.
Front	Select for front tier, deselect for back tier.
Cache Count	Number of caches in the selected server. This is not available in the Table view.
Cumulative	Select to show cumulative statistics for each node since the start of the node.
Current Activity Chart	Totals for Cache summed over all Storage Nodes. This is the default view. Toggle to Table view. Sort by: Objects shows the total number of objects in this cache (Object Count). Units shows the highest number of units before evictions occur. Ignore High Units removes High Units bars from view.
Current Activity Table	Totals for each Cache over all Storage Nodes. Toggle to Chart view. Sort by: Cache Abbreviated name of cache tier Front or back Hits Total number of successful gets Misses Total number of failed gets Gets Total requests for data from this cache Puts Total data stores into this cache Hit% Ratio of hits to gets Service Service Name Cache Full Name Full name of cache

Single Cache

Single Cache displays present detailed cache performance metrics for a single cache. Use the Single Cache displays to perform cache utilization analysis. The data in these displays can be sorted and viewed by service or cache.

- "Single Cache Summary" on page 109: Perform low level utilization analysis on a single cache.
- "Size Trends" on page 112: Trend chart displays size/capacity metrics.
- "Activity Trends" on page 115: Trend chart displays activity metrics.

- "Cache Detail Tables" on page 117: Table showing current detailed cache statistics by node.
- "Storage Manager Detail" on page 119: Table showing store manager metrics.
- "Node/Group Distribution" on page 121: Bar chart displays metrics showing distribution across cluster nodes or groups.
- "Front/Back Analysis" on page 123: Displays metrics for the front and back tiers of a selected cache.

Single Cache Summary

Use Single Cache - Summary display to do low level cache utilization analysis. Check the metrics for Size, Evictions and Misses to determine whether more capacity is needed. Cache Summary provides summary information about an individual cache.



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

Cluster	Select a cluster to display.
Service	Select a service to display.
Cache	Select a cache. Click the Detail button to get information specific to the selected cache.
Front	Select for front tier, deselect for back tier.
Туре	The type identifier string from the ServiceMBean (ReplicatedCache, DistributedCache, etc.).
Client Nodes	The number of cluster nodes that do not have storage enabled.
Storage Nodes	Select to display storage node data in the trend graphs of this display.
Туре	The type of cache.
Storage Nodes	The number of storage nodes in the cache.
Status	The high availability status of the service:
	• ENDANGERED: There is potential data loss in the cluster if a node goes offline.
	NODE-SAFE: There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.
	• MACHINE-SAFE: There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.
	• RACK-SAFE: There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.
	SITE-SAFE: There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.
Size	Units indicates memory usage for the back tier and number of objects for the front tier.
	Objects The number of objects in the selected cache. The value is the total across all storage nodes.
	Avg Size The average size of objects in the selected cache (in bytes if it is the back tier).
	Units The memory usage if back tier, or number of objects if front tier. The value is the total across all storage nodes.
	High Units Maximum memory, or number of objects allowed before Coherence starts to evict objects from the selected cache. The value is the total across all storage nodes.
	Low Units The level of memory, or number of objects to which Coherence will reduce the cache during the eviction process. The value is the total across all storage nodes.

Persistence	Type The persistence type for the cache. Possible values include: NONE , READ-ONLY , WRITE-THROUGH , and WRITE-BEHIND .
	Failures The number of write (cache store) failures, including load, store and erase operations. NOTE: This value is -1 if the persistence type is NONE.
	Wr Queue The size of the queue, in kilobytes, that holds data scheduled to be written to the cache store.
	Writes The number of objects (cache entries) written to the cache store.
Activity	Current: Use the Use Rates checkbox to toggle between two value types: Activity - Current (Rate) and Activity - Current (Delta) (as labeled in the display upon selection). When the Use Rates (checkbox) is NOT selected the Delta values are shown in the Activity - Current (Delta) fields and trend graphs. Delta is the difference in the value since the last sample. When the Use Rates (checkbox) is selected the Rate values are shown in the Activity - Current (Rate) fields and trend graphs. Rate is the value per second. The Rate value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the Rate value does not vary if the sample period changes (whereas the Delta value does vary). The Rate value enables you to directly compare rates on systems with different sample periods.
	Cumulative: The total since the service was started for the selected cache, or since statistics were reset.
	Gets The number of requests for data from this cache.
	Hits The number of successful gets.
	Misses The number of failed gets.
	Puts The number of data stores into this cache.
	Evictions The number of objects removed to make room for other objects.
Use Rates	Select Use Rates to show the Rate values in the Activity - Current (Rate) fields and trend graphs. Rate is the value per second. The Rate value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the Rate value does not vary if the sample period changes (whereas the Delta value does vary). The Rate value enables you to directly compare rates on systems with different sample periods.
	Deselect Use Rates to show the Delta values in the Activity - Current (Delta) fields and trend graphs. Delta is the difference in the value since the last sample.
Log Scale	Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Use zero for the Y axis minimum for all graphs.

Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last
3	7 Days , or display All Data . To specify a time range, click Calendar

	1 5 0	
Γ	Select or Enter Date and Time:	
	Destars to New	
	Restore to Now	
	Ok Apply Cancel	

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

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

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

Trend Graphs Use the Use Rates checkbox to toggle between two value types: Activity - Current (Rate) and Activity - Current (Delta) (as labeled in the display upon selection).

Objects The number of objects in the selected cache. The value is the total across all storage nodes.

TotalGets Total requests for data from this cache.

CacheMisses Total number of failed gets.

TotalPuts Total data stores into this cache.

EvictionCounts Number of objects removed from the cache to make room for other objects.

Size Trends

Size Trends provides a method of viewing the degree to which available cache size has been consumed. Under normal operations the cache will evict and reload objects into the cache. This will be displayed as a significant drop in the Units Used trend. However, if these drops are too frequent the application might not be performing optimally. Adding capacity and examining or modifying application usage patterns might be required. The data displayed here is a sum of all storage nodes in the cache filtered by the selected service and cache.

Try changing the High Units setting in the Cache Administration page to something like 100,000 and then see the effect on these trend charts.





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

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

Open the Alert Views - RTView Alerts Table display.

Cluster	Select a cluster to display.
Service	Select a service to display.
Cache	Select a cache. Click the Detail button to get information specific to the selected cache.
Front	Select for front tier, deselect for back tier.

Base at Zero Use zero for the Y axis minimum for all graphs.

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

Sele	ct or Enter Date and Time:
	Restore to Now
Ok	Apply Cancel

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

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

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

Activity Trends

Activity Trends provides a set of trend graphs that show the magnitude of the cache usage and the effectiveness of the implementation. If the overall effectiveness is not as desired, increasing capacity, preloading the cache and increasing the eviction time may result in improvements in cache hits. The data displayed here is a sum of all storage nodes in the cache filtered by the selected service and cache.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

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

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

Open the Alert Views - RTView Alerts Table display.

Cluster	Select a cluster to display.
Service	Select a service to display.
Cache	Select a cache. Click the Detail button to get information specific to the selected cache.
Front	Select for front tier, deselect for back tier.

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

Base at Zero Use zero for the Y axis minimum for all graphs.

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



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

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

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

Hits The number of successful gets from this cache. **Activity Metrics** for Selected

Total Gets Requests for data from this cache.

Cache Misses The number of failed gets by this cache.

Total Puts The number of data stores into this cache.

Evictions The number of objects removed from the cache to make room for other objects.

Write Queue The size of the queue, in kilobytes, that holds data scheduled to be written to the cache store.

Cache

Cache Detail Tables

This display presents detailed information about the contribution that each storage node makes to the cache. Select a node in the Statistics By Node for Selected Cache table to drill down to the "Node Summary" display for that node. The data displayed here is broken down for each storage nodes in the cache filtered by the selected service and cache.

🗲 n/a:De	emoCluster		Summa	ary Si	ngle (Cache -	Det	tail Tal	bles			05-Jul-2	016	12:15	¢	Data OK	+ 📀
Service: D	istributedCache	;	•	Cache:	foo					•	F	ront					
Objects	Units	LowUnits	HighU	nits	Hits	s De	lta	Miss	ses l	Delta	G	jets	De	elta	Pu	ıts	Delta
389,542	105,955,424	96,000,000	144,00	0.000	4,366	.039 2	282	1,38	3,079	193	5,7	49,118		475	3,50	0.956	310
											_						
AvgSize	Units	LowUnit	s	HighUr	nits	Hits	I	Delta	Misse	s De	elta	Gets		Delta	F	Puts	Delta
272	8,829,61	L8 8,000,	000	12,00	0,000	363,8	36	23	115,2	56	16	479,0	93	39	2	91,746	25
							_										
	Location		tier	Obje	ects	AvgSize		Un	its	Lo	wUn	its	Hig	ghUnit	s	Hit	s
StoreNode	1.SLHOST1		back	34	,541	27	2	9,:	395,152		8,00	0,000	1	.2,000,	000	1	L,669
StoreNode0	1.SLHOST2		back	37	,492	27	2	10,	197,824		8,000	0,000	- 1	<mark>2,000</mark> ,	000	183	3,778
StoreNode	01.SLHOST4		back	21	.,943	27	72	5,9	968,496		8,00	0,000	1	. 2,000,	000	734	4,269
StoreNode	04.SLHOST1		back	33	,379	27	72	9,0	079,088		8,00	0,000	1	2,000,	000	1	1,628
StoreNode0	04.SLHOST4		back	31	.,346	27	72	8,	526,112		8,00	0,000	1	.2,000,	000	658	3,192
StoreNode0	5.SLHOST2		back	32	,555	27	2	8,8	854,960		8,00	0,000	1	2,000,	000	202	2,794
StoreNode	5.SLHOST4		back	32	,889	27	2	8,9	945,808		8,00	0,000	1	.2,000,	000	770),567
StoreNode	5n.SLHOST1		back	39	,918	27	2	10,8	857,696		8,000	0,000	1	2,000,	000	1	L,650
StoreNode	5n.SLHOST2		back	33	,736	27	2	9,	176,192		8,00	0,000	1	2,000,	000	200),245
StoreNode	5n.SLHOST4		back	31	.,224	27	2	8,4	492,928		8,000	0,000	1	2,000,	000	764	1,372
StoreNode	8.SLHOST2		back	35	,390	27	2	9,0	626,080		8,00	0,000	1	2,000,	000	151	L,865
StoreNode	8.SLHOST4		back	25	,129	27	2	6,8	835,088		8,000	0,000	1	2,000,	000	695	5,010
•	111																

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

Cluster	Select a cluster to display.
Summary	Toggle between this display and Single Cache - Summary display.
Service	Select a service to display.
Cache	Select a cache. Click the Detail button to get information specific to the selected cache.
Front	Select for front tier, deselect for back tier.

Totals for Selected Cache	 Objects Number of objects in this cache. Units Total number of units (typically bytes) in this cache. LowUnits Low limit for cache evictions. HighUnits Highest number of units before evictions occur. Hits Total number of successful gets. Misses Total number of failed gets. Gets Total requests for data from this cache. Puts Total data stores into this cache.
Average for Selected Cache	 Objects Number of objects in this cache. AvgSize Average size of objects in this cache. Units Average number of units (typically bytes) in this cache. LowUnits Low limit for cache evictions. HighUnits Highest number of units before evictions occur. Hits Average number of successful gets. Misses Average number of failed gets. Gets Average requests for data from this cache. Puts Average data stores into this cache.
Statistics By Node for Selected Cache	The columns in this table, with the exception of Location , come from Cache and Node MBeans. Location is a unique identifier for each node and defined as: member_name.machine.rack.site . For details about attributes of these MBeans go to: http:// download.oracle.com/otn_hosted_doc/coherence/350/com/tangosol/ net/management/Registry.html.

Storage Manager Detail

This display presents detailed information about the Storage Manager. The data displayed here is queried from the Coherence StorageManagerMBean, filtered by the selected service and cache. Click on a row in the table to open the "Storage IndexInfo View" window.

← n/a:DemoCluster 🔽	Stora	age Manager D	etail	05-Jul-2016 12:16	🤹 💠 Data OK 🛛 💠
Service: DistributedCache	Cache: foo		•		
				(Select a No	ode to see IndexInfo
Location	EventsDispatched	EvictionCount	InsertCount	ListenerFilterCount	ListenerKeyCount
StoreNode01.SLHOST1	0	0	34,620	0	0
StoreNode01.SLHOST2	0	38,015	75,579	0	0
StoreNode01.SLHOST4	0	146,554	168,601	0	0
StoreNode04.SLHOST1	0	0	33,472	0	0
StoreNode04.SLHOST4	0	120,540	151,971	0	0
StoreNode05.SLHOST2	0	60,349	92,983	0	0
StoreNode05.SLHOST4	0	115,337	148,315	0	0
StoreNode05n.SLHOST1	0	0	40,006	0	0
StoreNode05n.SLHOST2	0	43,208	77,031	0	0
StoreNode05n.SLHOST4	0	130,074	161,389	0	0
StoreNode08.SLHOST2	0	44,379	79,849	0	0
StoreNode08.SLHOST4	0	121,196	146,401	0	0
•	III				4

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

Cluster Select a cluster to display.

Service Select a service to display.

Cache	Select a cache. Click the Detail button to get information specific to the selected cache.
Storage Manager Data	Location A unique identifier for each node. It is defined as member_name.machine.rack.site.
	EventsDispatched The total number of events dispatched by the Storage Manager since the last time the statistics were reset.
	EvictionCount The number of evictions from the backing map managed by this Storage Manager caused by entries expiry or insert operations that would make the underlying backing map to reach its configured size limit.
	InsertCount The number of inserts into the backing map managed by this Storage Manager. In addition to standard inserts caused by put and invoke operations or synthetic inserts caused by get operations with read-through backing map topology, this counter is incremented when distribution transfers move resources `into` the underlying backing map and is decremented when distribution transfers move data `out`.
	ListenerFilterCount The number of filter-based listeners currently registered with the Storage Manager.
	ListenerKeyCount The number of key-based listeners currently registered with the Storage Manager.
	ListenerRegistrations The total number of listener registration requests processed by the Storage Manager since the last time the statistics were reset.
	LocksGranted The number of locks currently granted for the portion of the partitioned cache managed by the Storage Manager.
	LocksPending The number of pending lock requests for the portion of the partitioned cache managed by the Storage Manager.
	RemoveCount The number of removes from the backing map managed by this Storage Manager caused by operations such as clear, remove or invoke.

Storage IndexInfo View

Click on a row in the Storage Manager Data table to open the Storage IndexInfo View window.

	Sto	orage IndexInfo View	Close
Cluster:	DemoCluster		
Service:	DistributedCach	e	
Cache:	foo		
Location:	StoreNode05.SL	HOST2	
Cardin	ality: 0	Total Indexes: 0	
•			۰.

Service	The name of the service.
Cache	The name of the cache.
Location Manager Data	The location of the node associated with the cache. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
(Index Table)	Each row in the table represents a unique index, where: Extractor = the index name. Ordered = true/false to indicate whether or not the data is sorted (false means the data is not sorted). Size = the number of entries in that cache whose value matches that extractor.

Node/Group Distribution

This display presents the distribution of cache activity across all storage nodes in the cluster. The buttons on the left may be used to select the metric by which all six bar charts are to be sorted. Note that the Gets, Hits, Misses, and Puts are shown in the same color as on the other Cache Analysis displays. The data displayed here is broken down for each storage nodes in the cache filtered by the selected service and cache.



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

Cluster	Select a cluster to display.
Service	Select a service to display.
Cache	Select a cache. Click the Detail button to get information specific to the selected cache.
Group By	Select the node group by which the data are totaled. Location A unique identifier for each node, defined as member_name.machine.rack.site. This is the default setting. Gets Requests for data from this cache. Hits Number of successful gets. Misses Number of failed gets. Puts Data stores into this cache. Mem% Calculated percent of memory used divided by total memory. K Units Units in thousand bytes.

Front/Back Analysis

← n/a:Der	noCluste	r	•	F	ront Bac	k Cache A	nalysis	05	-Jul-2016 12:27	🔹 Data OK	+ 📀
Service: Dis	tributedC	ache		Cache:	foo		•				
	Units: 0)	Cache Size: 0		High Units:	0	Low Units: 0			F	RONT
Cache Hits											
Cache Misses											
Total Puts											
Service											
Network											
	Misses	s: 0	Hits:	0	Total Gets:	0	Total Puts:	D			
	Total Get	s: 386	Hits:	174	Misses:	212	Total Puts: 3	07		CLUCCTA	BACK
Network									Node	: StoreNode(8.SLHOST4
Comico				_	_				Cach	e: foo	_
Service											
Cache Hits			-				_				_
October Minner											
Cache Misses											
Total Puts											
	Units: 1	08,951,776	Cache Size: 40	0,558	High Units:	144,000,000	Low Units: 96	,000,000			

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

- Cluster Select a cluster to display.
- Service Select a service to display.

Cache Select a cache. Click the Detail button to get information specific to the selected cache.

FRONT/BACK	Units: Front Number of objects. The value is the total across all storage nodes for the given tier.
	Back Memory usage. The value is the total across all storage nodes for the given tier.
	Cache Size: Total number of objects in the cache for the given tier (Front or Back). NOTE: Same value as Units for Front tier.
	High Units: Front Number of objects allowed before Coherence starts to evict objects from the selected cache. The value is the total across all storage nodes for the given tier.
	Back Maximum memory allowed before Coherence starts to evict objects from the selected cache. The value is the total across all storage nodes for the given tier.
	Low Units: Front Number of objects to which Coherence will reduce the cache during the eviction process. The value is the total across all storage nodes for the given tier.
	Back The level of memory to which Coherence will reduce the cache during the eviction process. The value is the total across all storage nodes for the given tier.
Cache Hits	Number of successful gets
Cache Misses	Number of failed gets
Total Puts	Data stores into this cache
Service	CPU usage (%) for the node.
Network	Front Sent Packet Failure Rate (%) for the node. Back Received Packet Failure Rate (%) for the node.
Misses	Number of failed gets.
Hits	Number of successful gets.
Total Gets	Total requests for data from this cache.
Total Puts	Total data stores into this cache.

All Nodes

All Nodes displays present high-level node performance metrics for the cluster. Use the All Nodes displays to quickly assess total utilization metrics for all nodes in the cluster.

- "All Nodes by Type/Host/Memory" on page 125: Heatmap of caches by service where size represents Max Memory and color represents percent of Memory Used.
- "All Nodes CPU" on page 126: Heatmap shows CPU utilization for all nodes in the cluster.
- "All Nodes Grid View" on page 127: Grid view showing information about all nodes.
- "Communication Issues" on page 128: Bar chart displays current communication issues for all nodes.
- "All Nodes Detail" on page 130: Table shows current detailed statistics for all nodes.
- "Invocation Service Detail" on page 132: Table shows invocation service detail for all nodes.

All Nodes by Type/Host/Memory

Heatmap of nodes organized by Type and Host: Size = Max Memory, Color = Percent of Memory Used.



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

Cluster	Select a cluster to display.
Nodes Type	Select the type of node to display: Storage Nodes, Process Nodes or All Types.
Memory Used%	Set the memory used percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
Heatmap of Nodes organized by Type/Host	A heatmap of memory usage per host.

All Nodes CPU

Heatmap shows CPU utilization for all nodes in the cluster organized by Type and Host: Size = Max Memory, Color = Percent of CPU Used.



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

Cluster	Select a cluster to display.
Node Type	Select the type of node to display: Storage Nodes, Process Nodes or All Types.
CPU Used%	Set the CPU used percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
Heatmap of Nodes organized by Type/Host	A heatmap of CPU usage per host.

All Nodes Grid View

This display shows a grid view of all of the nodes in the selected Node Type.

fn/a:DemoCluster	All Nodes - Object	Grid View	05-Jul-2016 12:33	< Data OK	+	
Node Type: All Types 🔽	Host: All Hosts	64				
	All Node Ty	pes				
0.0 %	ExtendNode01.SLHOST2 75 SLHOST2 254 0 2.0 %	ExtendNode01.SLHOS 56 SLHOST3 247 0 2.8 %	Τ3			A
ExtendNode01.SLHOST4 83 SLHOST4 254 0 1.6 %	ExtendNode02.SLHOST2 76 SLHOST2 254 0 1.6 %	ExtendNode02.SLHOS 86 SLHOST4 254 0 2.0 %	T4			
JmxNode01.SLHOST1 24 SLHOST1 247 5 17.8 %	OLS-ProcessNode-1.SLHOST2 78 SLHOST2 254 0 4.3 %	OLS-ProcessNode-1.S 88 SLHOST4 254 18 3.5 %	SLHOST4			
OLS-StoreNode-1.SLHOST2 2 77 SLHOST2 2 254 8 22,821 15.4 % 9,859 k	OLS-StoreNode-1.SLHOST4 87 SLHOST4 2 254 0 12,051 14.2 5,202 K	ProcessNode01.SLHO 42 SLHOST2 254 0 1.2 %	IST2			
ProcessNode01.SLHOST3 39 SLHOST3 247 0 2.0 %	ProcessNode01.SLH0ST4 37 SLH0ST4 254 0 1.6 %	ProcessNode02.SLHO 47 SLHOST2 254 0 1.2 %	IST2			
ProcessNode02.SLH0ST3 40 sLH0ST3 247 0 2.4 %	ProcessNode02.SLH0ST4 38 SLH0ST4 254 0 2.0 %	ProcessNode03.SLHO 48 SLHOST2 254 0 1.6 %	IST2			+

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

Cluster	Select a cluster to display.

Node Type Select the type of node to display: Storage Nodes, Process Nodes or All Types.

Host Select a host to display.

Heatmap of A heatmap of CPU usage per host.

Nodes organized by Type/Host The following icon is shown for each node in the cluster:



The icon describes the node:

- Location (StoreNodeO4.VMXP-7) A unique identifier for each node. It is defined as: member_name.machine.rack.site.
- Id (4)
- Host name or IP (vmpx-7)
- Max megabytes (247)
- Messages queued (0)
- Meter and label indicating the percent of memory utilization(12.1%)

For storage nodes, the following are also shown (in the lower right portion of the icon):

- Number of supported caches (13),
- Number of objects (32,944)
- Amount of memory used (8,935 K).

Communication Issues

This display presents detail information about communication issues by node or group. Both bar charts show the same data as the Packet Detail table. Click on a bar in either chart to drill down to the "Node Summary" display for that node.



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

Cluster	Select a cluster to display.
Node Count	Number of nodes in the cluster.
Packets Repeated Recently	Total number of repeated packets since the last update. The update rate is set by the Reporting Period.
Resent Recently	Total number or resent packets since the last update. The update rate is set by the Reporting Period.
Reporting Period	Select period varying from 30 Seconds to Last 7 Days, or display All Data.
Sort By	Select Packets Sent, Packets Received, Sent Failure Rate or Received Failure Rate.
Sent Failure	Packets failed to be sent by each node.
Rate/Received Failure Rate by Node/Group	Packets failed to be received by each node.
Packets Sent/ Received by Node/Group	Packets sent by each node. Packets received by each node.

All Nodes - Detail

This display presents detailed information about each node. This display includes information from the Coherence ClusterNodeMBean for both storage and processing nodes. Select a node in the All Node Data table to drill down to the "Node Summary" display for that node.

n/a:DemoCluster	•	All Nodes - Detail				5-Jul-2016 12:37	< Data O	к 🕂 🕜
Node Type: All Types 🔹	Host: All I	: All Hosts 🔽 64			Total	Avail MB: 14,1	08 Max:	16,103
			All Node T	ypes	(ME	Bean Detail Da	ta)	
Location	ld	Avail MB	Max MB	Pkts Sent	Delta	Pkts Rovd	Delta	Pkts I 🔺
StoreNode08.SLHOST4	27	56	254	20,680,021	1,881	19,769,956	1,759	1:
StoreNode08.SLHOST2	8	112	254	7,207,444	1,755	6,920,285	1,629	4
StoreNode07.SLHOST4	16	248	254	8,479,585	814	7,596,866	703	
StoreNode07.SLHOST2	7	249	254	2,954,996	811	2,652,154	693	=
StoreNode06.SLHOST4	15	245	254	14,413,751	1,289	13,550,004	1,187	
StoreNode06.SLHOST2	6	246	254	4,998,465	1,519	4,697,421	1,386	
StoreNode05n.SLHOST4	13	144	254	9,240,557	832	8,292,349	717	
StoreNode05n.SLHOST2	4	106	254	3,215,830	818	2,917,402	698	
StoreNode05n.SLHOST1	17	145	247	145,793	890	146,745	749	
StoreNode05.SLHOST4	14	112	254	9,045,804	734	8,096,728	619	(
StoreNode05.SLHOST2	5	149	254	3,130,043	799	2,826,880	694	
StoreNode04.SLHOST4	12	163	254	20,767,959	1,825	19,859,593	1,720	()
StoreNode04.SLHOST1	19	141	247	351,071	1,934	350,140	1,798	
StoreNode03.SLHOST4	11	249	254	8,452,999	915	7,564,521	800	
StoreNode03.SLHOST2	3	249	254	2,911,819	901	2,606,105	784	
StoreNode03.SLHOST1	22	236	247	151,367	890	132,011	751	
StoreNode02.SLHOST4	10	247	254	14,330,435	1,265	13,452,552	1,148	
StoreNode02.SLHOST2	2	245	254	4,873,858	1,357	4,569,689	1,226	
StoreNode02.SLHOST1	23	234	247	243,086	1,248	223,462	1,109	
StoreNode01.SLHOST4	9	97	254	9,003,123	710	8,202,539	600	
StoreNode01.SLHOST2	1	139	254	3,074,838	682	2,822,646	572	
StoreNode01.SLHOST1	18	169	247	145,132	848	147,525	718	-
< III								•

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

Cluster	Select a cluster to display.
Node Type	Select the type of nodes for which to display data: Storage Nodes, Process Nodes or All Types.
Host	Select the host for which to display data, or select All Hosts.
Node Count	Number of nodes for which data is currently displayed.
Total Avail MB	Total available memory of all nodes in the cluster.
Max	Total max memory of all nodes in the cluster.

All Node Types (MBean Detail Data) Location A unique identifier for each node. It is defined as: member_name.machine.rack.site.

• Id The short member id that uniquely identifies this member.

- Avail MB The amount of available memory for this node in MB.
- Max MB The maximum amount of memory for this node in MB.
- **Pkts Sent** The cumulative number of packets sent by this node since the node statistics were last reset.
- Delta The number of packets sent by this node since the last update.
- **Pkts Rcvd** The cumulative number of packets received by this node since the node statistics were last reset.
- Delta The number of packets received by this node since the last update.
- **Pkts Rptd** The cumulative number of duplicate packets received by this node since the node statistics were last reset.
- Delta The number of duplicate packets received by this node since the last update.
- **Pkts Resent** The cumulative number of packets resent by this node since the node statistics were last reset.
- Delta The number of packets resent by this node since the last update.
- Timestamp The date and time (in cluster time) that this member joined the cluster.
- **Pub Succ Rate** The publisher success rate for this node since the node statistics were last reset. Publisher success rate is a ratio of the number of packets successfully delivered in a first attempt to the total number of sent packets. A failure count is incremented when there is no ACK received within a timeout period. It could be caused by either very high network latency or a high packet drop rate.
- **Rec Succ Rate** The receiver success rate for this node since the node statistics were last reset. Receiver success rate is a ratio of the number of packets successfully acknowledged in a first attempt to the total number of received packets. A failure count is incremented when a re-delivery of previously received packet is detected. It could be caused by either very high inbound network latency or lost ACK packets.
- Member The member name for this node.
- Machine The machine name for this node.
- Rack The rack name for this node.
- Site The site name for this node.
- Process The process name for this node.
- **Uni Addr** The unicast address. This is the IP address of the node's DatagramSocket for point-topoint communication.
- **Uni Port** The unicast port. This is the port of the node's DatagramSocket for point-to-point communication.
- RoleName The role name for this node.
- **ProductEdition** The product edition this node is running. Possible values are: Standard Edition (SE), Enterprise Edition (EE), Grid Edition (GE).
- Send Queue The number of packets currently scheduled for delivery, including packets sent and still awaiting acknowledgment. Packets that do not receive an acknowledgment within the ResendDelay interval are automatically resent.

Packet Transmission Totals

- **Pkts Sent** Total cumulative packets sent by all nodes in the cluster since the node statistics were last reset.
- Delta Total packets sent by all nodes in the cluster since the last update.
- **Pkts Rcvd** Total cumulative packets received by all nodes in the cluster since the node statistics were last reset.
- **Delta** Total packets received by all nodes in the cluster since the last update.
- **Pkts Rptd** Total cumulative packets repeated by all nodes in the cluster since the node statistics were last reset.
- **Delta** Total packets repeated by all nodes in the cluster since the last update.

• **Pkts Resent** - Total cumulative packets resent by all nodes in the cluster since the node statistics were last reset.

Delta - Total packets resent by all nodes in the cluster since the last update.

Invocation Service Detail

This display presents detailed information about invocation services. The data displayed here is queried from the Coherence ServiceMBean filtered to only display services of type Invocation. Click on a node in the table to drill down to the "Node Summary" display for that node.

← n/a:DemoCluster 🔽	Invocation Service Detail				05-Jul-	2016 12:38 🗳	🔊 Data OK	+	0
Host: All Hosts	•								
	 Invoca	tion Service	e Detail b	v Node					
Location	name	Running	CPU %	Messages	Delta	Requests	Delta	Rec	*
ExtendNode01.SLHOST2	Management		0.1	183,053	26	2	0		
ExtendNode01.SLHOST3	Management		0	636,820	30		0		
ExtendNode01.SLHOST4	Management	×	0.3	558,849	48	2	0		
ExtendNode02.SLHOST2	Management	L L	0.5	193,035	50	2	0		
ExtendNode02.SLHOST4	Management	2	0.8	559,564	50	2	0		_
JmxNode01.SLHOST1	Management	V	11.5	841,145	4,790	839,383	4,790		=
OLS-ProcessNode-1.SLHOST2	Management	×	0.2	207,174	54	2	0		
OLS-ProcessNode-1.SLHOST4	Management	L L	0.2	599,986	52	2	0		
OLS-StoreNode-1.SLHOST2	Management	2	0.1	263,598	69	2	0		
OLS-StoreNode-1.SLHOST4	Management	L I	0.2	763,897	69	2	0		
ProcessNode01.SLHOST2	Management	M	0.6	207,681	52	2	0		
ProcessNode01.SLHOST3	Management	L I	0.2	691,792	62	2	0		
ProcessNode01.SLHOST4	Management	M	0.1	600,603	53	2	0		
ProcessNode02.SLHOST2	Management	L I	0.2	207,370	53	2	0		
ProcessNode02.SLHOST3	Management		0	691,971	59	2	0		
ProcessNode02.SLHOST4	Management	L I	0.2	600,699	53	2	0		
ProcessNode03.SLHOST2	Management		0.3	221,989	57	2	0		
ProcessNode03.SLHOST3	Management	V	0.3	733,684	66	2	0		
ProcessNode03.SLHOST4	Management	M	0.5	642,766	59	2	0		
ProcessNode04.SLHOST2	Management	V	0.3	235,858	61	2	0		
ProcessNode04.SLHOST3	Management	M	0.5	775,442	67	2	0		
ProcessNode04.SLHOST4	Management	L I	0.2	683,213	62	2	0		
ProcessNode05.SLHOST3	Management	M	0.2	689,541	59	2	0		
ProcessNode05.SLHOST4	Management	L I	0.3	599,289	52	2	0		
ProcessNode05n.SLHOST2	Management		0.5	235,663	61	2	0		
ProcessNode05n.SLHOST3	Management	L L	0.2	774,342	67	2	0		Ŧ
•								•	

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

Cluster

Select a cluster to display.

Host

Select the host for which to display data, or select All Hosts.

Invocation Service Information

• Location A unique identifier for each node. It is defined as: member_name.machine.rack.site.

- name The name of the invocation service.
- Running Indicates that the invocation service is running when checked.
- CPU% The percent (%) of CPU used by the node.
- Messages The number of messages issued by the service to the node in a given time period.
- **Delta** The number of messages received by the node since the last update.
- **Requests** The number of requests issued by the service to the node in a given time period.
- **Delta** The number of requests received by the node since the last update.
- **RequestAverageDuration** The average duration (in milliseconds) of an individual synchronous request issued by the service since the last time the statistics were reset.
- **RequestMaxDuration** The maximum duration (in milliseconds) of a synchronous request issued by the service since the last time the statistics were reset.
- RequestPendingCount The number of pending synchronous requests issued by the service.
- **RequestPendingDuration** The duration (in milliseconds) of the oldest pending synchronous request issued by the service.
- **RequestTimeoutCount** The total number of timed-out requests since the last time the statistics were reset.
- **RequestTimeoutMillis** The default timeout value in milliseconds for requests that can be timed-out (e.g. implement the com.tangosol.net.PriorityTask interface), but do not explicitly specify the request timeout value.
- TaskAverageDuration The average duration (in milliseconds) of an individual task execution.
- **TaskBacklog** The size of the backlog queue that holds tasks scheduled to be executed by one of the service pool threads.
- TaskCount The total number of executed tasks since the last time the statistics were reset.
- TaskHungCount The total number of currently executing hung tasks.
- TaskHungDuration The longest currently executing hung task duration in milliseconds.
- TaskHungTaskId The id of the of the longest currently executing hung task.
- **TaskHungThresholdMillis** The amount of time in milliseconds that a task can execute before it is considered hung. Note that a posted task that has not yet started is never considered as hung.
- TaskMaxBacklog The maximum size of the backlog queue since the last time the statistics were reset.
- TaskTimeoutCount The total number of timed-out tasks since the last time the statistics were reset.
- **TaskTimeoutMillis** The default timeout value in milliseconds for tasks that can be timed-out (e.g. implement the com.tangosol.net.PriorityTask interface), but do not explicitly specify the task execution timeout value.
- **ThreadAbandonedCount** The number of abandoned threads from the service thread pool. A thread is abandoned and replaced with a new thread if it executes a task for a period of time longer than execution timeout and all attempts to interrupt it fail.
- **ThreadAverageActiveCount** The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.
- ThreadCount The number of threads in the service thread pool.
- ThreadIdleCount The number of currently idle threads in the service thread pool.
- · HostName Name of the host machine on which the service resides.
- **Throughput** The amount of data (in kilobytes) that is transferred by the service to the node.

Single Node

Single Node displays present detailed node performance metrics for a single node. Use the Single Node displays to perform node utilization analysis.

- "Node Summary": Summary view showing details about a single node.
- "Service Trends": Trend graphs showing metrics on a selected node of a selected service. Allows you to visually compare the behavior of metrics over time, for a given node.
- "Node Detail": Tables showing metrics for Node, Cache, Invocation Service, Cache Service, and Storage Manager MBeans.
- "JVM Summary": Runtime, class loader, thread, OS and input arguments.
- "JVM Memory Trends": Heap and non-heap memory trends.
- "JVM GC Trends": Memory usage before and after garbage collection and Garbage Collector activity.
- "System Properties": Table of Java properties for a selected node.

Node Summary

This display presents summary information about an individual node.

st: All Hosts	s -	Location: Sto	preNode01.SLHOST2		-	ld: 1	
Member:	StoreNode01.SLH	IOST2					
Machine:	SLHOST2		Cache Na	me Tier	Objects	Units	
Rack:	n/a		near-foo2	back	35,455	10,731,760	
			foo2	back	35,529	10,232,352	i i
Site:	n/a		near-foo	back	35,444	10,207,872	
Role:	MyDemoDataMyD)emoData	Packet Commun	ication			
Process:	2964		Sent: 592	Resent %: 0	Send Ou	ieue: 0	
ni Address:	SLHOST2/192.168	.200.102	Rcvd: 552	Rptd %: 0			
nicast Port	8088						
			Memory				
CPU Count:	2		Memory Max MB: 254	Avail MB: 13	2 Used	%: 48.03	
CPU Count: Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13	2 Used	%: 48.03	
CPU Count: Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13 Base at Zero	2 Used	%: 48.03 5 Mins	•
CPU Count: Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13 Base at Zero	2 Used	%: 48.03	- [
CPU Count: Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13	2 Used	%: 48.03	ail Ri
Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13	2 Used Time Range:	%: 48.03	ail Ri ail R Jsed
CPU Count: Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13	2 Used	%: 48.03 5 Mins 5 Sent Fa Rcvd F Rcvd F Mem U CPU % CPU %	ail Ra ail R Jsed
Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13 Base at Zero 12:41:30 07/05 Sent Fail Rate: 0.0	2 Used	%: 48.03 5 Mins 5 Mins Sent Fa Rcvd F Mem U CPU %	ail Ra ail R Ised
CPU Count: Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13 Base at Zero	2 Used	%: 48.03 5 Mins 5 Mins 6 Revd F 8 Revd F 9 Mem U 0 CPU %	▼ [ail Ra ail R
Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13 Base at Zero Avail MB: 13 Base at Zero Avail MB: 13 Avail MB: 14 Avail	2 Used	%: 48.03 5 Mins 5 Mins 8 Sent Fa Rcvd F Mem U CPU %	▼ [ail Ra ail R Jsed
Start Time:	2 7/5/16 12:39 AM		Memory Max MB: 254	Avail MB: 13 Base at Zero Avail MB: 13 Base at Zero Avail MB: 13 Avail MB: 14 Avail	2 Used	%: 48.03	▼ [ail Ra ail R Jsed
Start Time:	2 7/5/16 12:39 AM	12:39:00	Memory Max MB: 254	Avail MB: 13 Base at Zero 12:41:30 07/05 Sent Fail Rate : 0.0 Rcvd Fail Rate : 0.0 Mem Used %: 48.0 CPU %: 0.1 12:41:00	2 Used Time Range:	%: 48.03	ail Ra ail R Jsed

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

Cluster	Select a cluster to display.
Detail	View "Node Detail" display.
Host	Select a host from the drop-down menu.
Location	Select a location from the drop-down menu. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
Id	The id for the selected node.
Member	The member name for this node.
Machine	The machine name for this node.
Rack	The rack name for this node.
Site	The site name for this node.
Role	The role name for this node.
Process	The process name for this node.
Uni Address	The unicast address. This is the IP address of the node's DatagramSocket for point-to-point communication.
Unicast Port	The unicast port. This is the port of the node's DatagramSocket for point-to-point communication.
CPU Count	Number of CPU cores for the machine this node is running on.
Start Time	The date and time that the selected node joined the cluster.
Cache Data	Cache Name Name of Cache. Tier Front or Back. Objects Number of objects. Units Number of units (typically bytes).
Packet Communication	Sent Cumulative number of packets sent by this node since the node statistics were last reset.
	Rcvd Cumulative number of packets received by this node since the node statistics were last reset.
	Resent% Cumulative number of packets resent by this node since the node statistics were last reset.
	Rptd% Cumulative number of packets repeated by this node since the node statistics were last reset.
	Send Queue The number of packets currently scheduled for delivery, including packets sent and still awaiting acknowledgment. Packets that do not receive an acknowledgment within the ResendDelay interval are automatically resent.
Memory	Max MB Total memory allocated. Avail MB Total memory available. Used% Percent of allocated memory being used.
Base at Zero	Use zero as the Y axis minimum for all graph traces.

Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last
5	7 Days, or display All Data. To specify a time range, click Calendar

 	·
Select or Enter Date and Time:	
Restore to Now	
Ok Apply Cancel]

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

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

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

- **Sent Fail Rate** Percentage of communication packages on this node that failed and needed to be resent.
- **Rcvd Fail Rate** Percentage of received communication packages that failed and needed to be repeated.
- Mem Used% Percent of memory used by the node.
- **CPU%** Percent of CPU used by the node.
Service Trends

Trend graphs showing metrics on a selected node of a selected service. Allows you to visually compare the behavior of metrics over time, for a given node.



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

Cluster	Select a cluster to display.
Service	Select a service to display.
Host	Select a host to display.
Location	Select a location to display I ocation is

Location Select a location to display. Location is a unique identifier for each node and defined as: member_name.machine.rack.site.

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

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

		-
S	elect or Enter Date and Time:	
	Restore to Now	
	Ok Apply Cancel	

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

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

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

Metrics for
Service selected
by LocationTrend chart displays the values of labeled Metrics for the selected Location
over the specified Time Range. Location is a unique identifier for each node
and defined as: member_name.machine.rack.site.

 $\ensuremath{\text{CPU\%}}$ CPU Utilization (as a percent) on the selected $\ensuremath{\text{Location}}$ (for example, node).

Requests Number of requests issued by the service in the measured period. **Messages** The number of messages for the given node in the measured interval.

Request Average Duration Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

Request Pending Count Number of pending requests issued by the service. **Task Backlog** Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads.

Active Threads Number of threads in the service thread pool, not currently idle.

Node Detail

This display presents detailed information about invocation services per node. The data on this display is queried from the Coherence MBeans. NOTE: For details on attributes of these MBeans go to: http://download.oracle.com/otn_hosted_doc/coherence/350/com/tangosol/ net/management/Registry.html.

4	n/a:DemoClust	ter		Summary	Single	e Node	e - Deta	il		05-JI	II-2016 1	2:44 🗳	🔊 Data OK 🕠	+ 💿
Hos	st: All Hosts		• L	ocation : St	oreNode01.	SLHOST	2					2	ld: 1	
	ld Avail MB	Max	MB P	kts Sent	Delta	Pkts	Rcvd	D	elta	Pkts Rp	td De	elta P	kts Resent	Delt
	1 14	41	254	3,105,267	729	2,8	850,223		668	1,0	94	0	1,256	;
•	III													- Þ.
	Name	Running	CPU %	6 Request1	FotalCoun	t Req	uests	Total	Message	es Me	ssages	Requ	uestAverag	JeDura
Ma	nagement	M	0	.3	2	2	0		323,6	511	8)		
•														
	Service		Running	Status	HA St	torage	CPU %	Re	equestTot	alCount	Requ	iests	Total Mess	sages
Dis	tributedCache		V	MACHINE-S	SAFE		0.	.3		378,458		96	1,33	8,919
•	III													- F.
	Service		(Cache Name	÷	Tier	Object	s	Hits	De	lta	Misse	es De	lta 🔺
Dis	tributedCache	1	near-foo			back	35,4	491	92,13	0	1	15,	671	3
Dis	tributedCache	h	near-foo2			back	37,4	413	96,27	8	4	15,	329	
Dis	tributedCache 🔪	<u> </u>	ioo2			back	35,5	587	119,80	5	12	28,	552	T
•			111											•
	Service		(Cache Name	÷	Event	sDispate	hed	Eviction	Count	Insert(Count	ListenerFi	lter 🔺
Dis	tributedCache		near-foo			_		0	_	4,440		39,931		11
Dis	tributedCache	1	ioo2					0		6,660	4	42,247		=
Dis	tributedCache	1	foo					0		38,015	-	77,698		~
								0		0 700		10 1 70		<u> </u>
4														P'

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

Cluster	Select a cluster to display.
Summary	View "Node Summary" display.
Host	Select a host.
Location	Select a location. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
Node MBean Data	This table contains data from the Node MBean for the selected node.

Invocation Service MBean	This table contains data from the Invocation Services MBean for the selected node.					
Data	StatusHA:					
	The high availability status of the service:					
	ENDANGERED: There is potential data loss in the cluster if a node goes offline.					
	NODE-SAFE: There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.					
	MACHINE-SAFE: There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.					
	RACK-SAFE: There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.					
	SITE-SAFE: There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.					
Cache Service MBean Data	This table contains data from the Cache Service and Node MBeans associated with the selected node, as well as the following data.					
Cache MBean Data	This table contains data from the Cache MBeans associated with the selected node.					
Storage Manager MBean Data	This table contains data from the Storage Manager MBeans associated with the selected node.					

JVM Summary

Runtime, class loader, thread, OS and input arguments. NOTE: Platform MBean information is available at: http://java.sun.com/javase/6/docs/api/java/lang/management/package-summary.html#package_description.

fn/a:DemoCluster		•	Singl	e Node - JVM Summary		05-Jul-2016 12:4	7 < 0	ata OK	+ •
Host: All Hosts		Location:	StoreNo	de01.SLHOST2		-		ld: 1	L
	Runtime	•		Operating	g Syster	n			
Start Time:	7/5/16 12	:31 AM		Operating System Name:	Window	vs XP			
Up Time:	0d 12:15			Version:	5.1				
				Architecture:	x86				
		lor		Available Processors:	2				
	Idss Ludi	ier		Percent CPU:	0.5				
Loaded	Classes:	2333							
Unloaded	Classes:	77		Total Swap Space Size	e: 4,09	6 MB			
				Free Swap Space Size	e: 2,63	0 MB			
Total Loaded C	lasses:	2410		Total Dhysical Memory	w 2.04	2 MR			
					y. 2,04				
C	ompilatio	on		Free Physical Memory	y: 716	MB			
Compilatio	on Time:	757 s		Committed Virtual Memory:	: 193	MB			
Complication									
	Threads	\$		VanualEGan					
Live	Threads:	16		-Dtangosol.coherence.mbeans=	=/sl-cust	om-mbeans.xm	l		
Daomon	[broade:	15		-Dcom.sun.management.jmxren	mote				
Daemon	nireaus.	15							
Peak	Threads:	16							



Cluster	Select a cluster to display.
Host	Select a host to display.
Location	Select a location to display. Location is a unique identifier for each node and defined as: member_name.machine.rack.site.
Id	This table contains data from the Node MBean for the selected node.
Runtime	Start Time The date and time that the JVM started. Up Time The uptime of the JVM.

Class Loader	Loaded Classes The number of classes that are currently loaded in the JVM.
	Unloaded Classes The total number of classes unloaded since the JVM started execution.
	Total Loaded Classes The total number of classes that have been loaded since the JVM started execution.
Compilation Time	The approximate accumulated elapsed time (in milliseconds) spent in compilation. If multiple threads are used for compilation, then this value is a summation of the approximate time that each thread spent in compilation. NOTE: Compilation Time monitoring may not be supported depending on the platform (for example, a Java virtual machine implementation).
Threads	Live Threads The number of live threads.
	Daemon Threads The number of live daemon threads.
	Peak Threads The peak live thread count since the Java virtual machine started or peak was reset.
Operating System	Operating System Name The operating system name.
	Version The operating system version.
	Architecture The operating system architecture.
	Available Processors The number of processors available to the JVM.
	Percent CPU Percent of CPU used by the JVM.
	Total Swap Space Size The value of the OperatingSystemMXBean's TotalSwapSpaceSize attribute.
	Free Swap Space Size The value of the OperatingSystem MXBean's FreeSwapSpaceSize attribute.
	Total Physical Memory The value of the OperatingSystemMXBean's TotalPhysicalMemorySize attribute
	Free Physical Memory The value of the OperatingSystemMXBean's FreePhysicalMemorySize attribute
	Committed Virtual Memory The value of the OperatingSystemMXBean's CommittedVirtualMemorySize attribute
Input Arguments	The list of JVM arguments in the RuntimeMXBean's InputArguments attribute.

JVM Memory Trends

Heap and non-heap memory trends. NOTE: Platform MBean information is available at: http://java.sun.com/javase/6/docs/api/java/lang/management/packagesummary.html#package_description.



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

Cluster	Select a cluster to display.
Host	Select a host to display.
Location	Select a location to display. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
Id	This table contains data from the Node MBean for the selected node.
Base at Zero	Use zero as the Y axis minimum for all graph traces.

Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar .						
	Select or Enter Date and Time:						
	Restore to Now						
	Ok Apply Cancel						
	By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: MMM dd, YYYY HH:MM . For example, Aug 21, 2011 12:24 PM .						
	Use the navigation arrows S b to move forward or backward one time period. NOTE: The time period is determined by your selection from the Time Range drop-down menu.						
	Click Restore to Now to reset the time range end point to the current time.						
Heap Memory	Maximum The value of the max field within the MemoryMXBean						
	Committed The value of the committed field within the MemoryMXBean HeapMemoryUsage attribute.						
	Used The value of the used field within the MemoryMXBean HeapMemoryUsage attribute.						
	Peak Tenured Used The value of the used field within the TenuredGen MemoryPoolMXBean PeakUsage attribute.						
Non-Heap Memory	Maximum The value of the max field within the MemoryMXBean NonHeapMemoryUsage attribute.						
-	Committed The value of the committed field within the MemoryMXBean NonHeapMemoryUsage attribute.						
	Used The value of the used field within the MemoryMXBean NonHeapMemoryUsage attribute.						
	Objects Pending Finalization The value of the MemoryMXBean ObjectPendingFinalizationCount attribute.						
	Verbose The value of the MemoryMXBean Verbose attribute.						
Garbage	name Name of the Garbage Collector MBean.						
Collection	LastGcInfo.Count The GcThreadCount from the Garbage Collector's LastGcInfo MBean.						
	LastGcInfo.Duration The Duration from the Garbage Collector's LastGcInfo MBean.						
	Operations Run Garbage Collector Executes the MemoryMXBean garbage collection operation, Reset Peak Usage Executes the TenuredGen resetPeakUsage operation.						

JVM GC Trends

Memory usage before and after garbage collection and Garbage Collector activity. NOTE: Platform MBean information is available at: http://java.sun.com/javase/6/docs/api/java/lang/ management/package-summary.html#package_description.





Cluster	Select a cluster to display.
Host	Select a host to display.
Location	Select a location to display. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
Id	This table contains data from the Node MBean for the selected node.
Garbage Collector	Select a Garbage Collector.
Max	Select to add the Max trace (graph will rescale if necessary).

- **Committed** Select to add the Committed trace (graph will rescale if necessary).
- **Base at Zero** Use zero as the Y axis minimum for all graph traces.

Time RangeSelect a time range from the drop down menu varying from 2 Minutes to Last7 Days, or display All Data. To specify a time range, click Calendar ____.



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

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

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

Memory Usage
(in MB) Before
and After
Garbage
CollectionMax The maximum amount of memory used by the node or nodes.Used - Before The amount of memory used by the node or nodes before
garbage collection.Used - Before The amount of memory used by the node or nodes before
garbage collection.

Duration The duration, in seconds, that memory is used by the node or nodes. **Duty** Cycle Percent of time spent by the node or nodes in garbage collection.

System Properties

Table of Java properties for a selected node.

← n/a:DemoCluster Single Nod	e - JVM System Properties 05-Jul-2016 12:53	< Data OK 🛛 💠 🕝
Host: All Hosts Cocation: StoreNoo	le01.SLHOST2	ld: 1
java.runtime.version 1.6.0_11-b03		
key	value	Connection
awt.toolkit	sun.awt.windows.WToolkit	DemoCluste
com.sun.management.jmxremote	true	DemoCluste
file.encoding	Cp1252	DemoCluste
file.encoding.pkg	sun.io	DemoCluste
file.separator		DemoCluste =
java.awt.graphicsenv	sun.awt.Win32GraphicsEnvironment	DemoCluste
java.awt.printerjob	sun.awt.windows.WPrinterJob	DemoCluste
java.class.path	;;";C:\rtvdemos\rtvoc_57c1\conf;C:\rtvdemos\rtvoc_57c2	DemoCluste
java.class.version	50.0	DemoCluste
java.endorsed.dirs	C:\Program Files\Java\jre6\lib\endorsed	DemoCluste
java.ext.dirs	C:\Program Files\Java\jre6\lib\ext;C:\WINDOWS\Sun\Jav	DemoCluste
java.home	C:\Program Files\Java\jre6	DemoCluste
java.io.tmpdir	C:\DOCUME~1\m\LOCALS~1\Temp\	DemoCluste
java.library.path	C:\WINDOWS\system32;.;C:\WINDOWS\Sun\Java\bin;C	DemoCluste
java.rmi.server.randomIDs	true	DemoCluste
java.runtime.name	Java(TM) SE Runtime Environment	DemoCluste
java.runtime.version	1.6.0_11-b03	DemoCluste
java.specification.name	Java Platform API Specification	DemoCluste
java.specification.vendor	Sun Microsystems Inc.	DemoCluste
java.specification.version	1.6	DemoCluste
java.vendor	Sun Microsystems Inc.	DemoCluste
java.vendor.url	http://java.sun.com/	DemoCluste
java.vendor.url.bug	http://java.sun.com/cgi-bin/bugreport.cgi	DemoCluste
java.version	1.6.0_11	DemoCluste T

Title Bar (possible features are):	🕼 Data OK Data connection state. Red indicates the Data
🗧 🏠 Open the previous and upper display.	Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the
Open an instance of this display in a new window.	data source is connected.
Open the online help page for this display. Menu Table open commonly accessed	might indicate the Monitor stopped running. Correct time
displays.	current and valid.
6,047 The number of items currently in the display.	Open the Alert Views - RTView Alerts Table display.

Cluster	Select a cluster to display.
Host	Select a host to display.
Location	Select a location to display. Location is a unique identifier for each node and defined as: member_name.machine.rack.site .
Id	This table contains data from the Node MBean for the selected node.
java.runtime.vers ion	The value of the RuntimeMXBeans's VmVersion attribute.
System Properties	This table contains the attribute/value pairs from the RuntimeMXBean's SystemProperties attribute.

Time Range Analysis

These displays allow you to compare data between two sets of time ranges.

- "Service Comparison" on page 148: Analyze service data for two sets of time ranges.
- "Cache Comparison" on page 149: Analyze cache data for two sets of time ranges.

Service Comparison

This display allows for analysis of service data for two sets of time ranges.

← n/a:[DemoCluster	-	TRA - S	ervice	Comparison	05-Jul-2	2016 12:55 < Data OK 🔶	0
Service:	DistributedCache	•	Sto	rage Nod	es 💿 Process Node	S		
Time Ra	nge Analysis:							
Time	Range 1:	2	Start:			End:		
Time	Range 2:		Start			End:		
-			-	_				
Time	Range 1: Data Bucket Tin	nestamp	Start	Jun 2	0, 2016 12:55:00 PM	End	Jul 5, 2016 12:54:22	2 PM
Time	Range 2: Data Bucket Tin	nestamp	Start	Jun 2	0, 2016 12:55:00 PM	End	Jul 5, 2016 12:54:22	2 PM
	Metric Name	Time R	ange 1 Valu	e	Time Range 2 \	/alue	Percentage Change	
DeltaMess	sages		444,7	34,253	44	14,734,253	0.0	00
DeltaRequ	uestTotalCount		123,8	83,421	12	23,883,421	0.0	00
DeltaTask	:Count			0		0		٠
TaskBack	log			22,701		22,701	0.0	00

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

Cluster	Select a cluster to display.
Service	Select a service to display.
Storage Nodes	Select to display storage node data in the trend graphs of this display.
Process Nodes	Select to display process node data in the trend graphs of this display.

Time Range	Time Range 1: Set Start and End times for Time Range 1				
Analysis	Time Range 2: Set Start and End times for Time Range 2				
	Time Range 1: Data Bucket Timestamp and Time Range 2: Data Bucket Timestamp displays the Start and End timestamps for the actual data buckets used in the comparison, since data may be compacted into buckets with different Start and End times from the specified values.				
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar				
	Select or Enter Date and Time:				
	Restore to Now				
	Ok Apply Cancel				
By default, the time range end point is the current time. To change the tim range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following forr MMM dd, YYYY HH:MM. For example, Aug 21, 2011 12:24 PM.					
	Use the navigation arrows S to move forward or backward one time period. NOTE: The time period is determined by your selection from the Time Range drop-down menu.				
	Click Restore to Now to reset the time range end point to the current time.				

Cache Comparison

This display allows for analysis of cache data for two sets of time ranges.

n/a:DemoCluster		05-Jul-20	16 12:57	< Data OK	+ 📀			
Service: DistributedCache	🔽 Cach	e: foo		-	Front			
Time Range Analysis:								
Time Range 1:	Star	t:			End:			
Time Range 2:	Sta	rt 📃			End:			
Time Range 1: Data Bucket Ti	mestamp S	tart	Jun 20. 2016 1:0	0:00 PM	End	Jul	5. 2016 12:	56:34 PI
			outreo,					
Time Range 2: Data Bucket Ti	nestamp S	tart	Jun 20, 2016 1:0	0:00 PM	End	Jul	5, 2016 12:	56:34 PI
Metric Name	Time Rang	e 1 Value	Time R	ange 2 Val	ue	Percent	age Chang	e
DeltaTotalGets		88,477,2	277	88,477,277		0		0.00
DeltaTotalPuts		51,678,7	42	51,678,742				0.00
DeltaCacheHits		62,646,428		62,646,428				0.00
DeltaStoreReads			0	0				*
DeltaStorewrites		0		0				~
DeltaStoreReadMillis		0		0				
DeltaStoreFailures			0		0			- V

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

Cluster	Select a cluster to display.
Service	Select a service to display.
Storage Nodes	Select to display storage node data in the trend graphs of this display.
Process Nodes	Select to display process node data in the trend graphs of this display.
Time Range Analysis	Time Range 1 : Set Start and End times for Time Range 1 Time Range 2 : Set Start and End times for Time Range 2 Time Range 1: Data Bucket Timestamp and Time Range 2: Data Bucket Timestamp displays the Start and End timestamps for the actual data buckets used in the comparison, since data may be compacted into buckets with different Start and End times from the specified values.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar

 ······································
Select or Enter Date and Time:
Restore to Now
Ok Apply Cancel

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

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

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

OC Administration

These displays allow you to manage your Oracle Coherence metrics, nodes and caches. Some of these displays might be read-only depending on your login credentials.

 "OC Metrics Administration" on page 151: Monitor information on metrics acquisition. Permits user to reset system metrics.

- "Cluster MBean Servers" on page 153: Access this display using the Cluster MBean Servers button in the "OC Metrics Administration" display. Permits user to find and choose a different MBean server.
- "Management Settings" on page 155: Monitor information about Coherence JMX management settings.
- "Alert Views" on page 162: Permits user to modify node parameters.
- "Alert Views" on page 162: Permits user to modify cache parameters.

OC Metrics Administration

This display allows various statistics to be reset, so that cumulative data can be visualized more meaningfully. It is read-only unless you are logged in as admin or super.





Cluster	Select a cluster to display.
Connection	The name of the JMX connection used to access the cluster data.
Reset Service Statistics	Click to reset the cumulative counts of the service statistics.
Reset Cache Statistics	Click to reset the cumulative counts of the cache statistics.
Reset Node Statistics	Click to reset the cumulative counts of the node statistics.
Reset Connection Statistics	Click to reset the cumulative counts of the connection statistics.
Total MBeans Queried	 Node MBeans Total number of node MBeans queried. Service MBeans Total number of service MBeans queried. Cache MBeans Total number of cache MBeans queried. Storage MBeans Total number of storage MBeans queried. JVM Platform MBeans Total number of JVM platform MBeans queried. Total number of MBeans queried.
RTView Memory Utilization for Metrics History	 By default, the Oracle Coherence Monitor stores several hours of data using inmemory tables. Cache Metrics, # Rows The number of table rows used by the Monitor to store cache metrics data. Cache Metrics, Memory Used (KB) The amount of memory (KB) used by the Monitor to store cache metrics data. Node Metrics, # Rows The number of table rows used by the Monitor to store node metrics data. Node Metrics, Memory Used (KB) The amount of memory (KB) used by the Monitor to store node metrics data. Node Metrics, Memory Used (KB) The amount of memory (KB) used by the Monitor to store node metrics data.
Cluster MBean Servers	Click to open the "Cluster MBean Servers" display which lists the currently detected remote JMX management enabled MBean Servers in the selected cluster. If your MBean server goes down, use this display to find and choose a different available MBean server.

Total RTView JMX
Query TimeTraces the t otal amount of time, in milliseconds, to query the monitoring
MBeans from Coherence.

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

Select or Enter Date and Time:
Restore to Now
Ok Apply Cancel

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

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

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

Cluster MBean Servers

Access this display using the **Cluster MBean Servers** button in the "OC Metrics Administration" display.

View a list of URLs for all currently detected remote JMX management enabled MBean Servers in a cluster. Information displayed includes the hostname and IP address of the cluster node, and the port used for remote JMX management.

n/a:DemoClu	uster	Cluster MBean Servers 11-Apr-2017 10:05			5 💠 Data OK 💠 🕜				
			Cluster MB	ean	Servers				
Connection =	nodeld	HostName	E IpAddress	Ξ	port	Ξ	Location	Ξ	InputArguments
DemoCluster		30 localhost	0.0.0			9971	JmxNode01.SLHOST1		-Xmx256m;-Dswin
	•			Ш					•

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

Cluster	Select a cluster to display.
Connection	The name of the JMX connection used to access the cluster data.
nodeld	The unique identifier for the MBean Server.
HostName	The name of the host for the MBean Server.
IpAddress	The IP address for the MBean Server.
port	The port number for the MBean Server.
Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site .
InputArguments	A list of JVM arguments in the Runtime JMX MBean's InputArguments attribute.
Expired	When checked, this connection is expired due to inactivity.

Management Settings

This display is read-only unless you are logged in as admin or super.

÷	n/a:DemoCluster	· 🗸	Mana	gement Settings	05-Jul-2016 13:00	< Data OK	+	۲
	JMX Connection:	DemoCluster						
	JMX Management	Settings (Current)		Enter New Settings:	type new value and	l Enter)		
	Refres	h Policy:refresh-expi	red	Refresh Policy:	refresh-ahead			
	Expir	ry Delay:1000		Expiry Delay:	0			
	Refresh O	n Query						
	Refres	sh Time:Jul 5, 2016 1:	:00:20 PM					
	Refres	h Count1,482,957	3,280					
	Refresh Exces	ss Coun <mark>t1,337</mark>	0					
	Refresh Timeo	ut Coun <mark>5</mark>	0					
	Refresh Predictio	on Coun <mark>t</mark> t	0					
			(Delta)					

 Title Bar (possible features are): Open the previous and upper display. Open an instance of this display in a new window. Open the online help page for this display. Menu , Table open commonly accessed displays. 	 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected. 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
0,0007 The number of items currently in the display.	Open the Alert Views - RIView Alerts Table display.

Cluster

Select a cluster to display.

JMX Connection The name of the JMX connection used to access the cluster data.

JMX Management

Settinas

Refresh Policy:

Select a refresh policy from the drop-down list.

refresh-expired Each MBean will be refreshed from the remote node when it is accessed and the expiry delay has passed from the last refresh (same functionality as in pre-3.4 Coherence releases. This option is the default setting and is best used when MBeans are accessed in a random pattern.

refresh-ahead MBeans are refreshed before they are requested based on prior usage patterns after the expiry delay has passed, reducing latency of management information with a minor increase in network consumption. This option is best when MBeans are accessed in a repetitive/programmatic pattern.

refresh-behind Each MBean will be refreshed after the data is accessed, ensuring optimal response time. However, note that the information returned will be offset by the last refresh time.

 ${\it refresh-onquery}$ Select this option if the refresh-on-query MBeanServer is configured.

Expiry Delay:

Duration (in milliseconds) that the MBeanServer will keep a remote model snapshot before refreshing.

Refresh on Query:

Specifies whether or not the refresh-on-query MBeanServer is configured. If so, then set the RefreshPolicy to refresh-onquery.

Refresh Time The timestamp when this model was last retrieved from a corresponding node. For local servers it is the local time.

Refresh Count* The total number of snapshots retrieved since the statistics were last reset.

Refresh Excess Count The number of times the MBean server predictively refreshed information and the information was not accessed. Delta values show the change in the counts within the most recent JMX retrieval period.

Refresh Timeout Count* The number of times this management node has timed out while attempting to refresh remote MBean attributes.

Refresh Prediction Count* The number of times the MBeanServer used a predictive (refresh-behind, refresh-ahead, refresh-onquery) algorithm to refresh MBean information.

Ti

di

Node Administration

This display allows the user to view and change settings for individual Nodes. It is read-only unless you are logged in as super. Click on the desired Node to select that Node. Change the data item in the bottom half of the display and press Return to make the change. All data on this display is queried from and set on the Coherence ClusterNodeMBean.

← n/a:DemoCluster 🔽	Node	Administration		05-Jul-2016 13:01	< Data OK	+ 📀
	Current Se (Select a No	ettings for all Nodes de to change settings)				
Location	BufferPublishSize Bu	fferReceiveSize BurstC	ount Bu	rstDelay Multio	astThreshold	d
ExtendNode01.SLHOST4	32	1428	0	10	2	25 =
ExtendNode02.SLHOST2	32	1428	0	10	2	25
ExtendNode02.SLHOST4	32	1428	0	10	2	25
JmxNode01.SLHOST1	32	1428	0	10	2	25
OLS-ProcessNode-1.SLHOST2	32	1428	0	10	2	25
OLS-ProcessNode-1.SLHOST4	32	1428	0	10	2	25
OLS-StoreNode-1.SLHOST2	32	1428	0	10	2	25
OLS-StoreNode-1.SLHOST4	32	1428	0	10	2	25
ProcessNode01.SLHOST2	32	1428	0	10	2	25
ProcessNode01.SLHOS13	32	1428	0	10	2	5
	32	1//28	0	10		- F
Selected Node(s): All		Select	All Node	S		
Burst Count: 0		Shutdown Node				
Burst Delay: 0		Ensure Cache Service:	Distrik	outedCache		
Logging Level: 0		Ensure Invocation Service:				
Logging Limit: 0		Multicast Threshold	d: 0			
Logging Format:		Resend Delay	<i>ı</i> : 0			
		Send Ack Delay	: 0			
		Traffic Jam Count	: 0			
		Traffic Jam Delay:	0			
le Bar (possible features are):		Server is not recei	connect ving dat	tion state. Rec ta or the Displ	l indicates t ay Server is	he Dat s not
Open the previous and upp	receiving data from	n the Da	ata Server. Gr	een indicate	es the	
Open an instance of this display	in a new window.		necced.			
Open the online help page for th	is display.	23-Mar-2017 12:04	Current	date and time	a. Incorrect	time
Menu 💌, Table open comr	might indicate the Monitor stopped running. Correct and green Data OK icon is a strong indication that o				t time data i	

Cluster Select a cluster to display.

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

JMX Connection The name of the JMX connection used to access the cluster data.

Open the Alert Views - RTView Alerts Table display.

Current Settings for All Nodes	Location A unique identifier for each node. It is defined as: member_name.machine.rack.site.
	BufferPublishSize The buffer size of the unicast datagram socket used by the Publisher, measured in the number of packets. Changing this value at runtime is an inherently unsafe operation that will pause all network communications and may result in the termination of all cluster services.
	BufferReceiveSize The buffer size of the unicast datagram socket used by the Receiver, measured in the number of packets. Changing this value at runtime is an inherently unsafe operation that will pause all network communications and may result in the termination of all cluster services.
	BurstCount The maximum number of packets to send without pausing. Anything less than one (e.g. zero) means no limit.
	BurstDelay The number of milliseconds to pause between bursts. Anything less than one (e.g. zero) is treated as one millisecond.
	MulticastThreshold The percentage (0 to 100) of the servers in the cluster that a packet will be sent to, above which the packet will be multicasted and below which it will be unicasted.
	ResendDelay The minimum number of milliseconds that a packet will remain queued in the Publisher`s re-send queue before it is resent to the recipient(s) if the packet has not been acknowledged. Setting this value too low can overflow the network with unnecessary repetitions. Setting the value too high can increase the overall latency by delaying the re-sends of dropped packets. Additionally, change of this value may need to be accompanied by a change in SendAckDelay value.
	SendAckDelay The minimum number of milliseconds between the queueing of an Ack packet and the sending of the same. This value should be not more then a half of the ResendDelay value
	TrafficJamCount The maximum total number of packets in the send and resend queues that forces the publisher to pause client threads. Zero means no limit.
	TrafficJamDelay The number of milliseconds to pause client threads when a traffic jam condition has been reached. Anything less than one (e.g. zero) is treated as one millisecond.
	LoggingLevel Specifies which logged messages will be output to the log destination. Valid values are non-negative integers or -1 to disable all logger output.
	LoggingLimit The maximum number of characters that the logger daemon will process from the message queue before discarding all remaining messages in the queue. Valid values are integers in the range [0]. Zero implies no limit.
	LoggingFormat Specifies how messages will be formatted before being passed to the log destination
	LoggingDestination The output device used by the logging system. Valid values are stdout, stderr, jdk, log4j, or a file name.
	nodeld The short Member id that uniquely identifies the Member at this point in time and does not change for the life of this Member.
	ProcessName A configured name that should be the same for Members that are in the same process (JVM), and different for Members that are in different processes. If not explicitly provided, for processes running with JRE 1.5 or higher the name will be calculated internally as the Name attribute of the system RuntimeMXBean, which normally represents the process identifier (PID).
Selected Node(s)	Lists the nodes selected in the table.
Select All Nodes	Click to select all nodes.
Shutdown Node	Stop all the clustered services running at this node (controlled shutdown). The management of this node will node be available until the node is restarted (manually or by programming).

Ensure Cache Service	Ensure that a CacheService for the specified cache runs at the cluster node represented by this MBean. This method will use the configurable cache factory to find out which cache service to start if necessary. Return value indicates the service name; null if a match could not be found.
Ensure Invocation	Ensure that an InvocationService with the specified name runs at the cluster node represented by this MBean.

Cache Administration

This display allows the user to view and change settings for individual caches. It is read-only unless you are logged in as super. Click on the desired cache to select that cache. Change the data item in the bottom half of the display and press Return to make the change. The data on this display is queried from and set on the Coherence CacheMBean.

← n/a:DemoCluster 🔽		Cache Administrati	on	05-Jul-2016	13:03 < Data	ок 💠 💮
Service: DistributedCache	Cache:	foo	•	Front		
Location	LowUnits	HighUnits I	BatchFactor	ExpiryDelay	FlushDelay	Queue 🔺
StoreNode01.SLHOST1	8,000,00	0 12,000,000	0.0	0	0	
StoreNode01.SLHOST2	8,000,00	0 12,000,000	0.0	0	0	
StoreNode01.SLHOST4	8,000,00	0 12,000,000	0.0	0	0	
StoreNode04.SLHOST1	8,000,00	0 12,000,000	0.0	0	0	=
StoreNode04.SLHOST4	8,000,00	0 12,000,000	0.0	0	0	
StoreNode05.SLHOST2	8,000,00	0 12,000,000	0.0	0	0	
StoreNode05.SLHOST4	8,000,00	0 12,000,000	0.0	0	0	
StoreNode05n.SLHOST1	8,000,00	0 12,000,000	0.0	0	0	
StoreNode05n.SLHOST2	8,000,00	0 12,000,000	0.0	0	0	
StoreNode05n.SLHOST4	8,000,00	0 12,000,000	0.0	0	0	~
•	111					4
Selected Node(s): All			Select All N	odes		
High Units: 0		Batch Fact	or: 0.0			
Low Units: 0		Expiry Del	ay: O			
		Flush Del	ay: 0			
		Queue Del	ay: 0			
		Refresh Fac	tor: 0.0			
		Requeue Thresho	old: 0			

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

Cluster Select a cluster to display.

Service Select the service to display.

Cache	Select the cache to display.
Front	Select for front tier, deselect for back tier.
Current Settings for all Nodes on	Location A unique identifier for each node. It is defined as: member_name.machine.rack.site.
Selected Cache	LowUnits The number of units to which the cache will shrink when it prunes. This is often referred to as a `low water mark` of the cache.
	HighUnits The limit of the cache size measured in units. The cache will prune itself automatically once it reaches its maximum unit level. This is often referred to as the `high water mark` of the cache.
	BatchFactor The BatchFactor attribute is used to calculate the `soft-ripe` time for write-behind queue entries. A queue entry is considered to be `ripe` for a write operation if it has been in the write-behind queue for no less than the QueueDelay interval. The `soft-ripe` time is the point in time prior to the actual `ripe` time after which an entry will be included in a batched asynchronous write operation to the CacheStore (along with all other `ripe` and `soft-ripe` entries). This attribute is only applicable if asynchronous writes are enabled (for example, the value of the QueueDelay attribute is greater than zero) and the CacheStore implements the storeAll() method. The value of the element is expressed as a percentage of the QueueDelay interval. Valid values are doubles in the interval [0.0, 1.0].
	ExpiryFactor The time-to-live for cache entries in milliseconds. Value of zero indicates that the automatic expiry is disabled. Change of this attribute will not affect already-scheduled expiry of existing entries.
	FlushDelay The number of milliseconds between cache flushes. Value of zero indicates that the cache will never flush.
	QueueDelay The number of seconds that an entry added to a write-behind queue will sit in the queue before being stored via a CacheStore. Applicable only for WRITE-BEHIND persistence type.
	RefreshFactor The RefreshFactor attribute is used to calculate the `soft-expiration` time for cache entries. Soft-expiration is the point in time prior to the actual expiration after which any access request for an entry will schedule an asynchronous load request for the entry. This attribute is only applicable for a ReadWriteBackingMap which has an internal LocalCache with scheduled automatic expiration. The value of this element is expressed as a percentage of the internal LocalCache expiration interval. Valid values are doubles in the interval[0.0, 1.0]. If zero, refresh-ahead scheduling will be disabled.
	Requeue Threshold The maximum size of the write-behind queue for which failed CacheStore write operations are requeued. If zero, the write-behind requeueing will be disabled. Applicable only for WRITE-BEHIND persistence type
	nodeld The node ID.
Selected Node(s)	Lists the nodes selected in the table.

Select All Nodes Click to select all nodes in the table.

RTView Cache Tables

View data that RTView is capturing and maintaining. Drill down and view details of RTView Cache Tables. Use this data for debugging. This display is typically used for troubleshooting with Technical Support.

Choose a cache table from the upper table to see cached data.

÷			RT	View C	ache	Tables		05-Jul-2016 13	3:09 < 🕎 Data C	ж 🕂 🕜
DataServer: <de< th=""><th>fault></th><th></th><th>R</th><th>rView (</th><th>Cache</th><th>Tables</th><th>Max</th><th>Rows:4000</th><th>Histo</th><th>ory Tables</th></de<>	fault>		R	rView (Cache	Tables	Max	Rows:4000	Histo	ory Tables
Ca	acheTable		TableT	ype		Rows	Co	lumns	Memor	у 🖍
JmxStatsTotals			current				1	4		441
OcBadCommunicati	ionNodes		current			14	10	6		14,999
OcCacheServiceSta	ats		current			8	38	58		62,666
OcCacheServiceTot	tals		current				8	26		4,441
OcCacheStats			current			17	72	80	2	206,148
OcCacheTotals			current			1	17	52		13,406
OcClusterOverview			current				1	7		791
OcClusterStats			current				1	19		14,103
OcExtendConnectio	ons		current			11	12	30		68,304
OcInvocationService	eStats		current			(63	60		62,252
OcInvocationService	eTotals 👆		current				1	26		2,841
OcJmxConnection			current				2	7		1,254
OalmyMamtData			ourront					15		1 754
			Oclr	nvocati	onSer	rviceStats			Rows:	63
TIME_STAMP	BackupCou	OwnedParti	OwnedParti	Refres	hTim	RequestAve	RequestMa	RequestPer	RequestPer I	Reques 🔺
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	85.5	109	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	8	16	0	0	-
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	7.5	15	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	16	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	15.5	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	24	32	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	31.5	47	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	39.5	63	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2	016 1	16	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5 2	016.1	8	16	0	0	v
 Ⅲ 										Þ

Title Bar (possible features are):

- 🔶 个 Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

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

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

Open the Alert Views - RTView Alerts Table display.

About

This display shows details about the Solution Package version and data sources available to your system.

RTView(R) Display Server - Oracle Co	oherence Monitor	
Version: 6.4.0.0 ALPHA		
Configuration: OCM 6 4 0 0 20160610 (000 21979-alpha 115	
Build Number: 000 21979		
Bullu Nulliber. 000.21975		
Available Data Sources:		
Alert		
Cache Datasource		
JMX		
LOG4J		
SQL		
XML		

Alert Views

These displays present detailed information about all alerts that have occurred in your system. These displays present performance data for your system. Enable and disable alerts in the "Alert Administration" display. Displays in this View are:

 "Alert Detail Table" on page 162: Time ordered list of all alerts that have occurred in the system.

Alert Detail Table

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

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

← Admin					Alert Detail Table	' 04-N	lov-2015 15:36 < Data OK 💠 🕐	
Alert Name Filter:	All Alert	Types			📕 🔲 Show Critical A	lerts Only 🛛 🔲 Sh	Show Cleared Alerts (214)	
Alert Text Filter:					Owner Filter: All	🗾 🔳 Sh	ow Acknowledged Alerts (1)	
Total Crit	otal Critical Warning						🌍 Alert Settings Conn OK	
31 2	4	13	Selecto	ne or mor	e alerts to enable acti	on buttons below)		
Time -	ID	Clr'd	Ack'd	Owner	Alert Name	Alert Index	^	
11/10/14 15:58:53	12150				BwProcessExecutionTime	slxp10(slapm)~domains	High Warning Limit exceeded, cu	
11/10/14 15:10:14	11993				BwEngineMemUsedHigh	slel4-64(slmon)~domaii	High Alert Limit exceeded, currer	
11/10/14 15:04:12	11969				BwServerFreeMemLow	slel4-64(slmon)	Low Warning Limit exceeded, cu	
11/10/14 14:23:12	11839				HostMemoryUsedHigh	myHawkDomain~slel4-	High Alert Limit exceeded, currer 🗏	
11/08/14 00:07:00	1007				BwEngineStopped	slapm(slapm)~domains	Engine has stopped	
11/08/14 00:07:00	1002				JvmNotConnected	localhost~domainslapm	Server disconnected	
10/31/14 14:01:36	1040828				HawkAlert	SLHOST6(domain6)~13	System Uptime changed to 0 dar	
10/28/14 16:38:01	1035056				HawkAlert	slapm(slapm)~2	System uptime changed to 14 da	
10/27/14 12:34:55	1031840				BwEngineStopped	slvmrh2(slapm)~domai	Engine has stopped	
10/27/14 12:34:55	1031839				BwEngineStopped	slvmrh2(slapm)~domai	Engine has stopped	
10/24/14 00:16:36	1015259				HawkAlert	SLHOST6(domain6)~12	Service Print Spooler is running.	
10/16/14 08:18:51	984247				HostMemoryUsedHigh	myHawkDomain~slhpu	High Warning Limit exceeded, cu	
10/03/14 15:50:05	943834				HawkAlert	SLHOST6(domain6)~11	Server Processes are at 59.0	
09/12/14 11:16:21	892842				BwEngineStopped	slvmware(slmon)~dom	Engine has stopped	
09/12/14 11:16:21	892841				BwEngineStopped	slvmware(slmon)~dom	Engine has stopped	
09/12/14 11:16:21	892840				BwEngineStopped	slvmware(slmon)~dom	Engine has stopped	
09/04/14 19:54:36	883519				HostMemoryUsedHigh	myHawkDomain~slvmr	High Alert Limit exceeded, currer 🔻	
•							4	
Selected Alert(s):								
				Ackno	owledge One Alert	Set Owner and Comn	nents See Details	

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

Row Color Code:

Tables with colored rows indicate the following:

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

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

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

Fields and Data

This display includes:

Alert Name Filter Select from a list of alert types or select All Alert Types. Filters limit display content and drop down menu selections to only those items that pass through the selected filter's criteria. Therefore if no items match the filter, you may see nothing in a given display and may not have any options available in the drop-down menu(s). NOTE: Filter selection is disabled on drill down summary displays.

Snow Critical Alerts Only	If selected, only currently critical alerts are shown in the table. Otherwise, all active alerts are shown in the table.					
Show Cleared Alerts	If selected, cleare	ed alerts are shown in the table.				
Alert Text Filter	Enter all or part of displays all alerts supported.	of the Alert Text to view specific alerts. For example, High selects and s that include High in the Alert Text. NOTE: Wild card characters are				
Owner Filter	Select the alert C	Dwner to show alerts for in the table.				
	All	Shows alerts for all Owners in the table: Not Owned and Owned By Me alerts.				
	Not Owned	Shows only alerts without Owners in the table.				
	Owned By Me	Shows only alerts for the current user in the table.				
Show Acknowle dged Alerts	If selected, ackno	owledged alerts are shown in the table.				
Total	Total number of alerts.					
Critical	Number of critical alerts.					
Warning	Total number of alerts that are currently in a warning state.					
Alert Settings Conn OK	The Alert Server Disconnected. Connected.	connection state:				

Alerts Table This table lists all active alerts for the current filters.

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

	Set Owner	and Comments	
ID: 28322	Source:		
Enter Owner: admin			
Enter Comment:			
Set Owner	Add Comment	Clear Comments	61

the Set Owner and Comments dialog.

See Details

Set Owner and Comments

Select an alert from the Current Alerts table and click to open the Set Owner and Comments dialog.

Select one or more alerts from the Current Alerts table and click to open

	Alert De	tail
Alert Time:	06/28/12 10:30:42	📕 Acknowledged
ID:	283221	🗾 Cleared
Name:	BwProcessExecution	Severity: 2
Index:	slapm(slapm)~domair	
Owner:		
Alert Text:	High Alert Limit current value:	exceeded,
Comments:		

Administration

These displays enable you to set alert thresholds, observe how alerts are managed, and view internal data gathered and stored by RTView (used for troubleshooting with SL Technical Support). Some of these displays may be read-only depending on your login. Displays in this View are:

- "Alert Administration" on page 166: Displays active alerts and provides interface to modify and manage alerts.
- "Metrics Administration" on page 173: Administer and configure data collection and processing options.
- "RTView Cache Tables" on page 174: View cached data that RTView is capturing and maintaining, and use this data use this for debugging with SL Technical Support.
- "RTView Cache Tables" on page 174: This display shows Monitor version details and the data sources available to your system.

Alert Administration

This section includes:

- "Tabular Alert Administration" on page 170
- "Setting Override Alerts" on page 171

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

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

Global Thresholds

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

The manner in which global alerts are applied depends on the CI Type. For example, the EMS Monitor has queue alerts, topic alerts and server alerts. When a queue alert is applied globally, it is applied to all queues on all servers. Likewise, a server alert applies to all servers, and a topic alert applies to all topics on all servers.

Override Thresholds

Setting override alerts allows you to set thresholds for a single resource (for example, a single server). Override alerts are useful if the majority of your alerts require the same threshold setting, but there are other alerts that require a different threshold setting. For example, you might not usually be concerned with execution time at a process level, but perhaps certain processes are critical. In this case, you can apply alert thresholds to each process individually.

To apply an individual alert you Index the Monitored Instance or resource. The Index Types available are determined by the CI Type. For example, the EMS Monitor lets you set an alert for a specific *topic* on a specific *server* (such as the PerServerTopic Index option), rather than for all topics on all servers.

	Alert Adminis	04-	Nov-2015 15:36 📢) Data OK 🔶 🕜		
Alert Filter: Clear	🍘 Alert Engine Enable	d Disable		🌍 Alert	Settings Conn OK	
Alert	Warning Level	Alarm Level	Duration	Alert Enabled	Override Count	
AcwInstanceCpuHigh	40	50	60		-1	
AcwInstanceDiskReadBytesHigh	10000	20000	30		-1	
AcwInstanceDiskReadOpsHigh	100	200	30		-1	
AcwInstanceDiskWriteBytesHigh	1000000	2000000	30		-1	
AcwInstanceDiskWriteOpsHigh	100	300	30		-1	
AcwInstanceNetworkReadBytesHigh	1000000	20000	30		-1	
AcwInstanceNetworkWriteBytesHigh	10000	20000	30		-1	
AmxServiceHitRateHigh	160	200	60	V	-1	
AmxServiceNodeFaultRateHigh	200	400	30		-1	
AmxServiceNodeHitRateHigh	75	100	60	~	-1	
AmxServiceNodeMovingAvgHitRateHigh	200	400	30		-1	
AmxServiceNodeMovingAvgResponseTimeHigh	200	400	30		-1	
AmxServiceNodeResponseTimeHigh	5	6	30		-1	
AmxServiceResponseTimeHigh	5	6	60		-1	
BirdExpired	NaN	NaN	0		-1	
BirdTooHigh	1600	2001	0		-1 👻	
	Settings for Sele	ected Alert				
Name: <select alert="" edit="" from="" one="" table="" the="" to=""> Warning Level: Duration (Secs.):</select>						
Description:		Alarm Le	vel:	Enable	d:	
				St	ave Settings	

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

Fields and Data This display includes:

Alert Filter	Enter the (case-sensitive) string to filter the table by the Alert table column value. NOTE: Partial strings can be used without wildcard characters. Press <enter></enter> or click elsewhere in the display to apply the filter.
Clear	Clears the Alert Filter entry.
Alert Engine Enabled	 Alerting is disabled. Alerting is enabled (by default).
Disable	Suspends all alerting.
Alert Settings Conn OK	The Alert Server connection state: Disconnected. Connected.

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

NOTE: To filter the alerts shown in the table by Solution Package, use the **\$rtvAlertPackageMask** substitution.

Alert	The name of the alert.
Warning Level	The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.
Alarm Level	The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.
Duration (Secs)	The amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution.
Alert Enabled	When checked, the alert is enabled globally.
Override Count	The number of times thresholds for this alert have been defined individually in the Tabular Alert Administration display.

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

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

Name	The name of the alert selected in the Active Alert Table.
Description	Description of the selected alert. Click Calendar for more detail.
Warning Level	Set the Global warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value.
	NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the warning to occur sooner, increase the Warning Level value. To set the warning to occur later, reduce the Warning Level value.
Alarm Level	Set the Global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value.
	NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value.
Duration	Set the amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution. This setting is global.
Enabled	Check to enable alert globally.
Save Settings	Click to apply alert settings.
Override Settings	Click to open the Tabular Alert Administration display to set override alerts on the selected alert.

Tabular Alert Administration

Set override alerts (override global alert settings). This display opens when you select an alert in the Alert Administration display and then select Override Settings.

For step-by-step instructions setting thresholds for individual alerts, see Setting Override Alerts.

(Tabular Alert Administration	10-N	ov-2014 09:35	🌍 Data (ж ^ ?
	Overrid	e Settings For Alert: TbeBackingStoreLoad	RateHigh	🏠 🚱	llert Settin	gs Conn OK
Index Type		Index Δ	Override Settings	Warning Level	Alarm Level	Alert Enabled
PerBECache	new51Cache~be_ger	n_Events_CreateAccount	~	80	95	~
Index Type: PerBEC	ache			-	_	
Index: new51C	ache~be_gen_Events	_CreateAccount	Add	Remove	Save	Settings
·	Unassigr	ned Indexes		Alert Se	ettings	
Conn	ection	beCacheName			000	
new51Cache		be_gen_Concepts_Account	War	ning Level:	80.0	
new51Cache		be_gen_Events_AccountOperations	A	larm Level:	95.0	
new51Cache		be_gen_Events_Debit	Alart Enchlady 🗹			
new51Cache		be_gen_Events_Deposit Alert Enabled:			iea: 🗖	
new51Cache		be_gen_EraudCriteria		oven	lae Setti	igs: 🖻
new51Cache		com_tibco_cep_runtime_model_element			Back	to Alerts

Fields and Data

This display includes:

Alert Settings The connection state. No servers are found.

Conn ÕK

One or more servers are delivering data.

Override Settings For Alert: (name)

This table lists and describes alerts that have override settings for the selected alert. Select a row to edit alert thresholds. The selected item appears in the **Index** field. Edit settings in the **Alert Settings** fields, then click Save Settings.

> Index Туре

Select the type of alert index to show in the Values table. Options in this drop-down menu are populated by the type of alert selected, which are determined by the CI Type. For example, the EMS Monitor has the following Index Types:

- PerServer: Alert settings are applied to a specific server.
- · PerQueue: Alert settings are applied to the queue on each server that has the queue defined.
- · PerServerQueue: Alert settings are applied to a single queue on a specific server.
- PerTopic: Alert settings are applied to the topic on each server that has the topic defined.
- PerServerTopic: Alert settings are applied to a single topic on a specific server.

	Index	The value of the index column.			
	Override Settings	When checked, the override settings are applied.			
	Alert Enabled	When checked, the alert is enabled.			
Index Type	Select the in example, to single serve alert selecte on the Pack	Select the index type. The index type specifies how to apply alert settings. For example, to a queue (topic or JVM, and so forth) across all servers, or to a queue on a single server. NOTE: Options in this drop-down menu are populated by the type of alert selected from the Alert Administration display. Index Types available depend on the Package installed.			
Index	The selected the Unassi g	d index column to be edited. This field is populated by the selection made in gned Indexes table.			
Unassigned Indexes	This table lists all possible indexes corresponding to the Index Type chosen in the drop- down list. Select a row to apply individual alert thresholds. The selected item appears in the Index field. Edit settings in the Alert Settings fields, then click Add .				
Add	Click to add	changes made in Alert Settings, then click OK to confirm.			
Remove	Click to rem confirm.	nove an alert selected in the Index Alert Settings table, then click OK to			
Save Settings	Click to save	e changes made to alert settings.			

Alert Settings Select a topic, server or queue from the **Unassigned Indexes** table and edit the following settings.

Warnin Level	g Set the warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value.		
	NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the warning to occur sooner, increase the Warning Level value. To set the warning to occur later, reduce the Warning Level value.		
	Click Save Settings to save settings.		
Alarm Level	Set the alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value. Click Save Settings to save settings.		
Alert Enabled	Check to enable the alert, then click Save Settings .		
Overrid Setting	 Check to enable override global setting, then click Save Settings. 		
Returns to the Administration - Alert Administration display.			

Back to Alerts

Setting Override Alerts

Perform the following steps to set an override alert. Index Types available depend on the CI Type. In this example, we use the EMS Monitor Package to illustrate.

NOTE: To turn on an alert, both Alert Enabled and Levels Enabled must be selected.

To turn on/off, change threshold settings, enable/disable or remove an alert on a single resource:

- 1. In the Alert Administration display, select an alert in the Active Alert Table and click Edit Index Levels. The Tabular Alert Administration display opens.
- 2. In the Tabular Alert Administration display, from the Index Type drop-down menu, select the Index type (options are populated by the type of alert you previously selected). For example, with the EMS Monitor, select PerServerQueue, PerServerTopic or PerServer. NOTE: If you select PerServerQueue or PerServerTopic, the alert settings are applied to the queue or topic on a single server.
- **3.** In the **Values** table, select the server to apply alert settings and click **Add**. In a few moments the server appears in the **Index Alert Settings** table.
- 4. In the Index Alert Settings table select the server.
- 5. In the Alert Settings panel (lower right), if needed, modify the Warning Level and Alarm Level settings.
- 6. In the Alert Settings panel, set the following as appropriate.

To turn on the alert for this index with the given thresholds:

Alert Enabled Select this option.

Levels Enabled Select this option.

To turn off the alert for only this index (global alert thresholds will no longer apply to this index):

Alert Enabled Deselect this option.

Levels Enabled Select this option.

To no longer evaluate this indexed alert and revert to global settings (or, optionally, Remove it if it is never to be used again):

Alert Enabled Not used.

Levels Enabled Deselect this option.

7. Click Save Settings. In a few moments the modifications are updated in the Index Alert Settings table.
Metrics Administration

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



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

RTView Internal Statistics	Select to open the RTView MBeans for Status and Timing Info display (in a separate window).			
RTView JMX	This table lists all JMX data objects.			
Query Statistics	Admin Query Key The dsString used for the data attachment to this data object.			
	Count The number of listeners for this data object. For example, graphical objects and function arguments.			
	Filter Objects The number of filtered data objects in this data object.			
	hasParent True if the data object is a filtered data object.			
	Last Exec Time The last time we queried for the metric associated with this data object.			

Time Range

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

Select or Enter Date and Time:
Restore to Now
Ok Apply Cancel

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

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

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

Trend GraphTraces the cumulative and maximum execution times, in seconds, for all Admin
Query Keys in the table.Max JMX Time The maximum execution time, in seconds, for all Admin Query
Keys in the table.Total JMX Time The cumulative execution time, in seconds, for all Admin
Query Keys in the table.

RTView Cache Tables

View data that RTView is capturing and maintaining. Drill down and view details of RTView Cache Tables. Use this data for debugging. This display is typically used for troubleshooting with Technical Support.

<			RT	View Cache	Tables	05-	Jul-2016 13:09	o 🔹 Data OK	+ 📀
DataServer: <de< th=""><th>efault></th><th></th><th>RT</th><th>View Cache</th><th>e Tables</th><th>Max Ro</th><th>ws:4000</th><th>History</th><th>Tables</th></de<>	efault>		RT	View Cache	e Tables	Max Ro	ws:4000	History	Tables
С	acheTable		TableTy	/pe	Rows	Colum	ins	Memory	A
JmxStatsTotals			current			1	4		441 _
OcBadCommunicat	tionNodes		current		14	0	6	14	,999 =
OcCacheServiceSt	ats		current		8	8	58	62	2,666
OcCacheServiceTo	itals		current			8	26	4	,441
OcCacheStats			current		17	2	80	206	6,148
OcCacheTotals			current		1	7	52	13	,406
OcClusterOverview			current			1	7		791
OcClusterStats			current			1	19	14	,103
OcExtendConnection	ons		current		11	2	30	68	,304
OcInvocationServic	eStats		current		6	3	60	62	2,252
OcInvocationServic	eTotals 👆		current			1	26	2	2,841
OcJmxConnection			current			2	7	1	,254 🖕
Oo ImvMamtData			ourront			1	1	1	754
			Oclr	vocationSe	rviceStats			Rows: 63	
TIME_STAMP	BackupCou	OwnedParti	OwnedParti	RefreshTim	RequestAve F	RequestMa Re	questPei R	equestPei Re	ques 🔺
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	85.5	109	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	7.5	15	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	16	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	15.5	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	24	32	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	31.5	47	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	39.5	63	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	16	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5 2016 1	8	16	0	0	The second secon
III									•

Choose a cache table from the upper table to see cached data.

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

DataServer Select a data server from the drop down menu.

Max Rows Enter the maximum number of rows to display in RTView Cache Tables.

History Select to include all defined history tables in RTView Cache Tables. Tables

RTView Cache Tables

This table lists and describes all defined RTView Cache Tables for your system. Cache tables gather Monitor data and are the source that populate the Monitor displays.

NOTE: When you click on a row in RTView Cache Tables a supplemental table will appear that gives more detail on the selected Cache Table.

CacheTable The name of the cache table.

TableTypeThe type of cache table:

	current	Current table which shows the current values for each index.	
	current_condensed	Current table with primary compaction configured.	
	history History table.		
	history_condensed	History table with primary compaction configured.	
Rows	Number of rows currently in the table.		
Columns	Number of columns currently in the table.		
Memory	Amount of space, in bytes, used by the table.		

About

This display shows Monitor version details and the data sources available to your system.

RTView(R) Display Server - Oracle Coherence Monitor
Version: 6400ALPHA
Configuration: OCM 6 4 0 0, 20160610, 000 21979-alpha, 115
Puild Number: 000 21979
Duid Number. 000.21373
Available Data Sources:
Alert
Cache Datasource
JMX
LOG4J
SQL XMI

CHAPTER 5 Report Generator

This section describes the Report Generator, and includes the following:

- "Overview" on page 177
- "Generating Reports from the Monitor" on page 177
- "Generating Reports From a Command Line" on page 178
- "Customizing Reports" on page 178

Overview

The Report Generator enables you to quickly see the state of any or all of your Monitor displays. The Monitor has two options for exporting reports to a printable PDF file from your display:

- Display Exports a screen capture of a display to a PDF file.
- Report Exports a screen capture of a display onto the first page of a PDF file, followed by at least one page for each table or object grid in the display. As many pages as are necessary to show all the data in each table or object grid are included in the report. This enables you to view all data in a table or object grid that you otherwise must use a scrollbar to see. If there are no tables or object grids in your display, you will only get a screen shot of the display.

You can customize your reports by simply editing the report configuration XML file. See "Customizing Reports" on page 178 for details.

You can generate reports from the Monitor application (see next.), or generate reports from a command line. See "Generating Reports From a Command Line" on page 178.

Generating Reports from the Monitor

Perform the following steps to start the Report Generator:

- 1. Right-click on the display and select Export.
- 2. Select Display or Report.

The Page Setup dialog opens.

3. Save the file to the directory of your choice.

NOTE: Performing **Export > Report / Display** from the file menu exports the display in the main window of the Monitor.

Generating Reports From a Command Line

Perform the following steps to start the Report Generator:

Windows

"Initialize a Command Prompt or Terminal Window", go to the **bin** directory and type: Win_RunOCMReportGen_DS.bat

UNIX/Linux

"Initialize a Command Prompt or Terminal Window", go to the directory where your displays reside and type: Unix_RunOCMReportGen_DS.sh

Customizing Reports

The Report Generator uses a report configuration file to create reports. You customize reports simply by editing the report configuration XML file. The report configuration file specifies which displays to generate a report for, as well as the report format, the report layout, the substitutions to use for the report, and the period of time to cover in the report.

The default report configuration file, ts_report.xml, is used if the script to start the Report Generator is executed without a -report argument.

The report configuration file can contain multiple reports. Each report is generated by a different display and output to a single PDF file.

This section includes:

- "The Report Generator Configuration File" on page 178
- "Report Generator Configuration File Parameters" on page 179
- "Sample Report Configuration File" on page 181

The Report Generator Configuration File

The **ts_report.xml** report configuration file generates a report for each type of display. You can customize your reports by simply adding or removing displays from the file as needed, and modifying any other parameters, as described in Report Generator Configuration File Parameters.

To edit the **ts_report.xml** report configuration file:

- 1. Navigate to the ocm\lib directory and locate the gmsjocm.jar file.
- 2. Extract the ts_report.xml file from gmsjocm.jar into the directory from which you run the Report Generator.

The following is an edited example of the **ts_report.xml** report configuration file. In this example, we specify to create a report for a single display, t**s_cache_summary.rtv**, to name both the report and the PDF file **newproduct_january2011**, to include a cover page, **ts_report_cover.rtv**, and to include a screen capture of the display followed by all of the data contained in the display (the **Report** option):

<rtvreport xmlns="www.sl.com" version="1.0">

```
<report name = "newproduct_january2011">
        <output filename = "newproduct_january2011"
               append_timestamp = "false"
               report_or_display = "report">
       </output>
       <section display = "ts_report_cover.rtv"
               subs = ""
               timeout = "4">
               <page_setup orientation = "landscape"/>
       </section>
<section display="ts_cache_summary.rtv"
               subs="$serviceName:OnLineStoreCache $cache:UserSession"
               timeout="2">
               <page_setup orientation="landscape" />
       </section>
</report>
</rtvreport>
The report configuration file must start with the following:
<?xml version="1.0" ?>
<rtvreport xmlns="www.sl.com" version="1.0">
And end with the following:
</rtvreport>
```

Report Generator Configuration File Parameters

The following tags are supported by the Report Generator configuration file:

Tag Description

report Specify a report definition. Your configuration file may contain multiple reports. **Attributes:**

- name The name of the report.
- **title** The title for the report. This tag is optional. If specified, this title will be used instead of the display name in the header of each page in the report.

section Specify the display and page setup information for this section. You may have multiple sections in a single report.

Available Attributes:

- **timeout** The time to wait (in seconds) for data sources to update before producing a report for this section.
- display The name of the Monitor display to use for this section.
- **subs** The substitution values to apply to the Monitor display. Specify initial substitutions for this display. Substitutions are optional and must use the following syntax: **\$subname:subvalue\$subname2:subvalue2**

If a substitution value contains a single quote, it must be escaped using a **/**: **\$filter:Plant=/'Dallas/'**

If a substitution value contains a space, it must be enclosed in single quotes. Do not escape these single quotes: **\$subname:subvalue \$subname2:'sub value 2'**

A substitution string cannot contain the following characters: :|.tabspace,;=<>'''& / \ { } [] ()

\$reportDuration - Use the \$reportDuration substitution to specify the time range for data displayed in charts. By default, the \$reportDuration substitution is set to **1 Hour** (to display one hour of data). Valid entries for the \$reportDuration substitution are:

- All Data
- 2 Mins
- 5 Mins
- 20 Mins
- 1 Hour
- 4 Hours
- 8 Hours
- 24 Hours
- Last 2 Days
- Last 7 Days

Example:

subs = "\$reportDuration:1 Hour"

output

- page_setup Specify the page setup options for this section.
 - Available Attributes:
 - orientation The values are portrait or landscape.
 - margin_left Sets the left margin. The values are in inches.
 - margin_right Sets the right margin. The values are in inches.
 - margin_top Sets the top margin. The values are in inches.
 - margin_bottom Sets the bottom margin. The values are in inches.

Specify the output options for this report.

Available Attributes:

- **filename** The base name of the generated report file. A PDF suffix is added automatically.
- append_timestamp The values are none, false, true, date_time or date_only. If date_time, append timestamp to file basename in the form yearmonthday_hourminsec, where hour is on a 24 hour clock. For example, a report with a filename of myreport, output on September 15, 2006 at 8:45am would be named myreport_20060915_084500.pdf. If date_only, append the timestamp in the form yearmonthday. For example, a report with a filename of myreport, output on September 15, 2006 at any time would be named myreport_20060915.pdf.
- **report_or_display** The kind of report to generate. The values are display or report. **display** exports a screen capture of a display. **report** exports a screen capture of a display onto the first page of the report, followed by at least one page for each table or object grid in the display. As many pages as are necessary to show all the data in each table or object grid are included in the report. This enables you to view all data in a table or object grid that you otherwise must use a scrollbar to see. If there are no tables or object grids in your display, you will only get a screen shot of the display.

Sample Report Configuration File

```
<rtvreport xmlns="www.sl.com" version="1.0">
<report name = "ocm_report1">
        <output filename = "ocm_report1"
                append_timestamp = "false"
               report_or_display = "display">
        </output>
        <section display = "ts_report_cover.rtv"
                subs = ""
               timeout = "4">
                <page_setup orientation = "landscape"/>
        </section>
        <section display = "ts_grid_overview.rtv"
                subs = ""
               timeout = "2">
                <page_setup orientation = "landscape"/>
        </section>
```

```
<section display = "ts_grid_overview_2.rtv"
               subs = ""
               timeout = "2">
                <page_setup orientation = "landscape"/>
        </section>
        <section display = "ts_extend_detail_fx.rtv"
               subs = ""
               timeout = "2">
                <page_setup orientation = "landscape"/>
        </section>
       <section display = "ts_grid_health_fx.rtv"
               subs = ""
               timeout = "2">
                <page_setup orientation = "landscape"/>
        </section>
        <section display = "ts_grid_status.rtv"
               subs = ""
               timeout = "2">
                <page_setup orientation = "landscape"/>
        </section>
</report>
</rtvreport>
```

CHAPTER 6 Additional Configurations

This section provides additional configuration information for the RTView Oracle Coherence Monitor, including:

- "Database Configuration," next For details about database requirements and configuration.
- "Configuring Servlets" on page 188 For details about configuring servlets.
- "Configuring Desktop Clients" on page 190 For details about configuring desktop clients.
- "Configure Alert Notification" on page 192 For details about configuring alert notification.
- "Log File Management on UNIX/Linux Systems" on page 197 For details about log file management on UNIX/Linux systems.
- "Role Based Security" on page 197 For details about configuring role based security.
- "Single Sign-On" on page 197- For details about configuring single sign-on.
- "High Availability Configurations" on page 198 For details about configuring high availability.
- "Custom Web Application Queries" on page 199 For details about the rtvquery servlet.
- "Invocation-only Configuration" on page 214 For details about configuring clusters that do not have storage nodes or caches.

Database Configuration

This section describes how to replace the default HSQLDB database with another database that is accessible via JDBC or ODBC. While the HSQLDB database may be sufficient for evaluation and testing, it is not recommended for use in production deployments.

NOTE: DB2 database page size must be a minimum of 8k.

You configure the database using property values in conjunction with property filters (which ensure that properties are only used to configure the appropriate applications). For details about properties and property filters, see "Monitor Properties" on page 227.

The Monitor requires a database that provides tables for the following information:

- Alert Thresholds: Each Monitor instance requires a database table that contains alert thresholds and enable/disable flags. The values in this table are used by the alert engine at runtime. If this table is not available, alerts will not function.
- Historical Data: If you use the Monitor Historian, it requires database tables in which to write data.

When using HSQLDB, the database capacity and space requirements described below should be applied to RAM for the system hosting HSQLDB because its tables are held in memory, not on disk.

The default Coherence monitoring configuration uses two database connections, one for alert threshold information (ALERTDEFS), and one for historical data storage via the Historian (RTVHISTORY). The default alert threshold table name is ALERTDEFS. The default history table names are OCM_CLUSTERSTATS, OCM_CACHETOTALS, OCM_STORAGETOTALS, OCM_CACHESERVICETOTALS, OCM_CACHESERVICETOTALS, OCM_NODESTATS. Substitutions are used for the database connection names and table names which are defined in the **rtview.properties** file, located in the **ocmon/projects/mysample/** directory.

Important Note: The default Historian configuration saves metrics only for the above named tables.

This section includes:

- "JDBC" on page 184
- "ODBC" on page 185
- "Database Capacity and Purging" on page 185
- "Defining Database Connections" on page 185
- "Modifying Database Connections" on page 186
- "Defining Database Tables" on page 187

JDBC

JDBC driver file for selected database.

Example: ojdbc14.jar

The location of the JDBC driver must be available on the JAVA classpath. The path to the driver can be added using the **sl.rtview.cp** property in the **rtview.properties** file.

JDBC access parameters such as user name, password, driver class and driver URL.

Example:

User Name: kermit

Password: frog

Driver: oracle.jdbc.driver.OracleDriver

URL: jdbc:oracle:thin:@dbhost:1521:rtv

collector.sl.rtview.sql.sqldbsqldb=ALERTDEFS kermit frog jdbc:oracle:thin:@myhost:1521:rtv oracle.jdbc.driver.OracleDriver - false false

collector.sl.rtview.sql.sqldbsqldb=RTVHISTORY kermit frog jdbc:oracle:thin:@myhost:1521:rtv oracle.jdbc.driver.OracleDriver - false false

Verify the following entries are appropriate for your environment:

- collector.sl.rtview.sql.sqldb= ALERTDEFS ...
- collector.sl.rtview.sql.sqldb=RTVHISTORY ...

NOTE: Examples are provided in the **database.properties** file, located in the **rtvapm/ common/dbconfig** directory

ODBC

In order for RTView to communicate with your database using an ODBC-JDBC bridge, you must have an ODBC driver for your database. Most databases that run on Microsoft Windows come standard with an ODBC driver. You must also register your database with ODBC before accessing it from RTView. The name specified for the ODBC data source name during the ODBC driver setup must be RTVHISTORY for the history table connection and ALERTDEFS for the alert parameters table connection.

Database Capacity and Purging

The Monitor Historian writes average values to the database every 5 minutes by default.

At the end of each interval, the last 5 minutes of real time data in the Data Server for the Cluster Stats, CacheTotals, Storage Totals, Cache Service Totals, Invocation Service Totals, NodeStats and Node Totals is averaged and written to the Historian database. One CacheTotals row is written for each cache and tier combination, and one NodeStats row is written for each cluster node. Each column value is an average for the most recent compaction interval.

Database capacity for the CACHETOTALS and NODESTATS tables can be estimated with the following formula:

Bytes/day written to historian tables = ((A * 420) + (B * 154)) * 288

Where:

A = the number of cache/tier combinations.

(one cache with both front and back tiers counts as 2 for database capacity estimates)

 \mathbf{B} = the number of cluster nodes.

Table purging, archiving, displacement or additional compaction is not provided in the default Historian configuration. Optionally, you can provide procedures or scripts to keep the database within capacity limits. If licensed, the Advanced Historian features for table displacement can be used.

Defining Database Connections

sql template files are provided to assist you in creating the correct tables in the database of your choice. The sql templates for the alert tables are in the **rtvapm\common\dbconfig** directory. The sql templates for the history tables are in the **rtvapm\ocmon\dbconfig** directory.

The database connections (referred to as SQL connections) used by the Monitor are defined in the **rtview.properties** file, located in the project directory (**mysample**):

HISTORIAN PROPERTIES

#

#

Define the RTVHISTORY DB

Note that hsqldb can have poor performance with big amounts of data.

We suggest to set an alternate database in production environments.

To specify your database:

- run the builder in a separate directory

- select Tools/Options/SQL; enter your database options

- select Save/No. The file OPTIONS.ini will be created

- find the line containing your database options and copy it here:

#

collector.sl.rtview.sql.sqldb=RTVHISTORY sa - jdbc:hsqldb:hsql://localhost:9103/rtvhistory org.hsqldb.jdbcDriver - false true

#

historian.sl.rtview.historian.driver=org.hsqldb.jdbcDriver

historian.sl.rtview.historian.url=jdbc:hsqldb:hsql://localhost:9103/rtvhistory

historian.sl.rtview.historian.username=sa

historian.sl.rtview.historian.password=

Modifying Database Connections

If you replace HSQLDB with another database, the database connections in the **rtview.properties** file must be modified for the database. Examples are provided in the **database.properties** file, located in the **rtvapm/common/dbconfig** directory. Use the appropriate commented out properties in the **database.properties** file to replace the default HSQLDB entries in the **rtview.properties** file.

Also, keep the same propfilter value (for example, **maincollector**, **collector**, **historian**) in the **rtview.properties** file when you modify database connection definitions. For details about properties, see "Monitor Properties" on page 227.

Encrypting Passwords

If you need to provide encrypted passwords (rather than expose server password names in a clear text file), do the following. If not, skip this step.

To provide encrypted passwords, use the "encode_string" utility:

"Initialize a Command Prompt or Terminal Window", and execute the following script where mypassword is your plain text password (this password will be encrypted).

encode_string sql mypassword

You then receive an encrypted password for copying and pasting into the password field. For example:

encrypted value: 013430135501346013310134901353013450134801334

Defining Database Tables

SQL template files are provided to assist you in creating the correct tables in the database of your choice. The sql templates for the alert tables are in the **rtvapm\common\dbconfig** directory. The sql templates for the history tables are in the **rtvapm\ocmon\dbconfig** directory. The standard SQL syntax is provided for each database, but requirements can vary depending on database configuration. Consult with your database administrator, or contact SL Technical Support, if you require assistance.

To define the database tables using the provided template files, perform the following steps.

- In the dbconfig directory, open the createtables_<tabletype>_<database>.sql file, where <database> corresponds to your type of database (for example, create_ocmon_history_tables_mysql.sql). Make the following edits as they apply to your database:
- **myinstance**: Replace with the database instance where you wish to create the table.
- mydbuser: Replace with the database user whose instance the tables will be created under.
- 2. Use the edited .sql file to create the database tables. The effective method for doing this will vary depending on your database and how that 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 createtables_<tabletype>_<database>.sql file into this tool.
- Import Interface Some database applications allow you to specify a .sql file containing SQL commands. You can use the createtables_<tabletype>_<database>.sql file for this purpose. In some cases it may be necessary to save the statements in that file into individual files.
- Third Party Application Third party tools are available for connecting to a variety of databases (RazorSQL, SQLMaestro, Toad, etc). 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.
- 3. Configure the Monitor to use the database tables you just created by editing the table names in your rtview.properties file, located in the ocmon\projects\mysample directory. You can use the default table names or use your own table names. The following is an example from a sample rtview.properties file:

```
#
#
Default OCM Cache History Tables
#
sl.rtview.sub=$OCMCLUSTERSTATS_TABLE:OCM_CLUSTER TATS
sl.rtview.sub=$OCMCACHETOTALS_TABLE:OCM_CACHETOTALS
sl.rtview.sub=$OCMSTORAGETOTALS_TABLE:OCM_STORAGETOTALS
sl.rtview.sub=$OCMCACHESERVICETOTALS_TABLE:OCM_CACHESERVICETOTALS
sl.rtview.sub=$OCMCACHESERVICETOTALS_TABLE:OCM_CACHESERVICETOTALS
```

```
sl.rtview.sub=$OCMINVOCATIONSERVICETOTALS_TABLE:OCM_INVOCATIONSERVICETOTAL
Ssl.rtview.sub=$OCMNODETOTALS_TABLE:OCM_NODETOTALS
sl.rtview.sub=$OCMNODESTATS_TABLE:OCM_NODESTATS
```

```
4. Save the file.
```

Configuring Servlets

A servlet container such as Tomcat is required for a Monitor Web deployment. A Monitor Web deployment makes use of two servlets that enable the servlet container to obtain data from Monitor server processes:

- Display Servlet: A Display Servlet allows the servlet container to communicate with a Monitor Display Servlet to satisfy the servlet requests for data used in a Web page. A Display Servlet requires the IP or host name of the system where the Monitor Display Server is running, and the port on which the Display Server is listening. The default Display Servlet, ocmon.war, uses localhost and Port 3361.
- Data Servlet: The data servlet allows a desktop client to connect to a Monitor Data Server via the servlet container. This technique allows a remote desktop client to obtain data from a Monitor Data Server when firewalls or network routing prevent a direct connection. A Data Servlet requires the IP or host name of the system where the Monitor Data Server is running, and the port on which the Data Server is listening. The default Data Servlet, ocmon_rtdata.war, uses localhost and Port 3381.

Monitor Displays and Data Servlets are configured using property files and built using batch files and shells scripts provided for that purpose which are found in the **projects/mysample** subdirectory of the Monitor installation.

NOTE: An initial or test deployment can use the WAR files in the **ocmon/webapps** directory, which are based on the default ports, as long as the default ports are available.

The display servlet is configured by the **rtvdisplay.properties** file and uses the **ocmon.war** file.

The data servlet is configured by the **servlet.properties** file and uses the **ocmon_rtvdata.war** file.

To reconfigure a servlet you edit the appropriate **.properties** file for the servlet, set the ServiceHost and ServicePort properties as appropriate for your deployment, update the **.war** files by running the **update_wars** script appropriate for your platform, rename the updated **.war** files as appropriate for your deployment and deploy the (renamed) **.war** files to your web server.

Example: The Monitor for Coherence cluster DevClusterA is deployed with Tomcat running on WebHost1. The Monitor Data Server and Display Server are running on DevHost1.

The following instructions describe how to reconfigure a servlet. Rather than using the default ports, the instruction use the following settings as an example:

- The Display Server is uses port **3367** instead of the default port 3361.
- The Data Server is using port **3387** instead of the default port 3381.

Windows

- 1. Open the **rtvdisplay.properties** file, located the **ocmon\project\mysample** directory, in a text editor and edit as follows to configure the display servlet:
- Change DisplayServerHost=localhost to DisplayServerHost=DevHost1
- Change DisplayServerPort=3361 to DisplayServerPort=3367
- 2. Save the file.

- **3.** Open the **servlet.properties** file, located the **ocmon\project\mysample** directory, in a text editor and edit as follows to configure the data servlet:
- Change ServiceHost=localhost to ServiceHost=DevHost1
- Change ServicePort=3381 to ServicePort=3387
- 4. Save the file.
- "Initialize a Command Prompt or Terminal Window", navigate to the ocmon\project\mysample directory and execute the following script to create the new .war files:

Example: update_wars.bat

- 6. Rename the .war files if required.
- Deploy the servlets by transferring them to Tomcat's webapps directory on WebHost1 (or by using the Tomcat Web Application Manager screen to deploy the .war files, if appropriate).

UNIX/Linux

- 1. "Initialize a Command Prompt or Terminal Window", and navigate to your project directory. Example: cd ocmon/project/mysample
- 2. Edit rtvdisplay.properties with the following changes to configure the display servlet:
- Change DisplayServerHost=localhost to DisplayServerHost=DevHost1
- Change DisplayServerPort=3361 to DisplayServerPort=3367
- 3. Save the file.
- 4. Open the **servlet.properties** file, located the **ocmon/project/mysample** directory, in a text editor and edit as follows to configure the data servlet:
- Change ServiceHost=localhost to ServiceHost=DevHost1
- Change ServicePort=3381 to ServicePort=3387
- 5. Save the file.
- 6. Execute the following script to create the new .war files.
 - Example: ./update_wars.sh
- 7. Rename the .war files if required.
- Deploy the servlets by transferring the .war files to Tomcat's webapps directory on WebHost1 (or by using the Tomcat Web Application Manager screen to deploy the .war files, if appropriate).

Configuring Desktop Clients

This section is optional.

Most Monitor installations use the full Web deployment which does not require desktop clients. However, the Monitor and RTView Builder desktop clients are sometimes useful with small test or development clusters. There are also times when the Builder is required to perform advanced configuration tasks.

The **startrtv** script referenced in these instructions starts processes in a Monitor configuration as specified in the "rtvservers.dat" configuration file. For details about **startrtv** and **rtvserver.dat**, see "Scripts" on page 217.

There are three options for using the Monitor and Builder desktop clients:

- Standalone Operation," next
- "Data Server Connection" on page 190
- "Data Servlet Connection" on page 191

Standalone Operation

In a standalone setup, the Monitor or Builder connects to the cluster either directly or via JMX.

- 1. Configure as described in "Configuring a JMX Connection" on page 16 or "Configuring a Direct Connection" on page 19. For details about choosing a data connection method, see "Configure a Data Connection" on page 15.
- 2. Start the desktop clients:

Windows:

"Initialize a Command Prompt or Terminal Window", and change directory (cd) to the projects directory.

Example: cd ocmon\projects\mysample

Start the Viewer by typing: start_rtv viewer

For details about start_rtv, see "Scripts" on page 217.

UNIX/Linux:

"Initialize a Command Prompt or Terminal Window", and change directory (cd) to the projects directory.

Example: cd ocmon/projects/mysample

Start the Viewer by typing: start_rtv.sh viewer

For details about start_rtv, see "Scripts" on page 217.

Data Server Connection

In a Data Server connection setup the Monitor or Builder connects to a Monitor Data Server. The connection is defined in the **rtview.properties** file, located in the **ocmon/projects/ mysample** subdirectory of the Monitor installation. 1. Open the **rtview.properties** file in a text editor. Locate the line that defines the **dataclient.sl.rtview.dataserver** property:

dataclient.sl.rtview.dataserver=//<dataserver_host>:<dataserver_port>

- **2.** Edit as follows:
- Change dataserver_host to the IP address or host name of the system hosting the Data Server.
- Change the dataserver_port to the port that the Data Server is listening on.
 Example: dataclient.sl.rtview.dataserver=//localhost:3381
- **3.** Use the following scripts to start the desktop clients:

Windows:

"Initialize a Command Prompt or Terminal Window", change directory (cd) to the projects directory.

Example: cd ocmon\projects\mysample

Start the Viewer by typing: start_rtv viewer_ds

UNIX/Linux:

"Initialize a Command Prompt or Terminal Window", and change directory (cd) to the projects directory.

Example: cd ocmon/projects/mysample

Start the Viewer by typing: start_rtv.sh viewer_ds

For details about **start_rtv**, see "Scripts" on page 217.

Data Servlet Connection

In a Data Servlet connection setup the Monitor or Builder connect to a Monitor Data Server using a Data Servlet hosted by a servlet container such as Tomcat. The connection is defined in the **rtview.properties** file, located in the **ocmon/projects/mysample** subdirectory of the Monitor installation.

1. Open the **rtview.properties** file in a text editor. Locate the line that defines the **dataclient.sl.rtview.dataserver** property:

dataclient.sl.rtview.dataserver=-remote:http://<http_host>:<http_port>/ ocmon_rtvdata

2. Make the following changes:

- Change http_host to the IP address or host name of the system hosting the servlet container.
- Change the **http_port** as needed to the port that the servlet container is listening on.
- Change ocmon_rtvdata to the name of the Data Servlet that connects to the Monitor Data Server (for example, DevClusterA_rtvdata).

Example: dataclient.sl.rtview.dataserver=remote:http://some_host:8080/ ocmon_rtvdata

3. Start the desktop clients:

Windows:

"Initialize a Command Prompt or Terminal Window", change directory (**cd**) to the projects directory.

Example: cd ocmon\projects\mysample

Start the Viewer by typing: start_rtv viewer_ds

For details about **start_rtv**, see "Scripts" on page 217.

UNIX/Linux:

"Initialize a Command Prompt or Terminal Window", change directory (cd) to the projects directory.

Example: cd ocmon/projects/mysample

Start the Viewer by typing: start_rtv.sh viewer_ds

For details about **start_rtv**, see "Scripts" on page 217.

Configure Alert Notification

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

By default, alerts execute a **.bat** script. 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).

There are two options for configuring Monitor alert notification:

"Using a Batch File or Shell Script" on page 192

This technique requires switching to an OS-specific set of alert definitions that execute the appropriate file type. Windows and UNIX alert definition files are provided with the Monitor. A sample batch file and a sample shell script are also provided which are customized as needed.

• "Using the Java Command Handler" on page 195

The Java source for the Monitor Java command handler is provided to facilitate customization.

When you uncomment the line in the script that prints alert data to standard output, both the scripts and the Java command handler output alert information to standard output. The alert output appears in the Data Server log file, or in the command window or shell in which the Data Server was started. 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 #####

Using a Batch File or Shell Script

A sample batch file, my_alert_actions.bat, and a sample shell script,

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

- "Windows Batch File," next
- "UNIX/Linux Shell Script" on page 193

Windows Batch File

- 1. Copy the **my_alert_actions.bat** file, located in the **common/bin** directory, into your project directory.
- 2. Open the **rtvapm.properties** file, located in the **rtvapm\common\conf** directory, in a text editor.
- **3.** Verify the following properties are set as follows:

command to execute for new alerts

sl.rtview.alert.notifiercommandnew=system cmd '\$alertActionScript.\$scriptEnding \$domainName +\$alertName+ "+\$alertIndex+" +\$alertID+ +\$alertSeverity+ +\$alertText'

command to execute on the first severity change

sl.rtview.alert.notifiercommandfirstsevchange=system cmd '\$alertActionScript.\$scriptEnding \$domainName +\$alertName+ "+\$alertIndex+" +\$alertID+ +\$alertSeverity+ +\$alertText'

- **4.** Open the **my_alert_actions.bat** file, located in your project directory, and uncomment the echo line (near the end of the file) to print alert information to standard output. Or, you can modify the script to execute an automated action (such as sending an email alert).
- 5. Restart the Data Server.

UNIX/Linux Shell Script

- 1. Copy the **my_alert_actions.sh** file, located in the **common/bin** directory, into your project directory.
- 2. Open the rtvapm.properties file, located in the rtvapm/common/conf/ directory, in a text editor.
- 3. Verify the following properties are set as follows:

command to execute for new alerts

sl.rtview.alert.notifiercommandnew=system cmd '\$alertActionScript.\$scriptEnding \$domainName +\$alertName+ "+\$alertIndex+" +\$alertID+ +\$alertSeverity+ +\$alertText'

command to execute on the first severity change

sl.rtview.alert.notifiercommandfirstsevchange=system cmd '\$alertActionScript.\$scriptEnding \$domainName +\$alertName+ "+\$alertIndex+" +\$alertID+ +\$alertSeverity+ +\$alertText'

4. Open the **sample.properties** file, located in your project directory, and copy/paste the following into it:

sl.rtview.cmd_line=-sub:\$scriptEnding:bat

sl.rtview.cmd_line=-sub:\$alertActionScript:my_alert_actions

- 5. Change the bat suffix to sh and my_alert_actions to ./my_alert_actions.
- 6. Save the sample.properties file.
- 7. Open the **my_alert_actions.sh** file, located in your project directory, and uncomment the echo line (near the end of the file) to print alert information to standard out. Or, you can modify the script to execute an automated action (such as sending an email alert).
- 8. Restart the Data Server.

Batch File or Shell Script Substitutions

The default **my_alert_actions** scripts use the substitutions described in the table below. When you customize the script, you can use a use substitution to get any of the columns in the alert table. To do this, modify the **sl.rtview.alert.notifiercommandnew** and **sl.rtview.alert.notifiercommandfirstsevchange** properties from Step 3 (above) to replace the default substitutions with the substitutions you want to use. You must make corresponding modifications to your script to use modified substitution values.

The substitution names map to the names of the columns in the alert table. Convert the column name to camel case and if it does not start with Alert, prepend alert to it. For example, to use the value of the **Alert Name** column, use **\$alertName**. To use the value of the **ID** column, use **\$alertID**. To use the value of the **Row Update Time** column, use **\$alertRowUpdateTime**. The following table contains the substitutions used by the default **my_alert_actions** scripts:

Argument	Description	Values
\$alert1d	This substitution specifies the unique ID for the alert. For example: alertId = 1004	Text or Numeric
\$alertIndex	This substitution specifies which source triggered the alert. With tabular objects, the first column of data is typically the Index column. The value in the Index column is a name that uniquely identifies each table row. The alertIndex uses the Index column name. For example, if the CapactityLimitAllCaches alert is configured to monitor all of your caches, and to trigger when any of the caches exceed the specified capacity threshold, the alertIndex indicates specifically which cache triggered the alert. With scalar objects, which do not have a table and therefore do not have a column (the useTabularDataFlag property is False), the alertIndex is blank. For example: alertIndex = MyCache01	Text or Numeric
\$alertName =	This substitution specifies the name of the alert. For example: alertName = CapacityLimitAllCaches	Values vary.
\$alertSeverity	 This substitution specifies the severity level of the alert. O: The alert limit has not been exceeded therefore the alert is not activated. 1: The alert warning limit has been exceeded. 2: The alert alarm limit has been exceeded. For example: alertSeverity = 1 	Numeric

\$alertText	This substitution specifies the text that is displayed when the alert executes.	Text
	For example: alertText = High Warning Limit exceeded, current value: 0.9452 limit: 0.8	
\$alertTime	This value is the time the alert was initially generated.	Text

Using the Java Command Handler

- 1. Open the custom_handlers.properties file, located in the rtvapm\common\conf\ directory, in a text editor.
- 2. Locate the following text in the custom_handlers.properties file and select it for copying:

#sl.rtview.alert.notifiercommandnew=system cust
'my_alert_notification.\$domainName.\$alertNotifyType.\$alertNotifyCol' \$alertNotifyTable

#sl.rtview.alert.notifiercommandfirstsevchange=system cust
'my_alert_notification.\$domainName.\$alertNotifyType.\$alertNotifyCol' \$alertNotifyTable

- **3.** Open the **sample.properties** file, located in your project directory, and paste the selected text into it.
- **4.** In the **sample.properties** file, uncomment the lines you just pasted to enable the Java command handler.
- **5.** Verify that the **rtvapm_custom.jar** file is built per the Customizing the Java Command Handler instructions.
- 6. Open the custom_handler.properties file and verify that the following line is the correct path to the rtvapm_custom.jar. If it is not the correct path, copy the line into the sample.properties file in your project directory and edit the path as needed.

sl.rtview.cp=./custom/lib/rtvapm_custom.jar

- 7. Save the sample.properties file.
- 8. Restart the Data Server using the following command line argument:

Windows -properties:%RTVAPM_HOME%/common/conf/custom_handlers UNIX

-properties:\$RTVAPM_HOME/common/conf/custom_handlers

Customizing the Java Command Handler

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

1. Open the RtvApmCommandHandler.java file.

- 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.
- Compile RtvApmCommandHandler.java and rebuild rtvapm_custom.jar using the supplied script (make_classes.bat or make_classes.sh) in the \projects\sample\custom\src directory.
- 5. Restart the Data Server using the following command line argument:

Windows

-properties:%RTVAPM_HOME%/common/conf/custom_handlers UNIX

-properties:\$RTVAPM_HOME/common/conf/custom_handlers

Java Command Handler Substitutions

When you customize the Java Command Handler, there is no need to modify the **sl.rtview.alert.notifiercommandnew** and

sl.rtview.alert.notifiercommandfirstsevchange properties in the custom_handlers.properties file. The entire alert table row is passed into the Java Command Handler for each alert that notifies so that all information regarding those alerts is available. The following substitutions are used:

Argument Description

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

Additional Optional Properties

- sl.rtview.alert.notifiercommandcleared Set this to the command to execute when an alert is cleared. By default, no command is configured. To execute a script, copy the notifiercommandnew line and replace \$alertActionScript with the name of the script you want to execute. To execute a custom Java command, see the example in common\conf\custom_handlers.properties.
- sl.rtview.alert.notifiercommandchanged Set this to the command to execute when a column in the alert table changes. To execute a script, copy the notifiercommandnew line and replace \$alertActionScript with the name of the script you want to execute. To execute a custom Java command, see the example in common confocustom bandlers properties. This must be used in conjunction with

common\conf\custom_handlers.properties. This must be used in conjunction with
the sl.rtview.alert.notifiercolumns property

sl.rtview.notifiercolumns - Set this to the name of one or more columns to execute the sl.rtview.alert.notifiercommandchanged notification when they change. For multiple columns, use a semi-colon delimited list. Note that this should be limited to the minimum number of necessary columns, preferably less than 5, as a large number of columns increases the persistence load on the central alert server.

Notification Persistence

To prevent duplication and missed notifications after restart or failover, you must configure the Data Server for alert persistence.

Log File Management on UNIX/Linux Systems

Most UNIX/Linux systems provide a utility for managing the log files produced by RTView applications.

For example:

- Solaris logadm
- Linux logrotate
- Macintosh daily script for OS version 10.5.1 and earlier, and newsyslog for OS version 10.5.1 and later

Consult your system administrator for details.

Role Based Security

Please refer to the following for details about role based security:

- For an introduction to role-based security, see: http://sldownloads.sl.com/docs/rtview/ 57/user/security_intro.html.
- For details about how to login, see: http://sldownloads.sl.com/docs/rtview/57/user/ security_login.html.
- For details about how to configure role-based security, see: http://sldownloads.sl.com/ docs/rtview/57/user/security_config.html.

NOTE: The Monitor **users.xml** and **roles.xml** files are in the Monitor JAR file gmsjocm.jar from which they must be extracted for customization.

Single Sign-On

Please refer to the following for details about single sign-on:

- For details about how to setup single sign-on, see the Setup Login section, located at: http://sldownloads.sl.com/docs/rtview/57/user/security_login.html.
- For details about Display Server single sign-on, see: http://sldownloads.sl.com/docs/ rtview/57/user/ds_single_sign.html.

 For details about Display Builder single sign-on, see Release Note E11674, located at: http://sldownloads.sl.com/docs/rtview/57/rn.html.

High Availability Configurations

High Availability (HA) mitigates single point of failure within your Monitor deployment 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.

A High Availability Data Server configuration is available for the Monitor. The **project/sample** directory provides an example of a HA Data Server with the Monitor. The example assumes the availability of two machines, PRIMARYHOST and BACKUPHOST, defined by environment variables of the same name.

The Monitor HA is configured by using the primary and backup configurations in "rtvservers.dat", with **start_rtv** in place of the default configuration.

The primary Data Server is run on PRIMARYHOST; the backup Data Server is run on BACKUPHOST; the other Monitor applications failover between the Data Servers as appropriate. Assuming the environment variables PRIMARYHOST and BACKUPHOST are set correctly, Monitor components on the primary machine are started as normal using the primary configuration (instead of the default configuration) with the **start_rtv** command. The backup Monitor Data Server on the backup machine is started using the backup configuration with the **start_rtv** command.

Start the primary Monitor components on the primary machine using the primary configuration with the **start_rtv** command. For example:

Windows

start_rtv primary

UNIX

start_rtv.sh primary

Start the backup Monitor Data Server on the backup machine using the backup configuration with the **start_rtv** command. For example:

Windows

start_rtv backup

UNIX

start_rtv.sh backup

The appropriate property files and property filters for the primary and backup Data Servers are defined in the rtvservers.dat file under the primary and backup configurations. The property values controlling HA, used by the Monitor Data Servers, are defined in the ha.properties file.

Custom Web Application Queries

The Data Server includes a rtvquery servlet that you can optionally install and run on your application server. The rtvquery servlet enables custom browser-based applications to query information from caches in the Data Server. The rtvquery servlet is useful for obtaining internal data from the Monitor--data that the Monitor collects and does not include in displays. You can then use the data in your own custom application. The data is returned as a response to a HTTP request, which allows the data to be incorporated into any client capable of dealing with RESTful service.

rtvquery Servlet

The rtvquery servlet communicates with the Data Server via socket. If you are using multiple Data Servers, you must configure and install a rtvquery servlet for each Data Server. If you want to install multiple rtvquery servlets on the same application server, each must have a unique name.

The rtvquery servlet is intended to display tables from the Data Server at relatively low data volumes. For high data volume displays, the RTView Viewer application or Thin Client is recommended.

cachedocs Directory

Refer to the **index.html** documentation, located in the **cachedocs** directory, for information about target caches for queries.

For rtvquery servlet examples, see the **rtvapm\rtview\custom\rtvquery-samples** directory.

To use the rtvquery servlet you must configure and install it. For details, see the following:

- "Configure rtvquery Servlet," next: This section describes how to configure and install the rtvquery servlet on your application server (for example, Tomcat). The rtvquery servlet comes with the Data Server. The files for configuration and installation the rtvquery servlet are located in the servlets\rtvquery directory
- "Install rtvquery Servlet" on page 200
- "Deploy Multiple rtvquery Servlets" on page 200: This section is intended for users with an application server (for example, Tomcat) and a standard working knowledge of servlet deployment.
- "Using rtvquery Servlet" on page 201: This section provides information you need to obtain the internal data from the Monitor using the rtvquery servlet.

Configure rtvquery Servlet

The rtquery properties file, located in the **servlets\rtvquery** directory, contains properties that determine servlet behavior. Edit this file to change any of the properties from their default values. The most important properties are **DataserverHost** and **DataServerPort**, which specify the hostname and port number of the Data Server to which the servlet is to connect. The default values are **localhost** and **3278**, respectively (**localhost** indicates that the servlet and Data Server are running on the same machine). See the comments in the **rtvquery.properties** file for a description of the other properties and their default values.

NOTE: If you have already installed the servlet on your application server, you can edit the **rtvquery.properties** file on your application server. If you edit this file in **servlets\rtvquery** you must remake and reinstall **rtvquery.war**, described in Installing rtvquery Servlet (next).

After modifying the **rtvquery.properties** file you may need to restart your application server.

The **web.xml** file should be edited only if it is necessary to change the servlet authentication. By default, authentication is disabled, so any client can submit queries to the servlet. To enable authentication edit the **web.xml** file. Uncomment the authentication section, located at the end of the file, then remake and redeploy the **rtvquery.war** file. When authentication is enabled, the browser prompts the user for login information when the first request for a session is sent to the servlet. The user must enter a username and password that are valid for the application server (for example, Tomcat). For details, see the comments in the **web.xml** file.

Install rtvquery Servlet

After you configure the servlet, you must install the servlet on your application server. The **servlets\rtvquery** directory contains an installation script to install the rtvquery servlet on the Tomcat instance included with RTView.

To install the rtvquery servlet:

"Initialize a Command Prompt or Terminal Window", go to the **servlets\rtvquery** directory and run the following scripts.

NOTE: These scripts require that you set the environment variable CATALINA_HOME to the location of your Tomcat installation directory, see the "Setup" on page 5 section for details.

- make_war This script creates a web archive (.war) named rtvquery.war.
- install_to_tomcat rtvquery This script installs the web archive rtvquery.war to your Tomcat server. NOTE: This script will shutdown and restart Tomcat and requires administrative permissions.

NOTE: If you are using an application server other than Tomcat, install the **rtvquery.war** file on your application server according to instructions given with that product.

Deploy Multiple rtvquery Servlets

One instance of the rtvquery servlet provides access to one Data Server. To provide access to multiple Data Servers via the rtvquery servlet, you have two options:

(1) Deploy one instance of the servlet on multiple application (web) servers: With this option, if the Data Server runs on a different host than the application server or uses a port other than 3278, you must edit the **DataServerHost** and **DataServerPort** properties in the rtvdisplay.properties file accordingly. Repeat this for each application server on which the servlet is deployed.

(2) Deploy multiple instances of the servlet on the same application server: With this option, for each servlet instance you must edit the rtvdisplay.properties file to specify a unique **DataServerHost** and/or **DataServerPort** value. Then, before deploying rtvquery.war file, rename the file to a unique name for that application server (for example,

rtvquery_dataserver1.war, rtvquery_dataserver2.war, and so forth). Use that filename when you install each .war file, for example:

install_to_tomcat rtvquery_dataserver1

The name you choose for each **.war** file is also used in the URL when accessing that specific servlet. For example:

```
http://somehost:8068/rtvquery_dataserver1/...
```

```
http://somehost:8068/rtvquery_dataserver2/...
```

Using rtvquery Servlet

This section provides the information you need to obtain internal data from the Monitor using the rtvquery servlet. The rtvquery servlet allows your custom client application to retrieve data from the Data Server via a REST interface. To use the rtvquery servlet you must Configure and Install the servlet.

This section is intended for users with an application server and a standard working knowledge of servlet deployment.

The rtvquery servlet supports queries for tables from either the RTView cache data source or the SQL data source. The query parameters are specified in the URL and the headers of an HTTP GET request. The client sends an HTTP GET to the servlet specifying the query parameters, and the servlet returns the query result in XML, JSON, or plain text format.

NOTE: Rather than formulating HTTP GETs and processing the responses, Ajax clients can make queries more simply by using the "JavaScript Library" included with the servlet.

cachedocs Directory

Refer to the index.html documentation, located in the cachedocs directory, for information about target caches for queries. This automatically generated documentation describes data that is available for queries such as:

- Caches that are available
- Table columns that are available
- Index column names
- Historical data available
- Compression method for historical data

NOTE: The rtvquery servlet is intended for use by custom browser-based applications to display tables from the RTView Data Server at relatively low data volumes. For high data volume displays, the RTView Viewer application or Thin Client is recommended.

This section includes:

- "Cache Table Query Format," next
- "SQL Table Query Format" on page 204
- "Response Formats" on page 204
- "Servlet Configuration Files" on page 207
- "JavaScript Library" on page 207
- "rtvQuery Class Fields" on page 207
- "rtvQuery Class Functions" on page 208

For rtvquery servlet examples, see the **rtvapm\rtview\custom\rtvquery-samples** directory.

Cache Table Query Format

If the rtvquery servlet is deployed at **http://host/rtvquery**, the URL for a cache query is:

http://host/rtvquery/cache/<cacheName>/<tableName>

For example, the URL to request the history table from a cache named **Production** is:

http://host/rtvquery/cache/Production/history

The cache and table names should be URL encoded if they contain special characters. For example, the URL to request the current table from a cache named CPU Usage is:

http://host/rtvquery/cache/CPU+Usage/current

In addition to the cache and table names, several optional parameters can be specified in the URL to refine the query. For example:

http://host/rtvquery/cache/<cacheName>/ <tableName>?<param1>=<value1>&<param2>=<value2>...

Parameter values must be URL encoded if they contain special characters.

Cache Query Optional Parameters

The following are optional parameters for cache queries. Each parameter corresponds to a filter option in a cache data source attachment in RTView. If a parameter is not specified, its default value is used in the query.

NOTE: The optional parameters can also be specified as custom headers in the HTTP GET request. For example, the name of the custom header for a query parameter named "P" is **x-sl-P**. The name of the custom header for specifying the time range is **x-sl-tr**. Unlike URL parameter values, if a query parameter is specified by a custom header, the value should not be URL encoded. Specifying query parameters in custom headers rather than URL parameters avoids the browser URL length limits (about 2000 characters in some IE versions). If the **fmt** parameter is not specified in the URL or in the **x-sl-fmt** header, the response format is determined by the value of the standard **http "Accept"** header.

Param Description

Default

- The names of table columns to include in the guery response, separated by cols semicolons (;). The names of the filter columns, separated by semicolons (;). none fcol The fcol and fval parameters are used together to specify one or more filters. Only rows that pass the filter(s) are included in the query result. The values of columns specified by fcol are compared with the values listed in the fval parameter. In the simplest case, fcol specifies the name of a single filter column, fval specifies a single filter value, and the result includes only rows in which the filter column matches the filter value. If multiple values are acceptable for the filter, they must be separated by commas in fval, and rows in which the filter column matches any of the filter values are included in the query result. If multiple filter columns are required, they should be separated by semicolons in fcol. The corresponding filter values for each filter column should be separated by semicolons in fval.
- fvalThe filter value(s) for each filter column. This parameter is used with the fcolnoneparameter. See the fcol parameter for further details.

tr

- The time range, in seconds. This parameter is valid only if **table=history**. The **30** tr, tb and te parameters are time parameters and return the following:
 - If only tr is specified, all rows with timestamps >= currentTime tr are returned.
 - If only tb is specified, all rows with **timestamps** >= **tb** are returned.
 - If tb and tr are specified, all rows in the range of **tb to tb + tr** are returned.
 - If only te is specified, all rows with **timestamps** <= **te** are returned.
 - If te and tr are specified, all rows in the range of te tr to te are returned.
 - If te, tr and tb are specified, tr is ignored and all rows with timestamps >= tb and <= te are returned.
- **tb** Specifies the begin time (in milliseconds since 1970) for the rows to be none retrieved. This parameter is valid only if **table=history**. See the tr parameter for details about behavior with other time parameters (tr and te).
- te Specifies the end time (in milliseconds since 1970) for the rows to be retrieved. **none** This parameter is valid only if **table=history**. See the tr parameter for details about behavior with other time parameters (tr and tb).
- **rp** Specifies the maximum number of rows per page to be returned. The rp and pn **none** parameters must be used together and have the following behavior:
 - If rp and pn are specified, the first row in the query result is row number pn * rp (where row zero is the top row of the table) and the last row in the result is row number pn * rp + rp.
 - If only rp is specified, pn defaults to zero.
 - If only pn is specified, it is ignored.
- **pn** Specifies the page number to be returned. The rp and pn parameters must be used together and have the following behavior: **0**
 - If rp and pn are specified, the first row in the query result is row number pn * rp (where row zero is the top row of the table) and the last row in the result is row number pn * rp + rp.
 - If only rp is specified, pn defaults to zero.
 - If only pn is specified, it is ignored.
- to Specifies the query timeout, in seconds.

15

fmt Specifies the desired response format: text, xml, js, or json. (For details, see xml "Response Formats").

If this parameter is not specified in the URL or in the x-sl-fmt header, the response format is determined by the value of the standard http "Accept" header.

Cache Query URL Examples

The following are cache query URL examples. NOTE: For brevity, the base URL (for example, **http://host/rtvquery/**) is omitted from the following examples.

 Get the Customer, Symbol and Purchase Price columns from the current table of a cache named trades, in XML format:

cache/trades/current?cols=Customer;Symbol;Purchase+Price&fmt=xml

 Get all columns of the most recent 5 minutes of data from the trades.history table, for a customer named John Doe, in JavaScript array format:

cache/trades/history?tr=300&fcol=Customer&fval=John+Doe&fmt=js

 Get all columns from the trades.current table, for customers named John Doe or Alice Chen and for symbol = IBM or GE, in json format:

cache/trades/

current?fcol=Customer;Symbol&fval=John+Doe,Alice+Chen;IBM,GE&fmt=json

 Get all columns from trades.history table, with indicated begin and end times, but no more than 1500 rows, in text format:

cache/trades/

history?tb=June+17,2010+10:00:00&te=June+17,2010+10:15:00&rp=1500&fmt =text

SQL Table Query Format

If the rtvquery servlet is deployed at http://host/rtvquery, the URL for an SQL query is:

http://host/rtvquery/sql/<dbName>?sql=<sqlQueryString>

where **dbName** is the name of the RTView database connection.

The SQL query string must be URL encoded. For example, the URL to perform the SQL query **select * from production_table** on an RTView database connection named **SampleDB** is:

http://host/rtvquery/sql/sampleDB?sql=select+*+from+production_table

Alternatively, the SQL query string can be omitted from the URL and specified in a custom http header named **x-sl-sql**.

SQL Query Optional Parameters

The following are optional parameters for SQL queries.

Parameter	Default Value	Description
maxrows	none	The maximum number of rows to be returned for the query.
to	15	The amount of time, in seconds, for the query to timeout.
fmt	xml	The query "Response Formats": text, xml, js, xmlrtv, or json.

Response Formats

This section describes supported response formats and response status for the rtvquery servlet. The supported response formats are named **text**, **xml**, **js**, **xmlrtv**, or **json**. This section includes:

- "XML Response Format" on page 204
- "JSON Response Format" on page 205
- "JavaScript Array Response Format" on page 205
- "XMLRTV Response Format" on page 206
- "Text Response Format" on page 206
- "Response Status" on page 206

XML Response Format

The following is an XML response format, where **<DataType>** is one of the following strings: **string**, **int**, **long**, **double**, or **date**.

<dataset>

<metadata>

```
<column name="column 1 name" type=DataType/>
```

... metadata for other columns ...

```
</metadata>
```

```
<data>
```

<row>

<column_1_name>row 1, column 1 value</column_1_name>

... data for other columns in row 1 ...

```
</row>
```

... data for other rows ...

</data>

</dataset>

JSON Response Format

The following is a JSON response format, where **<DataType>** is one of the following strings: **string**, **int**, **long**, **double**, or **date**.

```
{
```

"metadata":[

```
{ "name": "column 1 name", "type": DataType },
```

... metadata for other columns ...

],

"data":[

{"column 1 name": "row1, column1 value", ... data for other columns in row 1},

... data for other rows ...

]

}

JavaScript Array Response Format

The following is a JavaScript Array response format, where **<DataType>** is one of the following strings: **string**, **int**, **long**, **double**, or **date**. The first row in the array contains the column names, the second row contains the column data types, and the remaining rows are the data rows from the data table. This format is the most compact format.

```
[
[
["column 1 name", "column 2 name", ...],
[DataType, DataType, ...],
[value for column 1 in row 1, value for column 2 in row 1, ...],
```

[row 2 values, ...]

...

]

XMLRTV Response Format

The following is an XMLRTV response format, where **<DataType>** is one of the following strings: **string**, **int**, **long**, **double**, or **date**. The XMLRTV format is the traditional XML dataset format used in RTView.

<tc name="column 1 name" type=DataType index="false"/>

```
<tc name="column 2 name" type=DataType index="false"/>
```

...more column definitions ...

value for column 1 in row 1

```
value for column 2 in row 1
```

...

... other rows ...

Text Response Format

The following is a text response format, where **<DataType>** is one of the following strings: **string**, **int**, **long**, **double**, or **date**. The text format uses tabs to separate columns.

column 1 name <tab> column 2 name <tab> ...

row 1, col 1 value <tab> row 2, col 2 value <tab> ...

. . .

Response Status

The following describes the response status indicated by the integer value of the custom header **x-sl-status**, and the corresponding string value of the header **x-sl-status-text**.

x-sl-status	x-sl-status-text	Description
0	ОК	The query was successful.
-1	Not connected to data server.	The query failed and the response is empty.
-2	Query is missing one or more required parameters.	The query failed and the response is empty.
-3	No data received before timeout, query may be invalid.	The query failed and the response is empty.
-4	Error, reason unknown.	The query failed and the response is empty.
-5	Item not found.	The query failed and the response is empty. The cache query specified a cache table that does not exist.

Servlet Configuration Files

The rtvquery.war file contains two files that determine servlet behavior: web.xml and rtvquery.properties. Within rtvquery.war, the paths for these files are WEB-INF/web.xml and WEB-INF/classes/com/sl/rtvquery/rtvquery.properties, respectively.

rtvquery.properties defines the port number of the RTView Data Server to which the servlet connects. The default value is **3278**. Several other servlet properties are also defined in **rtvquery.properties**. See the comments in that file for a description of each.

web.xml should be edited only if it is necessary to change the servlet authentication. By default, authentication is disabled, so any client can submit queries to the servlet. To enable authentication, edit the **web.xml** file, uncommenting the authentication section at the end of the file, and remake and redeploy the rtvquery.**war** file. When authentication is enabled, the browser prompts the user for login information when the first request for a session is sent to the servlet. The user must enter a username and password that are valid for the application server (for example, Tomcat). For details, see the comments in the **web.xml** file.

The source for these two files are in **\servlets\rtvquery**. Use the **make_war.bat** and **make_war.sh** scripts to rebuild rtvquery.war after changing either file.

JavaScript Library

This section is intended for readers familiar with JavaScript, HTML and Ajax.

The rtvquery servlet includes a JavaScript library to simplify development of Ajax client applications. Examples of client applications that use the library are available in the **rtvapm\rtview\custom\rtvquery-samples** directory.

NOTE: Use of the JavaScript library is optional. Alternatively, an application could compose its own HTTP GET requests using the URL and header formats previously described, send them to the rtvquery servlet using XMLHttpRequest, and process the response itself.

The library is contained in a file named **rtvquery.js**. Assuming that the client application is deployed in a web directory that is a sibling of the directory in which the rtvquery servlet is deployed, the following line would typically be used to load the library into an HTML page:

<script src='../rtvquery/rtvquery.js'></script>

The library defines a single JavaScript class named rtvQuery. An instance of that class can be created as follows:

var rtvquery = new rtvQuery();

The rtvQuery constructor takes no arguments.

rtvQuery Class Fields

The following fields are defined by the rtvQuery class.



responseStatusText	String	The status of the last query. If the HTTP request failed, the value indicates the HTTP status (for example, 404 Not Found). If the HTTP request succeeded, the value indicates the query status: either OK for success or one of the error messages described for x-sI-status-text in "Response Status".
response	String or Object	The response result of the last query. If the HTTP request for the query failed, the value is undefined. If the query succeeded and the requested format is text, xml, or xmlrtv the value is a string. If the format is js or json the response is a JavaScript object. See "Response Status" for details.

rtvQuery Class Functions

The following functions are defined by the rtvQuery class.

startQuery(args) Function

Call this function to send a query to the Data Server. The function is returned immediately after the request is sent. The query result is returned asynchronously via the user-defined function specified by the **doneCB** field using arguments. The startQuery function expects a single argument containing the following fields.

Field	Description
baseURL	A string indicating the prefix to be prepended to the URL used to access the rtvquery servlet. For example, if the rtvquery servlet is located on the same application server as the calling application, the value would typically be /rtvquery . The default is "" .
format	A string indicating the desired response format: xml , json , js , xmlrtv , or text as described in "Response Formats". The default is xml .
timeout	The query timeout, in seconds. The default is determined by the rtvquery servlet properties file and is typically set to 15 seconds.
doneCB	The function to be called when the query result is received. There is no default value. The rtvQuery object that invoked the startQuery function is passed as the first (and only) argument to the doneCB function. The function can access the query result via the rtvQuery response* fields. See "rtvQuery Class Fields" for details.
noJSConver t	A boolean indicating if a js or json response should be parsed and converted to a JavaScript object. If false, the response is a string. The default is true. Typically this would only be set to false for debugging purposes.

startQuery For Cache Queries

The following fields in the startQuery argument pertain to cache queries. The values should not be URL encoded.

Field	Description
cache	The name of the RTView cache. There is no default value.
table	The name of the table. Typically, this is either current or history. The default is none.
columns	A string containing the names of the cache table columns. The default is $\star.$ For details, see the cols parameter.
filterColumns	A string containing the names of the cache table columns to be used to filter the result. There is no default value. For details, see the fcol parameter.
---------------	---
filterValues	A string containing the values that the filter column must match for a row to be included in the result. No default value. For details, see the fval parameter.
timeRange	The time range, in seconds, for a history query. The default is 30. For details, see the tr parameter.
timeBegin	The begin (minimum) time for a history query. There is no default value. For details, see the tb parameter.
timeEnd	The end (maximum) time for a history query. There is no default value. For details, see the te parameter.
rowsPerPage	The maximum number of rows to be returned. There is no default value. For details, see the rp parameter.
pageNumber	The page number. The default is ${f 0}.$ For details, see the pn parameter.

startQuery For SQL Queries

The following fields in the startQuery argument pertain to SQL queries. The values should not be URL encoded.

Field	Description
database	The RTView database name. There is no default value. For details, see "SQL Table Query Format" on page 204.
sql	The SQL query string. For details, see "SQL Table Query Format" on page 204.
maxRows	The maximum number of rows to be returned. There is no default value.

startQuery Example

The following HTML page calls startQuery to request the current table from a cache named prod_cache, in text format, and displays it in a text area component.

<html>

<head>

<title>Simple cache query using rtvquery servlet</title>

```
<script src='../rtvquery/rtvquery.js'></script>
```

<script>

// callback from 'Run Query' button

```
function doQuery ()
```

{

document.body.style.cursor = 'wait';

var ta = document.getElementById('ResultArea');

ta.value = 'Submitted query, waiting for response ...';

var rtvquery = new rtvQuery();

rtvquery.startQuery({

baseURL : '../rtvquery',

cache : 'prod_cache',

```
table : 'current',
           format : 'text',
           doneCB : function(rtvquery) {
                document.body.style.cursor = 'auto';
                if (rtvquery.responseStatus == 0) {
                      // query successful, show result
                      ta.value = rtvquery.response;
                } else {
                      // query failed, show error msg
                      ta.value = rtvguery.responseStatusText;
                }
           }
     });
}
</script>
</head>
<body>
<button onclick='doQuery()' id='queryButton'>Run Query</button>
<br><br>>
<textarea id='ResultArea' wrap='off' cols='80' rows='20' readonly></textarea>
</bodv>
</html>
```

getCacheNames(args) Function

Call this function to get the names of all RTView caches available from the Data Server. The function returns immediately after the request is sent. The cache names are returned asynchronously via the user-defined function specified by the doneCB argument.

The getCacheNames function expects a single argument containing the following fields.

- baseURL: A string indicating the prefix to be prepended to the URL used to access the rtvquery servlet. If the rtvquery servlet is located on the same application server as the calling application, the value typically would be ../rtvquery. The default value is "".
- doneCB: The function to be called when the result is received. It is called with two arguments. If the query fails, the first argument is null. If the query succeeds, the first argument is a JavaScript array whose first element is an array of the column names, and whose second element is an array of the column types. The column types are the strings int, string, etc. (see "Response Formats" on page 204). The rtvQuery object that invoked getCacheColumns is passed as the second argument to the function. The function can access the query result via the "rtvQuery Class Fields" fields.

getCacheNames Example

The following script calls to getCacheNames populates a drop-down list with the available cache names.

```
var namesQuery = new rtvQuery();
```

```
namesQuery.getCacheNames ({
     baseURL : '../rtvquery',
     doneCB: function(cacheNames, rtvquery) {
           var list = document.getElementById('cacheNameList');
           list.options.length = 0;
           if (!cacheNames || !cacheNames.length) {
                alert(rtvguery.responseStatus == 0?
                     'no caches found' :
                     rtvquery.responseStatusText);
           } else {
                for (var i=0; i < cacheNames.length; ++i) {
                     list.options[i] = new Option(cacheNames[i],
                           cacheNames[i]);
                }
           }
     }
});
```

getCacheColumns(args) Function

Call this function to get the name and type of the columns in a cache table. The function returns immediately after the request is sent. The column information is returned asynchronously via the user-defined function specified by the doneCB argument.

The getCacheNames function expects a single argument containing the following fields.

- baseURL: A string indicating the prefix to be prepended to the URL used to access the rtvquery servlet. Typically, if the rtvquery servlet is located on the same application server as the calling application, the value would be ../rtvquery. The default value is "".
- cache: The name of the RTView cache. There is no default value.
- table: The name of the table. Typically, this is either current or history.
- doneCB: The function to be called when the result is received. It is called with two arguments. If the query fails, the first argument is null. If the query succeeds, the first argument is a JavaScript array whose first element is an array of the column names, and whose second element is an array of the column types. The column types are the strings int, string, etc. (see "Response Formats" on page 204.) The rtvQuery object that invoked getCacheColumns is passed as the second argument to the function. The function can access the query result via the "rtvQuery Class Fields" fields.

getCacheColumns Example

The following script calls to getCacheColumns populates a drop-down list with the column names for the current table of a cache named prod_cache:

var colNamesQuery = new rtvQuery();

```
colNamesQuery.getCacheColumns ({
```

```
baseURL : '../rtvquery',
```

```
cache : 'prod_cache',
     table : 'current',
     doneCB: function(colInfo, rtvquery) {
           // collnfo[0] has column names, collnfo[1] has column types
           var colNames = colInfo ? colInfo[0] : null;
           var list = document.getElementById('cacheNameList');
           list.options.length = 0;
           if (!colNames || !colNames.length) {
                alert(rtvquery.responseStatus == 0?
                      'no columns found' :
                      rtvquery.responseStatusText);
           } else {
                for (var i=0; i < colNames.length; ++i) {
                      list.options[i] = new Option(cacheNames[i],
                            cacheNames[i]);
                 }
           }
     }
});
```

Encrypting Passwords for users.xml

This section describes how to encrypt Monitor passwords for use in the users.xml definition file using the encode_string utility. These instructions assume you have a project directory, as described in Monitor Setup. The **rtview.properties** file is located in your project directory.

The following procedures require the administrator password specified in your users.xml file. The users.xml file is located in the directory where the Monitor is started.

- Perform steps in "Windows"
- Perform steps in "UNIX"

NOTE: For details about using the encode_string utility for a specific data source, refer to the Data Sources/Application Options section for your data source in the RTView User Guide.

users.xml File

The user definition file must be an XML file and start with the following:

```
<?xml version="1.0"?>
```

```
<users xmlns="www.sl.com" >
```

and end with the following:

</users>

For example:

```
<?xml version="1.0"?>
```

```
<users xmlns="www.sl.com" >
```

<user>
<user</user>
<user</user>
<user</user>
<user</user>
<user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</user</us

</user>

```
</users>
```

To encrypt Monitor user passwords

Windows

1. "Initialize a Command Prompt or Terminal Window", and create a scratch directory. For example:

md temp1

2. Change directory to the scratch directory. For example:

cd temp1

- 3. Locate the rtvapm_ocmon.jar file in the ocmon/lib directory.
- **4.** Extract the **users.xml** file into the scratch directory you just created from the gmsjocm.jar file.

Change directory to your scratch directory (it must be the current directory) and type:

jar xvf ..\ocmon\lib\rtvapm_ocmon.jar users.xml

- 5. On the command line, type: encode_string mypassword where mypassword is the plain text password you want to encrypt. encode_string outputs the encrypted password. For example: encrypted value: 013430135501346013310134901353013450134801334
- 6. Copy and paste the encrypted value into the desired location in the users.xml file.
- 7. Copy users.xml to the projects/mysample directory on the appropriate host.

UNIX

1. "Initialize a Command Prompt or Terminal Window", and create a scratch directory. For example:

mkdir temp1

2. Change directory to the scratch directory. For example:

cd temp1

- 3. Locate the rtvapm_ocmon.jar file in the ocmon/lib directory.
- Extract the users.xml file into the scratch directory you just created from the rtvapm_ocmon.jar file.

Change directory to your scratch directory (it must be the current directory) and type:

jar xvf ../ocmon/lib/rtvapm_ocmon.jar users.xml

- 5. On the command line, type: encode_string mypassword where mypassword is the plain text password you want to encrypt. The encode_string utility outputs the encrypted password. For example: encrypted value: 013430135501346013310134901353013450134801334
- 6. Copy and paste the encrypted value into the desired location in the users.xml file.
- 7. Copy users.xml to the projects/mysample directory on the appropriate host.

Invocation-only Configuration

Clusters that do not have storage nodes or caches (*invocation-only* clusters) require additional configuration to be monitored. This section describes how to edit the **invocationonly.properties** file and then use it to enable the Data Server that is to monitor it.

NOTE: This configuration requires that only one invocation only cluster be monitored at a time by a suitably configured Data Server.

- 1. Open the **invocationonly.properties** file, located in the **sample** project directory, in a text editor.
- 2. Locate ACTUAL_CONN_NAME in the file, which resides here:

maincollector.sl.rtview.cache.config=oc_connection_dummy_cache_store.rtv \$conn: ACTUAL_CONN_NAME \$cache:OcStorageDataRaw \$file:ocmon_ts_constants.xml \$table:DummyOcStorageDataRaw

 Change ACTUAL_CONN_NAME to the value of the named connection for the sl.rtview.jmx.jmxconn property (which is used to connect to the invocation-only cluster).

NOTE: Consider using a descriptive name for the named connection, such as "MyInvocationCluster". Your **sl.rtview.jmx.jmxconn** property would then be: **sl.rtvapm.ocmon.jmxconn=MyInvocationCluster**.

4. Continuing with our example, your entry should look like this:

maincollector.sl.rtview.cache.config=oc_connection_dummy_cache_store.rtv \$conn: **MyInvocationCluster** \$cache: OcStorageDataRaw \$file: ocmon_ts_constants.xml \$table: DummyOcStorageDataRaw

- 5. Save the invocationonly.properties file.
- **6.** Start the Data Server that is to monitor the invocation-only cluster and execute the following command line argument:

-properties:invocationonly

NOTE: Alternatively, you can add **-properties:invocationonly** to the **rtvservers.dat** entry used to start the monitoring Data Server.

APPENDIX A Monitor Scripts

This section describes Monitor scripts. We recommend that all scripts be executed in a project directory containing the appropriate configuration files. For example, **ocmon/projects/mysample**.

This section includes:

- "Scripts" on page 217
- "rtvservers.dat" on page 223

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 the following section without the **.bat** suffix; on UNIX systems you must replace the **.bat** suffix with **.sh** for each command. For example, **rundata.sh**.

Name	Description		
dos2unix.sh	Replaces Windows line termination to Unix format. This script is not available as a .bat script.		
	Location: rtvapm/common/bin		
my_alert_actions.bat	Sample script to define actions for alerts.		
/ 511	Location: rtvapm/common/bin		
rtvapm_init.bat/sh Initializes a command window.			
	Format: rtvapm_init.bat		
	rtvapm (execute from the Monitor installation root directory)		
	Note: In UNIX, the script used to initialize a terminal window depends on whether you are in csh or rsh (e.g. Linux, Mac OS X). With a Bourne shell, open a terminal window, go to your Monitor installation directory and type:		
	./rtvapm_init.sh		
runa.bat/sh	Executes RTView Analyzer to extract the function chain of the chosen .rtv file. It returns a .pdf file with a graph of all the function chains in the file.		
	Format: runa.bat [rtvFileName].rtv rtvFileName - Any RTView file with function chains of relative complexity.		
	Location: rtvapm/common/bin		

Name	Description		
runb.bat/sh	Starts the Display Builder. We recommend that you use the start_rtv.bat script with the appropriate configuration and server. For example:		
	start_rtv builder start_rtv builder_ds (Append .sh on UNIX)		
	Format: runb.bat [-ds] [-bg] -ds - To use the currently running Data Server		
	-bg - Runs the Display Builder as a background process.		
	Location: rtvapm/common/bin		
rundata.bat/sh	Starts the Data Server. We recommend that you use the start_rtv.bat script with the appropriate configuration and server. For example:		
	start_rtv default dataserver-properties:sample (Append .sh on UNIX)		
	Format: rundata [-properties:Property File Name] [-propfilter:Property		
	-properties - flag used to apply a specific property file.		
	Property File Name - the name of the properties file to apply.		
	-propfilter - flag used to apply a property filter.		
	Property - the name of the property to apply.		
	Location: rtvapm/common/bin Example: rundata - properties: mycustom properties		
rundb.bat/sh	Starts the HSQLDB database. We recommend that you use the start_rtv.bat script with the appropriate configuration and server. For example:		
	start_rtv default database (Append .sh on UNIX)		
	Location: rtvapm/common/bin		
rundisp.bat/sh	Starts the Display Server. We recommend that you use the start_rtv.bat script with the appropriate configuration and server. For example:		
	start_rtv default displayserver-properties:sample (Append .sh on UNIX)		
	Location: rtvapm/common/bin		
runhist.bat/sh	Starts the Historian. We recommend that you use the start_rtv.bat script with the appropriate configuration and server. For example:		
	(Append .sh on UNIX)		
	Location: rtvapm/common/bin		
runv.bat/sh	Starts the Display Viewer. We recommend that you use the start_rtv.bat script with the appropriate configuration and server. For example:		
	start_rtv default viewer (Append .sh on UNIX) Location: rtvapm/common/bin		

Name	Description		
start_rtv.bat/sh	Starts processes of a named configuration as specified in the rtvservers.dat configuration file. A named configuration should include a Data Server, a Display Server or Viewer, an Historian and a Database. The start_rtv script only attempts to start processes if it detects they are not running. The action of starting processes can be applied to all named configurations, a single named configuration or a single process in a named configuration.		
	Note: The Display Viewer can be started using the start_rtv script, but cannot be stopped using the stop_rtv script. Stop the Viewer by closing its window.		
	To use additional arguments you should either specify a named configuration (to apply the argument to all processes in that configuration) or specify all (to apply the argument to all configurations).		
	-console (or -c) - Start the processes with a command window (which is useful for debugging and 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		
	all		
	Starts all named configurations specified in the rtvservers.dat file. The all argument applies the action to named configurations specified in the rtvservers.dat file.		
	Note: When multiple configurations are specified in the rtvservers.dat file and they have different project directory locations, the all argument processes all of the configurations. However, if the configurations have the same project directory locations, the all argument processes only the first configuration as the others are considered alternative configurations. For details, see the rtvservers.dat section (next).		
	Example: start_rtv.bat all (Append .sh on UNIX)		
	[Configuration Name] To start a single Monitor configuration specified in the rtvservers.dat file:		
	start_rtv.bat [Configuration Name] Configuration Name is the named configuration specified in the rtvservers.dat file. The action applies to all servers or clients specified in the configuration.		
	Example: start_rtv.bat web_deployment (Append .sh on UNIX)		

Name	Description			
	[Process Name] To start a single process in a Monitor configuration specified in the rtvservers.dat file:			
	start_rtv.bat [Configuration Name] [Process Name]			
	Process Name is the name of a process defined in the rtvservers.dat configuration file. For example, dataserver , displayserver , historian and database . The action applies only to that specific process in the configuration.			
	Example: start_rtv.bat default dataserver (Append .sh on UNIX)			
status_rtv.bat/sh	Returns the status all Monitor configurations specified in the rtvservers.dat configuration file. This action uses defined JMX ports for monitoring the status of the process. A named configuration should include a Data Server, a Display Server or Viewer, an Historian and a Database. The action can be applied to all named configurations, a single named configuration or a single process in a named configuration.			
	To use additional arguments you must either specify a configuration (to apply the argument to all servers in a configuration) or all (to apply the argument to each of the configurations).			
	This command returns status information upon execution. For example:			
	status_rtv default returns:			
	dataserver: Running PID 4696 Uptime 000:00:01:47 CPU 00:00:02 Heap 0.7% Clients 2			
	displayserver: Running PID 6340 Uptime 000:00:01:45 CPU 00:00:01 Heap 1.0% Displays 0			
	historian: Running PID 6108 Uptime 000:00:01:42 CPU 00:00:01 Heap 1.3% Connected true database: Running PID 6848 Uptime 000:00:01:39 CPU 00:00:00 Heap 0.4%			
	In the above example, note that the Data Server reports two clients, the Display Server and the Historian. Both the Display Server and the Historian were started with the -ds argument, which connects them to the Data Server. Note also that the Historian reports that it is connected to the database.			
	In the following example, status_rtv default reports that a configured port is in use but the process using the port does not appear to belong to named configuration:			
	dataserver: Data port xxx in use by PID yyy			
	displayserver: JMX port xxx in use by PID yyy			
	and a list of available configurations:			
	Usage: status_rtv config [server] or 'all'			
	Available configs:			
	rtvanm/common/hin			

Name	Description
	all
	Returns the status of all named configurations specified in the rtvservers.dat configuration file. The all argument applies the action to all named configurations specified in the rtvservers.dat configuration file.
	Note: When multiple configurations are specified in the rtvservers.dat file and they have different project directory locations, the all argument processes all of the configurations. However, if the configurations have the same project directory locations, the all argument processes only the first configuration as the others are considered alternative configurations. For details, see the rtvservers.dat section (next).
	Example: status_rtv.bat all (Append .sh on UNIX)
	[Configuration Name] Returns the status of a single named configuration specified in the rtvservers.dat configuration file:
	status_rtv.bat [Configuration Name]
	Configuration Name is the named configuration specified in the rtvservers.dat file. The action applies to all servers or clients specified in the configuration.
	Example:
	status_rtv.bat web_deployment (Append .sh on UNIX)
	[Process Name] Returns the status of a single process in a named configuration specified in the rtvservers.dat configuration file:
	status_rtv.bat [Configuration Name] [Process Name] Process Name is the name of a process defined in the rtvservers.dat configuration file. For example, dataserver, displayserver, historian and database. The action applies only to that specific process in the configuration.
	Example: status_rtv.bat default dataserver (Append .sh on UNIX)

Name	Description		
stop_rtv.bat/sh	Stops processes in a named configuration as specified in the rtvservers.dat configuration file. This action uses defined JMX ports for monitoring. A named configuration might include a Data Server, a Display Server or Viewer, an Historian and a Database. The action can be applied to all named configurations, a single named configuration or a single process in a named configuration.		
	To use additional arguments you must either specify a named configuration (to apply the argument to all processes in that configuration) or all (to apply the argument to all named configurations). If the Display Viewer is started using the start_rtv script, the Viewer cannot be stopped using the stop_rtv script. Stop the Viewer by closing its window.		
	Note: The HSQLDB server (if used) runs with a command window on Windows and cannot be stopped using the stop_rtv command. Stop the HSQLDB server by typing Ctrl-C in its command window.		
	This command returns status information upon execution. For example, stop_rtv.bat default returns:		
	dataserver: Stopped PID 4696 via JMX port 3368		
	If no JMX port is configured the stop_rtv command reports the following:		
	command. If the port is in use but the PID is not available (this might happen on HP- UX and some Linux systems) then the stop_rtv and status_rtv command will report the PID as "???", for example:		
	dataserver: Running PID ??? Uptime 000:00:00:37 CPU 00:00:01 Heap 1.3%		
	dataserver: Stopped PID ??? via JMX port 3368		
	When used without arguments, returns usage information and a list of available configurations. For example, stop_rtv returns:		
	Usage: stop_rtv [ConfigurationName] [ProcessName] or 'all' Location:		
	rtvapm/common/bin		
	all		
	Stops all named configurations specified in the rtvservers.dat configuration file. The all argument applies the action to all named configurations specified in the rtvservers.dat configuration file.		
	Note: When multiple configurations are specified in the rtvservers.dat configuration file and they have different project directory locations, the all argument processes all of the configurations. However, if the configurations have the same project directory locations, the all argument processes only the first configuration as the others are considered alternative configurations. For details, see the rtvservers.dat section (next).		
	Example: stop_rtv.bat all (Append .sh on UNIX)		
	[Configuration Name]		
	Stops a single named configuration specified in the rtvservers.dat configuration file:		
	stop_rtv.bat [Configuration Name]		
	Configuration Name is the named configuration specified in the rtvservers.dat configuration file. The action applies to all processes specified in the configuration.		
	Example:		
	stop_rtv.bat default (Append .sh on UNIX)		

Name	Description		
	[Process Name] Stops a single process in a named configuration specified in the rtvservers.dat configuration file:		
	stop_rtv.bat [Configuration Name] [Process Name]		
	Process Name is the name of a process in the named configuration. For example, dataserver , displayserver , historian and database . The action applies only to the process chosen from the named configuration.		
	Example:		
	stop_rtv.bat default dataserver (Append .sh on UNIX)		
update_wars.bat/sh	Script to regenerate war files when the configuration of the Solution Package has changed.		
	Location:		
	rtvapm/*mon/projects/sample		

rtvservers.dat

The **rtvservers.dat** text file, located in the ***mon/projects/mysample** directory, is used to manage your Monitor deployment and processes.

The **rtvservers.dat** configuration file contains one or more named configurations. A named configuration should include any of the following: Data Server, Historian, HSQLDB database, and either a Display Server (for a Web Deployment) or a Display Viewer (for a Desktop Deployment). The **rtvservers.dat** configuration file is used when the following scripts are executed:

- **start_rtv** Starts Monitor processes specified in the **rtvservers.dat** configuration file.
- **stop_rtv** Stops the Monitor processes specified in the **rtvservers.dat** configuration file.
- status_rtv Returns status information for processes specified in the rtvservers.dat configuration file.

NOTE: You can write the paths in the **rtvservers.dat** configuration file by exclusively using the forward-slash (/) notation for both Windows and UNIX systems. For example, if your project settings directory is located in a subdirectory below the location of your **rtvservers.dat** file, you write the path as **.**/subdirectory on both Windows and UNIX.

Single Configuration File

The following **rtvservers.dat** configuration file contains a single named configuration, which is called *default*, for a Web deployment:

- default . database rundb
- default . dataserver rundata
- default . historian runhist -ds
- default . displayserver rundisp -ds

NOTE: The rtvservers.dat configuration file must end with a new line.

In this example, to start the default configuration type: start_rtv defaultor start_rtv all. To start a single process from the configuration, type start_rtv[ConfigurationName] [ProcessName]. For example: start_rtv default displayserver. Each line has the following format consisting of four fields:

<configuration name=""></configuration>	The name of the configuration ("default" in this example).
<project directory="" location=""></project>	The project directory location, relative to the location of the rtvservers.dat configuration file (".", the current directory, in this example).
<process filter=""></process>	The filter that identifies the process, which is the property filter under which the JMX port for monitoring is defined. By default, this is the server name, such as dataserver , displayserver and historian .
<command/>	The script used to start the process. Valid values are: rundata : Starts the Data Server. runhist : Starts the Historian. rundisp : Starts the Display Server. rundb : Starts the HSQLDB Database. runv : Starts the Display Viewer. runb : Starts the Builder.

o a u					• •
< configuration	Name> <pro< td=""><td>Ject Directory</td><td>Location></td><td><pre><process flit<="" pre=""></process></pre></td><td>er> <command/></td></pro<>	Ject Directory	Location>	<pre><process flit<="" pre=""></process></pre>	er> <command/>

Multiple Configuration File

When multiple configurations are specified in the rtvservers.dat configuration file and they have different project directory locations, the **all** argument processes each of the configurations. However, if the configurations have the same project directory locations, the **all** argument processes only the first configuration as the others are considered alternative configurations. Alternative configurations allow you to alternate among several configurations for the same Monitor deployment.

For example, the following rtvservers.dat configuration file contains two configurations, *system1* and *system2*. Note that the project directory locations differ (../system1 and ../ system2, respectively).

- system1 ../system1 dataserver rundata
- system1 ../system1 historian runhist -ds
- system1 ../system1 displayserver rundisp -ds
- system2 ../system2 dataserver rundata
- system2 ../system2 historian runhist -ds
- system2 ../system2 displayserver rundisp -ds

Because the project directory locations differ, you can type **start_rtv all** to start both configurations. To start only the system1 configuration, type: **start_rtv system1**. To start a single process of the system1 configuration, type **start_rtv <Configuration Name> <Process Name>**. For example: **start_rtv system1 displayserver**.

The following illustrates an **rtvservers.dat** configuration file with two configurations: *desktop*, which is for a Desktop Deployment and *browser*, which is for a Browser Deployment. Note that the project settings directory locations are the same (**../system1** for both). The following **rtvservers.dat** configuration file allows you to switch between a Web Browser and a Desktop Deployment:

desktop ../system1 dataserver rundata-properties:desktop

desktop ../system1 historian runhist -ds-properties:desktop

desktop ../system1 viewer runv -ds-properties:desktop

browser ../system1 dataserver rundata-properties:browser

browser ../system1 historian runhist -ds-properties: browser

browser ../system1 displayserver rundisp -ds-properties: browser

where **desktop.properties** and **browser.properties** files reside in the system1 project directory containing all necessary settings for each deployment.

When the **rtvservers.dat** file contains several alternate configurations as this example does, the **all** argument processes only the first defined configuration, in this case, the desktop configuration. To execute the second configuration you type: **start_rtv browser**.

APPENDIX B Monitor Properties

This section describes properties that are available for the Monitor and how to configure them. This section includes:

- "Overview" on page 227
- "Property Format" on page 228: Describes property format, filters and naming conventions.
- "Properties" on page 229: Describes properties for modifying display behavior, such as drill-down targets.
- "Substitutions" on page 231: Describes substitutions available for modifying display behavior.

Overview

Monitor configuration is specified using a series of properties. Properties can be specified in the command line, in properties files, or in a properties database. However, the most convenient way from a maintenance perspective is to create your own properties files. You can override certain Monitor default settings by editing properties. There are several property files that you might edit to configure or optimize your Monitor deployment.

Modify properties in the **rtview.properties** file. These instructions assume you have a project directory, as described in "Create a Project Directory" on page 7. The **rtview.properties** file is located in your project directory.

Set properties using command line options in an initialized command window. For details about initializing a command window, see "Initialize a Command Prompt or Terminal Window" on page 16.

Options specified on the command line are applied last, therefore command line arguments override values saved in configuration files (such as **.properties** files). Also, in many cases the command-line option cannot be used as a property, or vice versa. For these reasons, we recommend that you use properties rather than command line options. To specify a property in the command line:

- Add a dash (-) at the beginning.
- Remove the prefix (sl.rtview.).
- Remove the "Property Filters" prefix if present (for example, displayserver.). For details, see "Property Filters" on page 228.
- Replace the colon (:) with an equals sign (=).

For example, the property **myprefix.sl.rtview.someflag=true** is **-someflag:true** as a command line option. If a command line argument contains a space or a semicolon, the entire argument must be enclosed in quotes (e.g.: **"-sub:\$data:my Data"**).

Property Format

This section describes the format for Monitor properties in the **rtview.properties** file. These instructions assume you have a project directory, as described in "Create a Project Directory" on page 7. All properties have the prefix **sl.rtview** which is followed by a property name = value pair: **sl.rtview.<property_name>=:<value>**. For example, to specify the Data Server port number globally:

sl.rtview.dataserver.port=3278

Property Filters

Filters are available to apply limit the scope to which a property is applied. Filters precede the **sl.rtview** property prefix followed by a period (.):

<property_filter>.sl.rtview.<property_name>=:<value>.

For example, to specify the Data Server port number to only proxy clients, we use the **proxyclient** filter: **proxyclient.sl.rtview.dataserver.port=3278**

The following Monitor property filters are predefined and apply automatically depending on what tool is being executed:

Filter	Description
agentsender	Applies the property to the Agent sender Data Server. For example:
	agentsender.sl.rtview.dataserver.port=3351
builder	Applies the property to the Display Builder. For example:
	builder.sl.rtview.stylesheet
collector	Applies the property to the Data Collection Server. For example:
	collector.sl.rtview.jmx.jmx_metrics_period=15000
dataserver	Applies the property to the Data Server. For example:
	dataserver.sl.rtview.dataserver.socket=true
displayserver	Applies the property to the Display Server. For example:
	displayserver.sl.rtview.displayserver.port=3079
historian	Applies the property to the Historian. For example:
	historian.sl.rtview.historian.driver=org.hsqldb.jdbcDriver
maincollector	This property filter applies to RTView EM installations. Applies the
	maincollector.sl.rtview.alert.persistAlerts=true
mainroadiyor	This property filter applies to RTView FM installations. Applies the
mainieceivei	property to the Agent main receiver Data Server. For example:
	mainreceiver.sl.rtview.jmxport=8911
proxyclient	Applies the property to the proxy client. For example:
	proxyclient.sl.rtview.dataserver.port=2078
rtvanalyzer	Applies the property to the RTView Analyzer. For example:
	rtvanalyzer.sl.rtview.stylesheet=rtv_default,rtv_flat
viewer	Applies the property to the Display Viewer. For example:
	viewer.sl.rtview.panelconfig=custom_panels.xml

You can define your own property filters and use them as prefixes in your properties files. To select a property filter on the command line use the **-propfilter** argument. For example, to apply the **AlertCollector** filter: **-propfilter**: **AlertCollector**

Properties

The following properties are available for the Monitor.

jmxsampleperiod	Specifies the time interval, in milliseconds, used for polling MBean attributes and operations executed in data attachments if no poll interval is specified in the data attachment. The default is 10000 (10 seconds).
	The jmxsampleperiod property can be used to balance the overhead of requesting the data with the latency of the results. To avoid overloading systems, request data at a rate no faster than it can be produced by the system being monitored. See the "Metrics Administration" on page 173 display to see the total time taken to obtain the JMX data.
	For details about using the jmxsampleperiod property when configuring node expiration modes, see "Node Expiration Mode Substitutions" on page 243.
	Because the Default Poll Interval is superseded by the General Update Period, the amount of time elapsed between MBean polls might be longer than the value entered. For example, if the General Update Period is 2000 milliseconds and the Default Poll Interval is 5000 milliseconds, MBean attributes and operations are polled every six seconds.
sl.rtview.alert.alertcomm and	Specifies to use the RTView EM Self Service Alert system in which the alertCommand is set at runtime by the sl.rtview.alert.alertcommand property. This property is used for Alert Notification.
sl.rtview.cache.config	Specifies a cache definition file. To load several cache definition files this property can be specified multiple times. New properties do not override previous values, they instead add one more file to the list.
	Example: collector.sl.rtview.cache.config=bird_cache.rtv
sl.rtview.cmd_line	Specifies options on the command line. When a new instance of this property is specified, the pre-existing values are not overridden. New properties concatenate its value to the command line. You can specify any valid command line option with a few exceptions. For example, you cannot add a new properties file with -properties, or add a property filter with -propfilter from within a property file. Example:
	displayserver.sl.rtview.cmd_line=-logfile:displayserver.log
sl.rtview.dataserver	Specifies the default Data Server to connect to. This setting must match the Data Server port setting specified for the dataserver.sl.rtview.dataserver.port property. The default is // localhost:3278 .
	Example:
	dataclient.sl.rtview.dataserver=//localhost:3278
	This property can also specify a named Data Server to connect to. The port used must match the port specified in the Data Server we are to connect to.
	Example:
	sl.rtview.dataserver=name=MISCMON-LOCAL;connect=// localhost:10123

si.rtview.dataserver.port	Specifies the Data Server port for client connections. This setting must match the data client connection port setting specified for the dataclient.sl.rtview.dataserver property. The default is 3278 .
	Example:
	dataclient.sl.rtview.dataserver=3278
sl.rtview.dsenable	Specifies the data source to enable.
	Example:
	proxyclient.sl.rtview.dsenable=cache
sl.rtview.historian.charli mit	Specifies the maximum number of characters per table column for the Historian (RTVHISTORY) database. The default is 255 .
	Example:
	historian.sl.rtview.historian.charlimit=255
sl.rtview.historian.comp actiontimerinterval	Specifies how often, in seconds, the aggregation engine checks for data to aggregate in the Historian (RTVHISTORY) database. The default is 5 . Data aggregation reduces the amount of aged data stored in the Historian cache table.
	Example:
	historian.sl.rtview.historian.compactiontimerinterval=5
sl.rtview.historian.comp actionverbose	Specifies for the Historian (RTVHISTORY) database whether and how to output to the console. The default is 0 . There are three options:
	0 - No information is output to the console.
	1 - Summary information is output to the console.
	2 - Debug-level information is output to the console.
	Example:
	historian.si.rtview.historian.compactionverbose=0
sl.rtview.historian.driver	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database.
sl.rtview.historian.driver	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example:
sl.rtview.historian.driver	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver
sl.rtview.historian.driver	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file.
sl.rtview.historian.driver sl.rtview.historian.index _history_tables	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true.
sl.rtview.historian.driver sl.rtview.historian.index _history_tables	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true. Example:
sl.rtview.historian.driver	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true. Example: historian.sl.rtview.historian.index_history_tables=true
sl.rtview.historian.driver sl.rtview.historian.index _history_tables sl.rtview.historian.nores et	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true. Example: historian.sl.rtview.historian.index_history_tables=true Specifies whether to clear the Historian (RTVHISTORY) database tables before storing new data. true does not clear the tables and false does clear the tables. The default is true .
sl.rtview.historian.driver sl.rtview.historian.index _history_tables sl.rtview.historian.nores et	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true. Example: historian.sl.rtview.historian.index_history_tables=true Specifies whether to clear the Historian (RTVHISTORY) database tables before storing new data. true does not clear the tables and false does clear the tables. The default is true . Example:
sl.rtview.historian.driver sl.rtview.historian.index _history_tables sl.rtview.historian.nores et	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true. Example: historian.sl.rtview.historian.index_history_tables=true Specifies whether to clear the Historian (RTVHISTORY) database tables before storing new data. true does not clear the tables and false does clear the tables. The default is true . Example: historian.sl.rtview.historian.noreset=true
sl.rtview.historian.driver sl.rtview.historian.index _history_tables sl.rtview.historian.nores et sl.rtview.historian.passw ord	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true. Example: historian.sl.rtview.historian.index_history_tables=true Specifies whether to clear the Historian (RTVHISTORY) database tables before storing new data. true does not clear the tables and false does clear the tables. The default is true . Example: historian.sl.rtview.historian.noreset=true Specifies the password for accessing the Historian (RTVHISTORY) database. When no characters are entered (for example, historian.sl.rtview.historian.password=) the password is blank, which is also a valid password. Valid values are according to the database engine.
sl.rtview.historian.driver sl.rtview.historian.index _history_tables sl.rtview.historian.nores et sl.rtview.historian.passw ord	Specifies the Historian (RTVHISTORY) database. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production Historian database. Example: sl.rtview.historian.driver=org.hsqldb.jdbcDriver NOTE: To enable the database, uncomment this property in the emcommon.properties file. Specifies whether to add indices when creating tables in the Historian (RTVHISTORY) database. true adds indices and false does not add indices. The default is true. Example: historian.sl.rtview.historian.index_history_tables=true Specifies whether to clear the Historian (RTVHISTORY) database tables before storing new data. true does not clear the tables and false does clear the tables. The default is true . Example: historian.sl.rtview.historian.noreset=true Specifies the password for accessing the Historian (RTVHISTORY) database. When no characters are entered (for example, historian.sl.rtview.historian.password=) the password is blank, which is also a valid password. Valid values are according to the database engine. Example:

sl.rtview.historian.verbo se	Specifies whether to print a line to the console for each record that is stored in the Historian (RTVHISTORY) database. true prints a line for each record and false does not. The default is false . Example: historian.sl.rtview.historian.verbose=false
sl.rtview.historian.url	Specifies the Historian (RTVHISTORY) database URL. Not enabled, by default. Use this property to change the default HSQLDB database (used for evaluation purposes) to a production database. Example: sl.rtview.historian.url=jdbc:hsqldb:hsql://localhost:9099/ rtvhistory NOTE: To enable the database, uncomment this property in the emcommon.properties file.
sl.rtview.historian.usern ame	Specifies the Historian (RTVHISTORY) database user login password. Not enabled, by default. Valid values are according to the database engine. Example: sl.rtview.historian.username=sa
sl.rtview.jmx.jmx_minre connecttime	Specifies the amount of time that elapses, in seconds, before JMX attempts to reconnect. Default is 30 . Example: sl.rtview.jmx.jmx_minreconnecttime=30
sl.rtview.jvm	Specifies command line options for the Java Virtual Machine. For example, to allocate additional heap memory. Examples: displayserver.sl.rtview.jvm=- Djava.awt.headless=truedataserver.sl.rtview.jvm=- Xmx768mdataserver.sl.rtview.jvm=-Xms128m NOTE: With the use of propfilters, Java applications can have different JVM configurations.
sl.rtview.sql.dbretry	Specifies the amount of time, in seconds, between attempts to reconnect to the database. The default is 40000 . Example: sl.rtview.sql.dbretry=40000
sl.rtview.stylesheet	Specifies a list of RTView Style definition files that determine the appearance of RTView displays (text sizes, fonts, colors and so forth). The default is rtv_darkstyles,rtv_flat . Example: sl.rtview.stylesheet=rtv_darkstyles,rtv_flat
sl.rtview.sub	Specifies to use a substitution. For details, see "Substitutions". Example: sl.rtview.sub=\$alertActionScript:my_alert_actions

Substitutions

The following substitutions can be set to modify your Monitor configuration and display behavior. There are different ways you can set a substitution, but the preferred way is to use the sl.rtview.sub property.

The Monitor is a configurable solution for monitoring Coherence clusters. The Monitor comes with default values for configuration options that determine Monitor behavior. Substitutions are a mechanism that allows you to configure Monitor behavior. At runtime, a defined substitution substitutes your own value for the preconfigured default. In this way, the runtime behavior of the Monitor can be configured.

Substitution Syntax

Substitutions are optional and require the following syntax:

sl.rtview.sub=<sub_name>:<sub_value>

For example:

sl.rtview.sub=\$OCMCLUSTERSTATS_TABLE:OCM_CLUSTERSTATS

If a substitution value contains a single quote, it must be escaped using a \boldsymbol{I} :

sl.rtview.sub=\$filter:Plant=/'Dallas/'

If a substitution value contains a space, it must be enclosed in single quotes. Do not escape these single quotes:

sl.rtview.sub=\$subname2:'sub value 2'

A substitution string cannot contain the following:

: | . tab space , ; = < > ' " & / \ { } [] ()

NOTE: The substitution string **\$value** is reserved for internal use.

This section explains how to use the following substitutions:

- "Database Substitutions," next: For configuring database connections and database tables.
- "Alert Substitutions" on page 235: For configuring alert behavior.
- "Filter Substitutions" on page 236: For configuring the JMX query so as to return only data of interest from the Coherence cluster.
- "Cache Substitutions" on page 237: For configuring cache behavior such as how data compaction is performed, how cache names are displayed, and whether to show expired nodes in Monitor displays. This section also describes "Node Expiration Mode Substitutions".
- "Cluster Substitutions" on page 244: For configuring cluster connection and data collection behavior, such as for multiple clusters or large clusters.

You configure a substitution by defining a value for it. To define the value for a substitution you edit the **rtview.properties** file. Typically, you configure substitutions using the sl.rtview.sub property, which is also the recommended method.

Database Substitutions

This section describes substitutions used to configure database connections and database tables. The table names and data connections specified in the substitutions must match the table names and data connections specified for your database configuration.

NOTE: The use of some persisted history value tables is optional. To prevent the use of such tables use the default substitution value of " (two single quotes) which prevents reading and writing of the given database table from the Monitor.

Substitution	Description
\$ALERTDEFS_DB	Use this substitution to specify the SQL connection to use to connect to the database containing alert threshold tables. The default is ALERTDEFS . Example:
	sl.rtview.sub=\$ALERTDEFS_DB:ALERTDEFS
\$ALERTDEFS_TABLE	Use this substitution to specify the database table containing threshold values for scalar alerts. The default is ALERTDEFS . Example: sl.rtview.sub=\$ALERTDEFS_TABLE:ALERTDEFS
\$OCMCACHESERVICEST ATS_TABLE	Use this substitution to specify the name of the persisted history value table for the Monitor cache service statistics. The use of this persisted history value table is optional and not enabled, by default (it has a value of ''). Example: sl.rtview.sub=\$OCMCACHESERVICESTATS_TABLE:OCM_CACHES FRVICESTATS_TABLE
\$OCMCACHESERVICETO TALS_TABLE	Use this substitution to specify the name of the database table containing persisted history values for the Monitor cache service totals. The default is OCM_CACHESERVICETOTALS . Example:
	sl.rtview.sub=\$OCMCACHESERVICETOTALS_TABLE:OCM_CACHE SERVICETOTALS
\$OCMCACHESTATS_TAB LE	Use this substitution to specify the name of the database table containing persisted history values for the Monitor cache statistics. The use of this persisted history value table is optional and not enabled, by default (it has a value of ''). Example:
	sl.rtview.sub=\$OCMCACHESTATS_TABLE:OCM_CACHESSTATS_T ABLE
\$OCMCACHETOTALS_TA BLE	Use this substitution to specify the name of the database table containing persisted history values for the Monitor cache totals. The default is OCM_CACHETOTALS . Example:
	sl.rtview.sub=\$OCMCACHETOTALS_TABLE:OCM_CACHETOTALS
\$OCMCLUSTERSTATS_T ABLE	Use this substitution to specify the name of the persisted history value table for the Monitor cluster statistics. The default is OCM_CLUSTERSTATS . Example:
	sl.rtview.sub=\$OCMCLUSTERSTATS_TABLE:OCM_CLUSTERSTAT S
\$OCMEXTENDCONNECTI ONS_TABLE	Use this substitution to specify the name of the persisted history value table for the Monitor extend connections. The use of this persisted history value table is optional and not enabled, by default (it has a value of "). Example:
	sl.rtview.sub=\$OCMEXTENDCONNECTIONS_TABLE:OCM_EXTEN DCONNECTIONS

\$OCMINVOCATIONSERV ICESTATS_TABLE	Use this substitution to specify the name of the persisted history value table for the Monitor invocation service statistics. The use of this persisted history value table is optional and not enabled, by default (it has a value of "). Example: sl.rtview.sub=\$OCMINVOCATIONSERVICESTATS_TABLE:OCM_I NVOCATIONSERVICESTATS
\$OCMINVOCATIONSERV ICETOTALS_TABLE	Use this substitution to specify the name of the persisted history value table for the Monitor invocation service totals. The default is OCM_CLUSTERSTATS. OCM_INVOCATIONSERVICETOTALS. Example: sl.rtview.sub=\$OCMINVOCATIONSERVICETOTALS_TABLE:OCM_ INVOCATIONSERVICETOTALS
\$OCMJMXMGMTDATA_T ABLE	Use this substitution to specify the name of the persisted history value table for the Monitor JMX management data. The use of this persisted history value table is optional and not enabled, by default (it has a value of ''). Example: sl.rtview.sub=\$OCMJMXMGMTDATA_TABLE:OCM_JMXMGMTDAT A
\$OCMJMXSTATSTOTALS _TABLE	Use this substitution to specify the name of the persisted history value table for Monitor JMX statistic totals. The default is OCM_JMXSTATSTOTALS. Example: sl.rtview.sub=\$OCMJMXSTATSTOTALS_TABLE:OCM_JMXSTATST OTALS
\$OCMJVMGCINFO_TABL E	Use this substitution to specify the name of the persisted history value table for Monitor JVM garbage collection information. The use of this persisted history value table is optional and not enabled, by default (it has a value of ''). Example: sl.rtview.sub=\$OCMJVMGCINFO_TABLE:OCM_JVMGCINFO
\$OCMJVMMEMORYPOOL _TABLE	Use this substitution to specify the name of the persisted history value table for Monitor JVM memory pool data. The use of this persisted history value table is optional and not enabled, by default (it has a value of ''). Example: sl.rtview.sub=\$OCMJVMMEMORYPOOL_TABLE:OCM_JVMMEMOR YPOOL
\$OCMJVMOPERATINGSY STEM2_TABLE	Use this substitution to specify the name of the persisted history value table for Monitor JVM operating system data. The default is OCM_JVMOPERATINGSYSTEM2. Example: sl.rtview.sub=\$OCMJVMOPERATINGSYSTEM2_TABLE:OCM_JVM OPERATINGSYSTEM2
\$OCMNODESTATS_TABL E	Use this substitution to specify the name of the persisted history value table for Monitor node statistics. The default is OCM_NODESTATS . Example: sl.rtview.sub=\$OCMNODESTATS_TABLE:OCM_NODESTATS
\$OCMNODETOTALS_TAB LE	Use this substitution to specify the name of the persisted history value table for Monitor node totals. The default is OCM_NODETOTALS . Example: sl.rtview.sub=\$OCMNODETOTALS_TABLE:OCM_NODETOTALS

\$OCMPROXYSERVICEST ATS_TABLE	Use this substitution to specify the name of the persisted history value table for Monitor proxy service statistics. The use of this persisted history value table is optional and not enabled, by default (it has a value of "). Example: sl.rtview.sub=\$OCMPROXYSERVICESTATS_TABLE:OCM_PROXY SERVICESTATS
\$OCMPROXYSERVICETO TALS_TABLE	Use this substitution to specify the name of the persisted history value table for Monitor proxy service totals. The use of this persisted history value table is optional and not enabled, by default (it has a value of "). Example: sl.rtview.sub=\$OCMPROXYSERVICETOTALS_TABLE:OCM_PROX YSERVICETOTALS
\$OCMSTORAGESTATS_T ABLE	Use this substitution to specify the name of the persisted history value table for Monitor storage statistics. The use of this persisted history value table is optional and not enabled, by default (it has a value of ''). Example: sl.rtview.sub=\$OCMSTORAGESTATS_TABLE:OCM_STORAGESTATS_TABLESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TABLE:OCM_STORAGESTATS_TAB
\$OCMSTORAGETOTALS_ TABLE	Use this substitution to specify the name of the persisted history value table for Monitor storage totals. The default is OCM_STORAGETOTALS. Example: sl.rtview.sub=\$OCMSTORAGETOTALS_TABLE:STORAGETOTALS
\$RTVHISTORY_DB	Use this substitution to specify the name of the SQL connection to use for the database containing persisted history value tables (the named SQL connection is also defined in the rtview.properties file). The default is RTVHISTORY . Example: sl.rtview.sub=\$RTVHISTORY_DB:RTVHISTORY

Alert Substitutions

This section describes substitutions that are used to configure behavior of alerts described in the following table.

Substitution	Description
\$alertActionScript	Specifies the name of the script to execute for an alert command, without the extension. This name is combined with the value of \$scriptEnding to form the complete name of the script.
	Example:
	sl.rtview.cmd_line=-sub:\$alertActionScript:my_alert_actions
\$AVERAGE_MEMORY _TIME_WINDOW	Use this substitution to specify the average memory time window (the time range over which available memory is averaged) for the OcAvailableMemoryLowNodeSpike alert. The default is 86400 seconds (24 hours).
	Example:
	sl.rtview.sub=\$AVERAGE_MEMORY_TIME_WINDOW:86400
\$BAD_COMMUNICATI ON_NODES_TIME_RA NGE	Use this substitution to specify the time range for the OcBadCommunicationNodesInTimeRange alert. The default is 300 seconds (5 minutes). Example:
	sl.rtview.sub=\$BAD_COMMUNICATION_NODES_TIME_RANGE:300

\$domainName	Specifies a domain name to be used by the alert commands. Use this substitution on any Data Server that generates alerts to identify the source of the alert. Example: sl.rtview.sub=\$domainName:mvdomainname
\$NODES_DEPARTED_ TIME_WINDOW	Use this substitution to specify the time window (the time range over which departed nodes are monitored) for the OcDepartedNodesPercentage alert. The default is 300 seconds (5 minutes). Example: sl.rtview.sub=\$NODES_DEPARTED_TIME_WINDOW:300
\$scriptEnding	Specifies the suffix of the script called for an alert command. Typically, it is set to bat on Windows systems and sh on Linux. The default is bat . Example: sl.rtview.cmd_line=-sub:\$scriptEnding:bat

Filter Substitutions

This section describes substitutions that are used to filter the JMX query returning data from the Coherence cluster. Reducing the amount of data returned can improve Monitor performance in cases where returning all data is too much. Filter substitutions specify what data to return in a JMX query (rather than what data to exclude) and subsequently display. Filter substitutions can be used to return all relevant data (when the filter is *) or a subset of data that matches the filter (for example, when the filter is

service=DistributedCache,name=foo,*). Data can also be filtered to include a specific value.

For details about JMX specifications, see http://docs.oracle.com/javase/6/docs/technotes/ guides/jmx/JMX_1_4_specification.pdf.

Substitution Description

\$cacheFilter	Use this substitution to modify a basic Monitor Cache query. The purpose of this substitution is to reduce the amount of Cache MBean data gathered from the cluster and subsequently displayed by the Monitor, thereby improving Monitor performance. The default is * (asterisk), which returns all Cache MBean data. To illustrate, the following examples contain <u>underlined</u> text to indicate where the \$cacheFilter substitution modifies the Monitor Cache query. The following is the basic Cache query used by the Monitor which is modified by the value of the \$cacheFilter substitution variable: Coherence:type=Cache,\$cacheFilter 0 * -1 *-
	Examples.
	si.rtview.sub=\$cacheFilter:*
	and produces this query:
	Concrence:type=cache, 0 ^ -1 ^-
	data (from the DistributedCache service Cache named foo):
	sl.rtview.sub=\$cacheFilter:service=DistributedCache,name=foo,*
	and produces this query:
	Coherence:type=Cache, <u>service=DistributedCache,name=foo,*</u> 0 * -1 *-
\$storageFilter	Use this substitution to modify the basic Monitor StorageManager query. The purpose of this substitution is to reduce the amount of StorageManager MBean data gathered from the cluster and subsequently displayed by the Monitor, thereby improving Monitor performance. The default is * (asterisk), which returns all StorageManager MBean data.
	To illustrate, the following examples contain <u>underlined</u> text to indicate where the \$storageFilter substitution modifies the Monitor StorageManager query.
	The following is the basic StorageManager query used by the Monitor which is modified by the value of the \$storageFilter substitution variable:
	Coherence:type=StorageManager, <u>\$storageFilter</u> 0 * -1 *-
	Examples:
	The following storage filter substitution returns all StorageManager MBean data:
	sl.rtview.sub=\$storageFilter:*
	and produces this query:
	Coherence:type=StorageManager, <u>*</u> 0 * -1 *-
	The following storage filter substitution returns a subset of StorageManager MBean data (from the DistributedCache service Cache named foo):
	sl.rtview.sub=\$storageFilter:service=DistributedCache,cache=foo, *
	and produces this query:
	Coherence:type=StorageManager, <u>service=DistributedCache,cache</u> =foo,* 0 * -1 *-

Cache Substitutions

This section describes substitutions that are used to configure cache behavior. For details about Monitor caches in the cluster that persist data to the database, see the index.html documentation located in the **cachedocs** directory. This documentation describes settings for the cache such as persisted columns, default table sizes and compaction rules.

Substitution	Description
jvmCondenseRowsIn terval	Use this substitution to reduce the amount of in-memory data stored in a JVM cache table via in-memory condensing of historical data. Specifies the time interval used for JVM cache history condensing. The default is 300 seconds (5 minutes). Raw values for this interval are condensed into a single value representing the interval, on a per-column basis.
	Specify a value using the following format:
	NNu
	where \mathbf{NN} is a number and u is a single character. Valid characters are:
	w weeks (7 days)
	d days
	h hours
	m minutes
	s seconds
	For example, to specify a ten minute interval:
	10m
	If only a number is entered, it is assumed to be seconds.
	Example:
	si.rtview.sub=\$jvmCondenseRowsInterval:300
jvmCondenseRowsR awDataTimeSpan	Use this substitution to specify the time span of raw JVM historical data held in-memory before in-memory condensing is applied. The raw data is kept in the JVM cache history table and, if enabled, its history_combo table. By default, this is enabled. The default is 1200 seconds (20 minutes).
	Specify a value using the following format:
	NNu
	where \mathbf{NN} is a number and u is a single character. Valid characters are:
	w weeks (7 days)
	d days
	h hours
	m minutes
	s seconds
	For example, to specify a ten minute interval:
	10m
	If only a number is entered, it is assumed to be seconds.
	Example:
	sl.rtview.sub=\$jvmCondenseRowsRawDataTimeSpan:1200
\$cacheNameFormat	Use this substitution to modify how cache names are shown in Monitor displays.
	By default, \$cacheNameFormat is set to 4*24 which displays the initial 4 characters followed by "" if the name has more than 24 characters, then up to 14 remaining characters, followed again by "" if the name has more than 28 characters.
	You can change the value of $cacheNameFormat$ to N * M , where N is the number of initial characters to display, and M is the number of ending characters to display. In the following example the initial 4 characters of the cache name are displayed, up to 24 ending characters are displayed, and additional characters are elided and replaced by ""
	Example:
	sl.rtview.sub=\$cacheNameFormat:4*24

Use this substitution to reduce the amount of data stored in the Historian table. Data compaction achieves this by aggregating stored data as the \$ocmCompactionRul data ages. By default, data compaction is enabled, with settings suitable 65 for most use cases. When data compaction is not enabled, data must be reduced manually by backing up or deleting archived data. This substitution specifies to aggregate the number of data points and the time intervals for doing so. The default is 1d - ;1w 5m ;1M 15m (see detailed description, below). Compaction is specified using a semi-colon separated list in the following format: \$ocmCompactionRules:'NNu<waitperiod> -;NNu<firstaggregationrule> ;NNu<secondaggregationrule>' where **NN** is a number and u is a single character. Valid characters are as follows: w weeks (7 days) d days **m** minutes M months Using the ocmCompactionRules default settings, for example: sl.rtview.sub=\$ocmCompactionRules:'1d - ;1w 5m ;1M 15m', no compaction occurs for data less than 24 hours old -- a 1 day wait period specified by the first rule: 1d -. During this time data is stored 3600 points per hour (every second). When data is 1 day old, compaction begins at 5 minute intervals for the next week, specified by the second rule: **1w 5m**. During this time the data is aggregated at a compaction level of 12 points per hour (60 minutes divided by 5 minutes). When the data is 8 days old (1 week + 1 day), compaction occurs at 15 minute intervals for the next month, specified by the third rule: 1M 15m. During this time the data is aggregated at a compaction level of 4 points per hour (60 minutes divided by 15 minutes). When that data is 38 days old (1 month + 1 week + 1 day), the data is stored in the Historian table at the compaction level of 4 points per hour. Data compaction increases the length of time between trend graph data

Data compaction increases the length of time between trend graph data points as the data ages. You can modify compaction settings by editing the ocmCompactionRules substitution in the **rtview.properties** file. For example, if you need to further reduce the amount of stored Historian data you might increase the compaction level sooner--in the second rule--from 4 points per hour to 1 point per hour. The 4 points per hour compaction level is the maximum recommended as trend graphs plot gaps when the level is above this.

Conversely, if you need more data points to be visible in trend graphs, you might decrease the compaction level from 4 points per hour to 8 points per hour.

Example:

sl.rtview.sub=\$ocmCompactionRules:'1d - ;1w 5m ;1M 15m'

\$ocmCondenseRows Interval	Use this substitution to reduce the amount of in-memory data stored in a cache table via in-memory condensing of historical data. Specifies the time interval used for Monitor cache history condensing. The default is 300 seconds (5 minutes). Raw values for this interval are condensed into a single value representing the interval, on a per-column basis. Specify a value using the following format: NNu where NN is a number and u is a single character. Valid characters are: w weeks (7 days) d days h hours m minutes s seconds For example, to specify a ten minute interval: 10m If only a number is entered, it is assumed to be seconds. Example: s.
\$ocmCondenseRows RawDataTimeSpan	Use this substitution to specify the time span of raw Monitor historical data held in-memory before in-memory condensing is applied. The raw data is kept in the Monitor cache history table and, if enabled, its history_combo table. By default, this is enabled. The default is 1200 seconds (20 minutes). Specify a value using the following format: NNu where NN is a number and u is a single character. Valid characters are: w weeks (7 days) d days h hours m minutes s seconds For example, to specify a ten minute interval: 10m If only a number is entered, it is assumed to be seconds. Example: sl.rtview.sub=\$jvmCondenseRowsRawDataTimeSpan:1200
\$ocmHistoryTimeSpa n	Use this substitution to specify, in seconds, the number of days of history data to load at startup. This substitution can be used to limit the SQL result set. The default is 1296000 (15 days). Example: sl.rtview.sub=\$ocmHistoryTimeSpan:1296000

ocmMaxNumberOfHi storyRowsLarge	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 300000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store.
	Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory.
	Cache size categories with default values are as follows:
	Small : 100000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsSmall substitution to modify the maximum number of rows.
	Medium: 200000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsMedium substitution to modify the maximum number of rows.
	Large : 300000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsLarge substitution to modify the maximum number of rows.
	A higher number of rows typically shortens response times and makes more history data available, while more memory is consumed. A lower number of rows typically lengthens response times as history data not in- memory is read from the SQL database.
	Example:
	sl.rtview.sub=\$ocmMaxNumberOfHistoryRowsLarge:300000
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store.
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store. Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory.
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store. Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory. Cache size categories with default values are as follows:
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store. Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory. Cache size categories with default values are as follows: Small : 100000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsSmall substitution to modify the maximum number of rows.
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store. Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory. Cache size categories with default values are as follows: Small : 100000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsSmall substitution to modify the maximum number of rows.
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store. Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory. Cache size categories with default values are as follows: Small : 100000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsSmall substitution to modify the maximum number of rows. Medium : 200000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsMedium substitution to modify the maximum number of rows.
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store. Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory. Cache size categories with default values are as follows: Small : 100000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsSmall substitution to modify the maximum number of rows. Medium : 200000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsMedium substitution to modify the maximum number of rows. Large : 300000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsLarge substitution to modify the maximum number of rows. A higher number of rows typically shortens response times and makes more history data available, while more memory is consumed. A lower number of rows typically lengthens response times as history data not in-memory is read from the SQL database.
ocmMaxNumberOfHi storyRowsMedium	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 200000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store. Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory. Cache size categories with default values are as follows: Small : 100000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsSmall substitution to modify the maximum number of rows. Medium : 200000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsMedium substitution to modify the maximum number of rows. Large : 300000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsLarge substitution to modify the maximum number of rows. A higher number of rows typically shortens response times and makes more history data available, while more memory is consumed. A lower number of rows typically lengthens response times as history data not in-memory is read from the SQL database.

ocmMaxNumberOfHi storyRowsSmall	Use this substitution to size in-memory storage of history data. This substitution is typically helpful in multi-cluster monitoring, where a cache is used to hold data from multiple clusters. The default is 100000 . This substitution is one of three substitutions that are used for the same purpose but for different cache sizes. By default, caches that store history data are categorized by size (as small, medium or large) according to the expected maximum number of history rows they store.
	Determine the size category of a cache by referring to the number of rows specified for Max Number Of History Rows in the index.html documentation, located in the cachedocs directory.
	Cache size categories with default values are as follows:
	Small : 100000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsSmall substitution to modify the maximum number of rows.
	Medium : 200000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsMedium substitution to modify the maximum number of rows.
	Large : 300000 maximum number of history rows. Use the ocmMaxNumberOfHistoryRowsLarge substitution to modify the maximum number of rows.
	A higher number of rows typically shortens response times and makes more history data available, while more memory is consumed. A lower number of rows typically lengthens response times as history data not in- memory is read from the SQL database.
	Example:
	sl.rtview.sub=\$ocmMaxNumberOfHistoryRowsSmall:100000
ode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode.
ode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays).
ocmRowExpirationM ode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays). When enabled (1) only active node counts are included in the total number of nodes in the system. Expired nodes are included in displays that show all nodes and the expired nodes are highlighted in red. Trend graphs stop updating expired nodes at the time of departure. When displays show selectable nodes (heatmaps, table rows, grids and drop-down lists) the total of selectable nodes is shown - active nodes and expired nodes which are highlighted in red. Also, node drop-down lists include the suffix [X] for departed nodes. Single node displays have a red background for expired nodes.
ocmRowExpirationM ode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays). When enabled (1) only active node counts are included in the total number of nodes in the system. Expired nodes are included in displays that show all nodes and the expired nodes are highlighted in red. Trend graphs stop updating expired nodes at the time of departure. When displays show selectable nodes (heatmaps, table rows, grids and drop-down lists) the total of selectable nodes is shown - active nodes and expired nodes which are highlighted in red. Also, node drop-down lists include the suffix [X] for departed nodes. Single node displays have a red background for expired nodes. When not enabled (3) only active nodes are included in the total number of nodes in the system and expired nodes are not shown in displays (they are not considered part of the system). where:
ocmRowExpirationM ode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays). When enabled (1) only active node counts are included in the total number of nodes in the system. Expired nodes are included in displays that show all nodes and the expired nodes are highlighted in red. Trend graphs stop updating expired nodes at the time of departure. When displays show selectable nodes (heatmaps, table rows, grids and drop-down lists) the total of selectable nodes is shown - active nodes and expired nodes which are highlighted in red. Also, node drop-down lists include the suffix [X] for departed nodes. Single node displays have a red background for expired nodes. When not enabled (3) only active nodes are included in the total number of nodes in the system and expired nodes are not shown in displays (they are not considered part of the system). where: 1 - Specifies to mark and show expired nodes in displays and allow them to
ocmRowExpirationMode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays). When enabled (1) only active node counts are included in the total number of nodes in the system. Expired nodes are included in displays that show all nodes and the expired nodes are highlighted in red. Trend graphs stop updating expired nodes at the time of departure. When displays show selectable nodes (heatmaps, table rows, grids and drop-down lists) the total of selectable nodes is shown - active nodes and expired nodes which are highlighted in red. Also, node drop-down lists include the suffix [X] for departed nodes. Single node displays have a red background for expired nodes. When not enabled (3) only active nodes are included in the total number of nodes in the system and expired nodes are not shown in displays (they are not considered part of the system). where: 1 - Specifies to mark and show expired nodes in displays and allow them to be selectable.
ocmRowExpirationM ode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays). When enabled (1) only active node counts are included in the total number of nodes in the system. Expired nodes are included in displays that show all nodes and the expired nodes are highlighted in red. Trend graphs stop updating expired nodes at the time of departure. When displays show selectable nodes (heatmaps, table rows, grids and drop-down lists) the total of selectable nodes is shown - active nodes and expired nodes which are highlighted in red. Also, node drop-down lists include the suffix [X] for departed nodes. Single node displays have a red background for expired nodes. When not enabled (3) only active nodes are included in the total number of nodes in the system and expired nodes are not shown in displays (they are not considered part of the system). where: 1 - Specifies to mark and show expired nodes in displays and allow them to be selectable. 3 - Specifies not to mark and show expired nodes are not part of the system.
ocmRowExpirationMode	Use this substitution with the ocmRowExpirationTime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays). When enabled (1) only active node counts are included in the total number of nodes in the system. Expired nodes are included in displays that show all nodes and the expired nodes are highlighted in red. Trend graphs stop updating expired nodes at the time of departure. When displays show selectable nodes (heatmaps, table rows, grids and drop-down lists) the total of selectable nodes is shown - active nodes and expired nodes which are highlighted in red. Also, node drop-down lists include the suffix [X] for departed nodes. Single node displays have a red background for expired nodes. When not enabled (3) only active nodes are included in the total number of nodes in the system and expired nodes are not shown in displays (they are not considered part of the system). where: 1 - Specifies to mark and show expired nodes in displays and allow them to be selectable. 3 - Specifies not to mark and show expired nodes are not part of the system. Use the ocmRowExpirationTime substitution to specify the amount of time, in seconds, after which a node is considered expired.
ocmRowExpirationM ode	Use this substitution with the ocmRowExpiration lime and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode. Use this substitution to make expired nodes visible and selectable in Monitor displays. The default is 3 (to not mark and show expired nodes in displays). When enabled (1) only active node counts are included in the total number of nodes in the system. Expired nodes are included in displays that show all nodes and the expired nodes are highlighted in red. Trend graphs stop updating expired nodes at the time of departure. When displays show selectable nodes (heatmaps, table rows, grids and drop-down lists) the total of selectable nodes is shown - active nodes and expired nodes which are highlighted in red. Also, node drop-down lists include the suffix [X] for departed nodes. Single node displays have a red background for expired nodes. When not enabled (3) only active nodes are included in the total number of nodes in the system and expired nodes are not shown in displays (they are not considered part of the system). where: 1 - Specifies to mark and show expired nodes in displays and allow them to be selectable. 3 - Specifies not to mark and show expired nodes are not part of the system. Use the ocmRowExpirationTime substitution to specify the amount of time, in seconds, after which a node is considered expired. Example:

ocmRowExpirationTi me	Use this substitution with the ocmRowExpirationMode and ocmRowExpirationTimeForDelete substitutions to configure the Node Expiration Mode.
	Specifies the amount of time, in seconds, after which a node is considered expired when data updates are not received from it. The default is 25 (seconds).
	Best practices dictate to allow at least two JMX updates to detect an expired node. Less than two updates might give a false positive. If node data is missing from one sample, the second sample can confirm it, making a false positive unlikely. To ensure a minimum of two JMX updates, set the ocmRowExpirationTime to 2.5 x the current JMX MBean sampling interval. For example, if the JMX MBean sampling interval is 10 seconds, set the ocmRowExpirationTime substitution to ocmRowExpirationTime:25 .
	Also note that if the ocmRowExpirationTime is set to 3 (or more) x the current JMX MBean sampling interval, it will take at least three (or more) updates after no data is received from a node before a node is marked expired. Therefore, a higher setting can increase the latency in detecting expired nodes.
	The JMX MBean sampling interval is specified by the collector.sl.rtvapm.ocmon.jmxsampleperiod property.
	Example:
	sl.rtview.sub=\$ocmRowExpirationTime:25
ocmRowExpirationTi meForDelete	Use this substitution with the ocmRowExpirationMode and ocmRowExpirationTime substitutions to configure the Node Expiration Mode.
	Specifies the amount of time, in seconds, after which an expired node is no longer shown in displays. The default is 25 (seconds).
	Example:
	sl.rtview.sub=\$ocmRowExpirationTimeForDelete:25

Node Expiration Mode Substitutions

When nodes expire, by default they are no longer selectable, nor are they shown, in Monitor displays. However, under certain circumstances it might be beneficial to display them, and control how long expired nodes are shown in Monitor displays. There are three possible modes you can configure for expired nodes:

- Mode 1: Expired nodes are not shown in displays (the default)
- Mode 2: Expired nodes are shown and selectable in displays indefinitely. Expired nodes
 persist as expired nodes until they rejoin the cluster. If there is a large population of
 expired nodes, consider Mode 3.
- Mode 3: Expired nodes are shown and selectable in displays for a specified time, then they are removed from displays at a user-specified time. This option enables you to manage the clutter of expired nodes – with the time window with which you wish to investigate them.

NOTE: Expired nodes that rejoin the cluster are no longer considered expired, and thus are displayed and selectable.

To change the node expiration mode you configure three substitutions, ocmRowExpirationMode, ocmRowExpirationTime and ocmRowExpirationTimeForDelete.

Example Mode 1: Expired nodes not shown in displays

sl.rtview.sub=\$ocmRowExpirationMode:3

sl.rtview.sub=\$ocmRowExpirationTime:25

sl.rtview.sub=\$ocmRowExpirationTimeForDelete:25

Where:

\$ocmRowExpirationTime is 2.5 times the jmxsampleperiod in seconds

\$ocmRowExpirationTimeForDelete is the same value as ocmRowExpirationTime (nodes are deleted as they expire and are thus not displayed)

Example Mode 2: Expired nodes shown and selectable in displays indefinitely

sl.rtview.sub=\$ocmRowExpirationMode:1

sl.rtview.sub=\$ocmRowExpirationTime:25

sl.rtview.sub=\$ocmRowExpirationTimeForDelete:25

Where:

\$ocmRowExpirationMode is 1

\$ocmRowExpirationTime is 2.5 times the jmxsampleperiod in seconds

\$ocmRowExpirationTimeForDelete is ignored in this mode

Example Mode 3: Expired nodes shown and selectable in displays for a specified time

sl.rtview.sub=\$ocmRowExpirationMode:3

sl.rtview.sub=\$ocmRowExpirationTime:25

sl.rtview.sub=\$ocmRowExpirationTimeForDelete:86400

Where:

\$ocmRowExpirationMode is 3

\$ocmRowExpirationTime is 2.5 times the jmxsampleperiod in seconds

\$ocmRowExpirationTimeForDelete is the amount of time, in milliseconds, expired nodes are displayed. This value must be longer than \$ocmRowExpirationTime. A value of 86400 would display expired nodes for 24 hours.

collector.sl.rtvapm.ocmon.jmxsampleperiod

It is helpful to understand jmxsampleperiod property when configuring node expiration modes. jmxsampleperiod is a property used to control the rate at which JMX MBean attributes are polled. It can be used to balance the overhead of requesting the data, with the latency of the results. To avoid overloading systems, request data at a rate no faster than it can be produced by the system being monitored. See the Metrics Administration display to see the total time taken to obtain the JMX data.

The jmxsampleperiod property specifies the time interval, in milliseconds, for polling MBean attributes and operations executed in data attachments if no poll interval is specified in the data attachment. The default is 10000 (10 seconds). This attribute is specified in the **rtview.properties** file, located in the Monitor **projects/mysample** directory.

Because the Default Poll Interval is superseded by the General Update Period, the amount of time elapsed between MBean polls might be longer than the value entered. For example, if the General Update Period is 2000 milliseconds and the Default Poll Interval is 5000 milliseconds, MBean attributes and operations are polled every six seconds.

Cluster Substitutions

This section describes substitutions that are used to configure cluster behavior.

\$coherenceGlobalDomain
Use this substitution to fetch data from "super size" clusters. Specifies the global domain name for JMX Queries. The default is Coherence. Use the default value of Coherence to fetch data from Coherence MBeans. NOTE: This feature requires additional system management for the cluster that is not included with the Monitor. For information, contact SL Corporation, at info@sl.com.

Example:

sl.rtview.sub=\$coherenceGlobalDomain:Coherence

\$coherenceLocalDomain: Use this substitution to fetch data from "super size" clusters. Specifies the local domain name for JMX Queries. The default is Coherence. Use the default value of Coherence to fetch data from Coherence MBeans. NOTE: This feature requires additional system management for the cluster that is not included with the Monitor. For information, contact SL Corporation, at info@sl.com.

Example:

sl.rtview.sub=\$coherenceLocalDomain:Coherence

APPENDIX C JMX Connection Options

The RTView Monitor application collects capacity and performance metrics from an operational Coherence Cluster using standard JMX protocols. These metrics are made available to developers and support personnel for analysis and alerting using RTView desktop applications, Web browser clients, or passively via event-triggered alerts.

There are several modes by which the Monitor may connect to a Coherence cluster using JMX. With RTView, users have a choice as to which mode to use, either of which may be relevant or appropriate depending on the monitoring requirement. This is especially important in a situation where users are called on to monitor and manage multiple disparate clusters. This section includes:

- "Connection to Cluster Using JMX Remote Port or RMI URL" on page 247
- "Optimizing Data Retrieval Using JMX Tables" on page 248
- "Direct Connection to Cluster as a Coherence Management Node" on page 250

Connection to Cluster Using JMX Remote Port or RMI URL

In this mode, the Monitor makes a connection to a remote JMX port or RMI URL exposed by a node in the cluster that has been configured as a Coherence "management" node on startup. This node must also have defined its JMX remote port or RMI URL using standard JMX configuration properties and may include a requirement for secure user authentication.

Once connected, the Monitor begins querying all (or a subset) of the MBeans from the Coherence management node at a regular interval.

NOTE: The management node may exist on the same machine as the Monitor; the "remote" designation simply means that the JMX connection is made to MBeans instanced in a separate process from the Monitor.

The information required for the Monitor to connect in this manner is minimal, only the host and port, or RMI URL. Typically, this makes it quick and easy to begin monitoring a cluster, a particular advantage in development environments where clusters come and go on a regular basis. There is no need to configure, then start and stop an agent in order to monitor the cluster.

Another advantage of remote JMX collection is that you do not have to install anything in the cluster or in a production environment – often the cluster itself is running behind a firewall and the monitor does not have easy access to the data. As long as a management node in the cluster exposes JMX MBeans, the connection process can be completely hands off.

A third advantage to this mode is that the Monitor makes no Coherence API calls, meaning that there is a next-to-zero chance of corrupting or crashing the cluster through improper configuration. The rate at which the JMX data are queried can be easily tuned so as to put a minimal monitoring load on the management node in the cluster and on the cluster itself.

Additionally, by having a Coherence management node in the cluster, it can act as backup in case the monitoring system itself goes down.

One disadvantage of the remote JMX connection is that its performance can degrade as the number of monitoring MBeans grows with the complexity of the cluster. A simple measure of cluster complexity is the product of number of nodes (N) times the number of caches supported by the cluster (C). Practical experience has shown that a cluster consisting of 150 nodes and 10 caches (N * C = 1500) can be adequately monitored using the remote JMX connection. Clusters larger than this can benefit from the direct connection mode described in the next section.

Clusters larger than this can benefit from the "Optimizing Data Retrieval Using JMX Tables" mode, or the "Direct Connection to Cluster as a Coherence Management Node" node. The JMX Tables approach has higher performance than the raw JMX approach, but requires custom MBeans to be deployed in the Coherence cluster. The Direct Connection approach has higher performance than JXM Tables but has tradeoffs in the form of access to all of the important cluster configuration parameters, and having the Monitor join the cluster as a management node.



Optimizing Data Retrieval Using JMX Tables

An option is available to speed up retrieval of Coherence MBean information (over JMX) by providing the aggregated MBean data in tabular form by using custom MBeans. By using custom MBeans the data is aggregated within the cluster and transmitted in the form of tabular data, rather than as individual attributes. This reduces the time taken to query the data.

This option is useful when monitoring large clusters (clusters with a large number of nodes, caches and/or services) using JMX, where the volume of data retrieved can affect the time taken to retrieve all the data, and thus limit the sampling rate for monitoring data.

Enabling this requires (unlike default JMX monitoring) that the custom MBeans (contained in a jar) are deployed and registered on all nodes in the cluster, and the monitoring is configured to query the custom MBeans.

The Oracle Coherence Documentation describes registering custom MBeans in a declarative manner in detail: https://docs.oracle.com/cd/E18686_01/coh.37/e18682/custom_mbeans.htm#COHMG4712.

To use this option:

- Configure the monitored Coherence cluster to use JMX Tables custom MBeans. Add the ocjmxtables.jar to the classpath of the cluster members. And set Dtangosol.coherence.mbeans=/sl-custom-mbeans.xml for the cluster members JVM's.
- Configure your Monitoring system to use JMX Tables. Configure your monitoring system to use JMX as normal. And edit the rtview.properties file to use the maincollector.sl.rtview.cmd_line=-ocjmxtables property for the monitoring system.

Requirements:

- The Custom MBeans must be found at run time. You must place the library that contains the MBeans in the classpath of the Coherence nodes/members, including the JMX management-enabled member.
- The custom MBeans (contained in a jar) must be deployed and registered on all nodes in the cluster, and the monitoring configured to query the custom MBeans.
- The Custom MBeans must be specified using a MBean Configuration Override File.
- The Custom MBeans (CacheTable, ServiceTable, StorageManagerTable) are contained in the jar ocjmxtables.jar, located in the rtvapm/ocmon/lib directory of the Monitor installation. This jar file must added to the classpath of the Coherence members to be monitored. This may require that the jar be copied to a location that is visible to all the Coherence members. This may vary based on your deployment. It may prove convenient to copy it to where the Coherence jars are deployed, so they can use the same classpath root.
- The tangosol.coherence.mbeans system property specifies an MBean configuration override file to be used instead of the default custom-mbeans.xml override file. The MBean configuration file to use is sl-custom-mbeans.xml, contained at the root of the ocjmxtables.jar. Thus when the ocjmxtables.jar is added to the Coherence members classpath, it can be specified by setting the tangosol.coherence.mbeans system property for the Coherence cluster members to reference it thus: Dtangosol.coherence.mbeans=/sl-custom-mbeans.xml.

The above should be applied to all Coherence cluster members so that the **tangosol.coherence.mbeans** system property is set to **/sl-custom-mbeans.xml**.

If you have configured your Coherence cluster correctly, you should be able to connect to the cluster using JConsole, and see in addition to the previous Cache, Service, and StorageManager MBeans the new custom CacheTable, ServiceTable, and StorageManagerTable MBeans.

After you configure your Monitor system to use the Custom MBeans and configure your monitoring system to use JMX as normal, uncomment the following line in the **rtview.properties** file:

JMX TABLES

#

Uncomment the line below to use the JMX tables custom mbeans

#maincollector.sl.rtview.cmd_line=-ocjmxtables

This sets the **-ocjmxtables** command line argument to be passed to the maincollector program (typically this is the Data Server), and the log file will then contain the following text at startup:

... using OC JMX Tabular Data

And at runtime, the previous JMX queries (as seen in the **JMX Metrics Administration** display in the **MBean Query Key** column of the **RTView JMX Query Statistics** table):

- * Coherence:type=Cache,* 0 * -1 *-
- * Coherence: type=Cluster 0 * -1 *-
- * Coherence: type=Service, * 0 * -1 *-

become the following:

- * Coherence: type=CacheTable, * 0 CacheTable -1 *-
- * Coherence: type=ServiceTable, * 0 ServiceTable -1 *-
- * Coherence: type=StorageManagerTable, * 0 StorageManagerTable -1 *-

The JMX queries should also have a reduced execution time leading to a reduced total (JMX Query) Execution time.

Direct Connection to Cluster as a Coherence Management Node

In this mode, the Monitor itself joins the cluster and establishes itself as a management node. As a management node, it is configured with local data storage disabled so that it does not store any cache data and serves only as a monitoring node. In this role, it creates the JMX MBean server in-process and collects JMX monitoring data from other Coherence node using fast internal Coherence protocols.

The primary advantage for this mode is speed. In practice, this performance improvement can range from 2 to 10 times faster, depending a number of factors, in particular the network configuration environment.

However, there are tradeoffs. In order to use the direct connection mode, one must have access to all of the important cluster configuration parameters that are used by other nodes in the cluster. These include the Coherence override file, or specific settings like cluster name, well-known address, multicast ports, and Coherence mode. Having limited access to this information can make the configuration process time-consuming.



APPENDIX D Alert Definitions

This section describes alerts for Oracle Coherence and their default settings. This section includes:

- "Alert Thresholds" on page 253
- "Alert Types" on page 253
- "Monitor Alerts" on page 254

Monitor alerts enable you to be aware of components that are unresponsive, performing slowly, generating errors or returning invalid information. By default, Monitor alerts are disabled.

Alert Thresholds

Adequate values for warning and alert thresholds depend on your environment. You might modify alert thresholds if you:

- Receive a lot of alerts for an alert type right after you enable it. In this case, consider increasing the threshold value.
- Receive no alerts for an alert type. In this case, consider decreasing the threshold value.
- Have two different clusters (for example, one for sales business rules and one for inventory management rules) with very different usage patterns. In this case, consider using Alert Override Settings for one of the clusters to apply different thresholds.

How you set alert thresholds also depends on your goals. For example, you might want to:

- Alert IT when user experience is threatened because of capacity constraints.
- Notify IT of unusual activity that is totally independent of system capacity, given that the usage would rarely push the metric above the capacity threshold.

For details about enabling alerts and setting alert thresholds, see the *RTView Enterprise Monitor ® User's Guide*.

Alert Types

To configure thresholds or enable/disable alerts, you configure the Alert Type (also referred to as the Alert Name) such as **OcAvailableMemoryLowNodeSpike**. A single alert type is applied to multiple sources--nodes or caches--depending on the kind of alert type. Alert types issue alerts for any source in the cluster that exceeds the specified thresholds for the alert type. Each alert issued contains a unique identifier for the source that indicates the alert origin.

For example, the **OcAvailableMemoryLowNodeSpike** alert type applies to all nodes in the cluster, therefore the **OcAvailableMemoryLowNodeSpike** threshold settings are the same for all nodes in the cluster. When the **OcAvailableMemoryLowNodeSpike** alert type is disabled, no **OcAvailableMemoryLowNodeSpike** alerts are issued for any node in the cluster. When the **OcAvailableMemoryLowNodeSpike** alert type is enabled, **OcAvailableMemoryLowNodeSpike** alerts are issued for any node in the cluster that exceeds a threshold specified for the **OcAvailableMemoryLowNodeSpike** alert type. Each alert issued by the **OcAvailableMemoryLowNodeSpike** alert type contains the unique identifier for the source node.

Likewise, the **OcEndangeredAllCaches** alert type applies to all caches in the cluster, therefore the **OcEndangeredAllCaches** threshold settings are the same for all caches in the cluster. When the **OcEndangeredAllCaches** alert type is disabled, no

OcEndangeredAllCaches alerts are issued for any cache in the cluster. When the OcEndangeredAllCaches alert type is enabled, **OcEndangeredAllCaches** alerts are issued for any cache in the cluster that exceeds a threshold specified for the **OcEndangeredAllCaches** alert type. Each issued alert contains the unique identifier for the source cache.

Monitor Alerts

The following describes the available Monitor alerts you can enable and their default settings.

OcAvailableMemoryLo wCluster	A single alert is executed if the average percent memory used over max memory of all nodes in the cluster exceeds the specified thresholds.
OcAvailableMemoryLo wNode	For each node in the cluster, an alert is executed if the percent memory used over max memory available for that node exceeds the specified thresholds.
OcAvailableMemoryLo wNodeSpike	For each node in the cluster, an alert is executed if the percent memory used exceeds the specified threshold for the percent above average memory used in the previous 24 hours. For example, if the threshold is set to 50% of total memory used, and the average memory consumption on a particular node for the previous 24 hours is 40%, an alert will be executed if current memory usage exceeds 60% of the total.
	NOTE: The 24 hour time span (86400 seconds) is controlled by the \$AVERAGE_MEMORY_TIME_WINDOW substitution.
	The warning default setting is 115 (percent) of the previous 24 hours and the alarm default setting is 125 (percent) of the previous 24 hours.
	By default the alert is disabled.
OcBadCommunication Cluster	A single alert is executed if the average communication failure rate of all nodes in the cluster exceeds the specified thresholds.
OcBadCommunication Node	For each node in the cluster, an alert is executed if the communication failure rate for that node exceeds the specified thresholds.
OcBadCommunication NodesInTimeRange	Executes a single warning and a single alert if the percentage of nodes in a cluster exceeds the specified threshold for the BadCommunicationNode alert within a time range specified.
	To specify the time range, modify the \$BAD_COMMUNICATION_NODES_TIME_RANGE
	substitution.
	The default time range setting is 5 minutes (300 seconds), the warning default setting is 40 (percent) and the alarm default setting is 50 (percent).
	By default the alert is enabled.

OcCacheHitPercentage Low	This alert is executed when the current Hit% (total current hits/total current gets) is below the specified threshold for a sampling period and the specified cache(s).
OcCacheQueueSizeHig h	A single alert is executed when the CacheQueueSize for all nodes in the cluster exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is 100 (total objects), Alarm is 200 (total objects) and Duration is 60 (total objects).
OcCacheRateCacheMis sesHigh	Executes when the Misses per second exceed the specified threshold and duration. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	By default the alert is disabled with the following default settings: Warning is 1000 , Alarm is 2000 and Duration is 0 (seconds). Before enabling this alert, you MUST change the default settings to values that are suitable for your environment.
OcCacheRateStoreRea dsHigh	Executes when the cache StoreReads rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	By default the alert is disabled with the following default settings: Warning is 1000 , Alarm is 5000 and Duration is 0 (seconds). Before enabling this alert, you MUST change the default settings to values that are suitable for your environment.
OcCacheRateStoreWri tesHigh	Executes when the cache StoreWrites rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	By default the alert is disabled with the following default settings: Warning is 1000 , Alarm is 5000 and Duration is 0 (seconds). Before enabling this alert, you MUST change the default settings to values that are suitable for your environment.
OcCacheRateTotalGets High	Executes when the cache total gets rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	By default the alert is disabled with the following default settings: Warning is 1000 , Alarm is 5000 and Duration is 0 (seconds). Before enabling this alert, you MUST change the default settings to values that are suitable for your environment.

OcCacheRateTotalPuts High	Executes when the cache DeltaTotalPuts rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered. This is a Key Metrics alert that is available with the RTView Enterprise
	By default the alert is disabled with the following default settings: Warning is 1000 , Alarm is 5000 and Duration is 0 (seconds). Before enabling this alert, you MUST change the default settings to values that are suitable for your environment.
OCCacheSizeHigh	Executes when the number of objects in a cache exceeds the specified threshold. By default the alert is disabled with the following default settings: Warning is 1000 (count), Alarm is 5000 (count) and Duration is 60 (seconds).
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	NOTE: If you want to know when the size of a specific cache exceeds specific thresholds, it might be preferable to use the Per Cache or Per Storage Class override settings, allowing you set specific thresholds for specific caches.
OCCacheSizeLow	Executes when the number of objects in a cache goes below the specified threshold. By default the alert is disabled with the following default settings: Warning is 1000 (count), Alarm is 5000 (count) and Duration is 60 (seconds).
	NOTE: If you want to know when the size of a specific cache goes below specific thresholds, it might be preferable to use the Per Cache or Per Storage Class override settings, allowing you set specific thresholds for specific caches.
OcCapacityLimitAllCac hes	An alert is executed if the percent cache used over cache capacity for any cache in the cluster exceeds the specified thresholds. There is one highWarning and one highAlert threshold. For example, if there are 3 caches in a cluster, where:
	cache1 val = 95
	cache2 val = 100
	and the CapacityLimitAllCaches highWarning is 80 and highAlert is 90 , one high alert is executed.
OcCapacityLimitCache	Executes when the average CPU usage for the cluster / storage class exceeds the specified thresholds and durations. This alert has a per cluster and a per (cluster) storage class override. This alert appears in the Other Category when executed.
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	By default the alert is disabled with the following default settings: Warning is 95 (percent), Alarm is 95 (percent) and Duration is 60 (seconds).
OcClusterNodesRcvdF ailureRateHigh	Executes when the average network/packet received failure rate for the cluster/storage class exceeds the specified thresholds and durations. The metrics are averaged across all nodes of a storage class in a cluster.
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	This alert has a per cluster and a per (cluster) storage class override. Note that this alert appears in the Network Category when executed.
	By default the alert is disabled with the following default settings: Warning is 95 (percent), Alarm is 95 (percent) and Duration is 60 (seconds).

OcClusterNodesSentFa ilureRateHigh	Executes when the average network/packet sent failure rate for the cluster / storage class exceeds the specified thresholds and durations. The metrics are averaged across all nodes of a storage class in a cluster.
	This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.
	This alert has a per cluster and a per (cluster) storage class override. Note that this alert appears in the Memory Category when executed.
	By default the alert is disabled with the following default settings: Warning is 95 (percent), Alarm is 95 (percent) and Duration is 60 (seconds).
OcDepartedNode	For each node in the cluster, an alert is executed if the time a node is absent from the cluster exceeds the specified thresholds. When the departed node rejoins the cluster, the alert is cleared.
OcDepartedNodesPerc entage	This scalar alert executes a single warning and a single alert if the percentage of nodes departed from the cluster exceeds the specified thresholds within the specified time periods. The percentage is measured against the total number of nodes in the cluster, including both running and departed nodes.
	The time period is set in the rtview.properties file using the \$NODES_DEPARTED_TIME_WINDOW substitution. The time period can also be overridden using the command line interface. For example, the following sets a time window of 300 seconds:
	-sub: \$NODES_DEPARTED_TIME_WINDOW: 300
	The time period default setting is 600 (10 minutes), the warning default setting is 90 (percent) and the alarm default setting is 95 (percent).
	By default the alert is disabled.
OcEndangeredAllCach es	This alert is executed if the StatusHA for the cache service is NODE_SAFE (high warning) or ENDANGERED (high alert).
OcEndangeredCache	For each node in the cluster, an alert is executed if the StatusHA value is ENDANGERED. By default the alert is disabled.
OcExtendConnectionB yteBacklogHigh	This limits alert executes a single warning and a single alert if the OutgoingByteBacklog for a Proxy Extend Connection exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is 1000 (bytes), Alert is 5000 (bytes).
OcHATargetFailed	This alert executes when the distributed service target status (HATarget) is not met. The HATarget value is determined using the PartitionAsignment MBean in Coherence Versions 12 and above. In prior Coherence versions, the default value of MACHINE-SAFE is used. The default value can be overridden by setting the substitution variable \$ocmDefaultHATarget to the desired value.
OcHighGCDutyCycleNo de	This scalar alert executes a single warning and a single alert if a node exceeds the specified duty cycle threshold (the percent of time spent in Garbage Collection).
	By default the alert is enabled with the following default settings: Warning is 10 (percent), Alarm is 20 (percent) and Duration is 10 seconds.
OcHighPendingReques tNode	A single alert is executed if the RequestPendingCount amount exceeds the specified threshold. This alert allows for setting the warning level, alarm level and duration.
	By default the alert is disabled.

OcHighTaskBacklogNo de	A single warning and a single alert are executed if the number of backlogged tasks exceeds the specified user threshold. This alert allows for setting the warning level, alarm level and duration.
	The default setting executes a warning if the number of backlogged tasks exceeds 10 , and executes an alert if the number of backlogged tasks exceeds 20 .
	By default the alert is disabled.
OcHighThreadAbando nedNode	A single alert is executed if the Coherence Thread Abandoned Count amount exceeds the specified threshold. This alert allows for setting the warning level, alarm level and duration.
	The default setting executes a warning and an alert if the Thread Abandoned Count amount exceeds 0 . The default duration setting is 60 .
	By default the alert is enabled.
OcJmxProcessingTime	This alert is executed if the sum of time for JMX queries and all data processing functions exceeds the specified threshold for the jmxsampleperiod property. By default the alert is disabled with the following default settings: Warning is 80 (percent), Alarm is 90 (percent) and Duration is 0 (seconds).
	NOTE: The OcJmxProcessingTime alert does not support overrides. For that alert the Override Count is displayed as -1 .
OcLongGCDurationNo de	A single warning and a single alert are executed if any of the last garbage collection times exceed the specified duration.
	The default setting executes a warning if the duration exceeds 1 second, and executes an alert if the duration exceeds 2 seconds.
	It is possible for GC times to exceed the specified duration and NOT execute an alert. This is possible if it occurs between the alert duration time and an alert condition time.
	For example, if your alert duration is 60 seconds, and there is also an alert condition set at 27 seconds into that 60 seconds, the following scenarios could occur (where XX:XX:XX is Hours:Minutes:Seconds):
	Scenario 1:
	12:00:00 GC amount is below the specified threshold. No alert executed.
	12:00:27 GC amount exceeds the specified threshold. Alert ignored for now.
	12:01:00 C amount is below the specified threshold. No alert executed.
	Scenario 2:
	12:00:00 GC amount is below the specified threshold. No alert executed.
	now.
	12:01:00 GC amount remains above the specified threshold. Alert executed.
	By default the alert is enabled.
OcLowClientNodeCoun t	This alert executes if the total number of nodes being monitored, including storage enabled nodes, client nodes, and management (JMX) nodes, exceeds the specified threshold. When the count returns to above to above the threshold (departed nodes rejoin the cluster), the alert is cleared.
	By default the alert is disabled.
OcLowStorageNodeCo unt	This alert executes if the total number of storage nodes in the cluster exceeds the specified threshold. When the count returns to above to above the threshold (departed nodes rejoin the cluster), the alert is cleared.
	By default the alert is disabled.

OcLowTotalNodeCount This alert executes if the total number of client nodes being monitored exceeds the specified threshold. When the count returns to above to above the threshold (departed nodes rejoin the cluster), the alert is cleared. By default the alert is disabled. This alert is executed if the percent of memory used on a node after garbage collection exceeds the specified threshold. By default the alert is OcMemoryUsedPercen tageAfterGC disabled with the following default settings: Warning is **70** (percent), Alarm is **80** (percent) and Duration is **30** (seconds). For each node in the cluster, an alert is executed if the StatusHA value is **OcNodeSafeCache** NODE-SAFE. By default the alert is disabled. This alert is executed if a JMX connection remains disconnected after a **OcNoJmxConnection** specified duration of time. The default duration of time is **60** seconds. By default, this alert is enabled. This tabular alert executes a single warning and a single alert for each cache in the cluster if the cache object count delta increases and exceeds **OcObjectCountDeltaU** pCache the specified threshold. In addition to setting the warning and alarm levels, this alert also allows for setting the duration for each cache. When this alert is selected in the Active Alert Table, the Per Cache Alert Setting box is displayed (rather than the scalar alert box). By default the alert is disabled. This tabular alert executes a single warning and a single alert for each **OcObjectCountDeltaD** cache in the cluster where the cache object count delta decreases and ownCache exceeds the specified threshold. In addition to setting the warning and alarm levels, this alert also allows for setting the duration for each cache. When this alert is selected in the Active Alert Table, the Per Cache Alert Setting box is displayed (rather than the scalar alert box). By default the alert is disabled. **OcProxyNodeByteBac** This limits alert executes a single warning and a single alert if the OutgoingByteBacklog for a Proxy Node exceeds the specified threshold. klogHigh This is often indicates overloaded capacity on an individual proxy node. By default the alert is disabled with the following default settings: Warning is **100** (bytes), Alert is **50** (bytes). For each node in the cluster, an alert is executed if the Send Queue for that node exceeds the specified thresholds. By default the alert is **OcSendQueueSize** disabled with the following default settings: Warning is 100 (seconds), Alarm is 200 (seconds) and Duration is 60 (seconds). This alert is executed if the number of StoreFailures exceeds the specified **OcStoreFailure** threshold. By default the alert is disabled with the following default settings: Warning is 1 (second), Alarm is 10 (seconds) and Duration is 30 (seconds). This alert is executed if the current average read per millisecond (total current StoreReadMillis/total current StoreReads) exceeds the specified **OcStoreReadMillisHigh** threshold for a sampling period and the specified cache(s).

Alert Definitions

APPENDIX E Limitations

iPad Safari Limitations

- In the iPad settings for Safari, JavaScript must be ON and Block Pop-ups must be OFF. As of this writing, the Thin Client has been tested only on iOS 4.3.5 in Safari.
- The iPad does not support Adobe Flash, so the Fx graph objects (obj_fxtrend, obj_fxpie, obj_fxbar) are unavailable. The Thin Client automatically replaces the Fx graph objects with the equivalent non-Fx object (obj_trendgraph02, obj_pie, obj_bargraph). Note that the replacement objects behave the same as the Fx objects in most cases but not in all. In particular, obj_trendgraph02 does not support the sliding cursor object nor the legendPosition property. Custom Fx objects are not supported on the iPad.
- The Thin Client implements scrollbars for table objects and graph objects. However, unlike the scrollbars used on desktop browsers, the scrollbars used on the iPad do not have arrow buttons at each end. This can make it difficult to scroll precisely (for example, row by row) on objects with a large scrolling range.
- At full size, users may find it difficult to touch the intended display object without accidentally touching nearby objects and performing an unwanted drill-down, sort, scroll, and so forth. This is particularly true of table objects that support drill-down and also scrolling, and also in panel layouts that contain the tree navigation control. In those cases, the user may want to zoom the iPad screen before interacting with the Thin Client.
- If the iPad sleeps or auto-locks while a Thin Client display is open in Safari, or if the Safari
 application is minimized by clicking on the iPad's home button, the display is not updated
 until the iPad is awakened and Safari is reopened. In some cases it may be necessary to
 refresh the page from Safari's navigation bar.

Because the iPad uses a touch interface there are differences in the Thin Client appearance and behavior in iOS Safari as compared to the conventional desktop browsers that use a cursor (mouse) interface, such as Firefox and Internet Explorer. These are described below.

- Popup browser windows: An RTView object's drill-down target can be configured to open a display in a new window. In a desktop browser, when the RTView object is clicked the drill-down display is opened in a popup browser window. But in iOS Safari 4.3.5, only one page is visible at a time, so when the RTView object is touched a new page containing the drill-down display opens and fills the screen. The Safari navigation bar can be used to toggle between the currently open pages or close them.
- Mouseover text: When mouseover text and drill-down are both enabled on an RTView object (for example, a bar graph), in iOS Safari the first touch on an element in the object (for example, a bar) displays the mouseover text for that element and the second touch on the same element performs the drill-down.
- Resize Mode and Layout: By default, the Display Server runs with resizeMode set to crop. In crop mode, if a display is larger than the panel that contains it only a portion of the display is visible. In a desktop browser, scrollbars become available to allow the user to scroll to view the entire display. In iOS Safari, scrollbars do not appear but the display can be scrolled by dragging two fingers inside the display. (Dragging one finger scrolls the entire page, not the display).

If the Display Server is run with **resizeMode** set to **scale** or **layout**, the display is resized to fit into the panel that contains it. If a desktop browser is resized after a display is opened, the display is resized accordingly. On the iPad, the Safari browser can only be resized by reorienting the iPad itself, between portrait mode and landscape mode.

The panel layout feature is supported in the Thin Client. However, unlike a desktop browser which resizes to match the layout size, the size of Safari is fixed. So if the Display Server is run with **resizeMode** set to **crop** or **scale** mode, there may be unused space at the edges of the display(s) or, in **crop** mode, the panels and displays may be cropped.

This means that **layout** mode should be used for best results on the iPad. For layout mode to be most effective, displays should use the **anchor** and **dock** object properties. Please see RTView documentation for more information.

 Scrolling: The Thin Client implements scrollbars for table objects and graph objects. The scrollbars are activated by dragging with one finger.

If an RTView display is viewed in **crop** mode and is too large to be displayed entirely in Safari, scrollbars do not appear (as they would in a desktop browser) but the display can be scrolled by dragging with two fingers inside the display.

Scrollbars do not ever appear in a text area control. If the text area contains more text than is visible, use the two finger drag in the text area to scroll the text.

Regardless of the size of a listbox control, it can only display a single item (typically, the selected item). When the listbox is touched, the list of items appear in a popup list. In other words, on iOS Safari the listbox control and the combobox control behave identically.

Context menu: The Thin Client context menu is opened by a right mouse button click in a desktop browser. It is opened in iOS Safari by touching any location on a display and holding that touch for 2 seconds. The menu appears in the top left corner of the display, regardless of where the display is touched. The items **Export Table to Excel**, **Drill Down**, and **Execute Command** are not included on the context menu in Safari. All other items are available. The **Export Table to HTML** item is enabled if a table object is touched (unless the table object's drillDownTarget is configured to open another display). After an **Export to PDF/HTML** is performed, the exported content opens on another page in Safari. From there, the content can either be opened by another application (for example, the iBooks application opens PDF) and emailed, or it can be copied ands pasted into an email.

APPENDIX F 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

** JEval 0.9.4 is licensed under the Apache License Version 2.0

Apache License

Version 2.0, January 2004

http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean thepreferred form for making modifications, including but not limited to softwaresource code, documentation source, and configuration files.

"Object" form shall mean anyform resulting from mechanical transformation or translation of aSource form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in orattached to the work (an example is provided in the Appendixbelow)

"Derivative Works" shall meanany work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original workof authorship. For the purposes of this License, Derivative Works shallnot include works that remain separable from, or merely link (or bindby name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean anywork of authorship, including the original version of the Work and anymodifications or additions to that Work or Derivative Worksthereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entityauthorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, orwritten communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailinglists, source code control systems, and issue tracking systems that aremanaged by, or on behalf of, the Licensor for the purpose of discussingand improving the Work, but excluding communication that isconspicuously marked or otherwise designated in writing by the copyrightowner as "Not a Contribution."

"Contributor" shall meanLicensor and any individual or Legal Entity on behalf of whom a Contribution hasbeen received by Licensor and subsequently incorporated within theWork.

2. Grant of Copyright License. Subject tothe terms and conditions of this License, each Contributor herebygrants to You a perpetual, worldwide, non-exclusive, no-charge,royalty-free, irrevocable copyright license to reproduce, prepareDerivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Sourceor Object form.

3. Grant of Patent License. Subject to theterms and conditions of this License, each Contributor herebygrants to You a perpetual, worldwide, non-exclusive, no-charge,royaltyfree, irrevocable (except as stated in this section)patent license to make, have made, use, offer to sell, sell, import, andotherwise transfer the Work, where such license applies only to thosepatent claims licensable by such Contributor that are necessarilyinfringed by their Contribution(s) alone or by combinationof their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against anyentity (including a cross-claim or counterclaim in alawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement,then any patent licenses granted to You under this License forthat Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in anymedium, with or without modifications, and in Source or Objectform, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative 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/orrights consistent with this License. However, in accepting suchobligations, You may act only on Your own behalf and on Your soleresponsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributorharmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License isdistributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANYKIND, either express or implied. See the License for the specific languagegoverning permissions and limitations under the License.

**TreeMap Algorithms v1.0 is used without modifications and licensed by MPL Version 1.1. The source for TreeMap Algorithms can be obtained from http://www.cs.umd.edu/hcil/treemap/

MOZILLA PUBLIC LICENSE

Version 1.1

1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code andModifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the SourceCode notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains anypart of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including allmodules it contains, plus any associated interface definitionfiles, scripts used to control compilation and installation of anExecutable, or source code differential comparisons against eitherthe Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed orarchival form, provided the appropriate decompression or de-archivingsoftware is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complyingwith all of the terms of, this License or a future version of thisLicense issued under Section 6.1. For legal entities, "You"includes any entity which controls, is controlled by, or is under common controlwith 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 more than fifty percent (50%) of the outstanding shares orbeneficial ownership of such entity.

- 2. Source Code License.
- 2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a worldwide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections2.2(a) and 2.2(b) are effective on the date Contributorfirst makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code thatContributor 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 othersoftware (except as part of the Contributor Version) or otherdevices; or 4) under Patent Claims infringed by Covered Code in theabsence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or imposeany terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rightshereunder. 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 tocreate that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly orindirectly, from Original Code provided by the Initial Developerand including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or relateddocumentation in which You describe the origin or ownership of the Covered Code.

3.4. Intellectual Property Matters

(a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations.

Contributor represents that, exceptas disclosed pursuant to Section 3.4(a) above, Contributorbelieves that Contributor's Modifications are Contributor'soriginal creation(s) and/ or Contributor has sufficient rights togrant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation, 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 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 beenmet for that Covered Code, and if You include a notice stating thatthe Source Code version of the Covered Code is available under theterms of this License, including a description of how and whereYou have fulfilled the obligations of Section 3.2. The noticemust be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describerecipients' rights relating to the Covered Code. You may distribute theExecutable version of Covered Code or ownership rights under a licenseof Your choice, which may contain terms different from thisLicense, provided that You are in compliance with the terms of this Licenseand that the license for the Executable version does not attempt tolimit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executableversion under a different license You must make it absolutely clearthat any terms which differ from this License are offered by Youalone, not by the Initial Developer or any Contributor. You herebyagree to indemnify the Initial Developer and every Contributorfor any liability incurred by the Initial Developer or such Contributoras a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all ofthe Covered Code due to statute, judicial order, or regulationthen You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the codethey affect. Such description must be included in the LEGAL filedescribed in Section 3.4 and must be included with all distributions of theSource Code. Except to the extent prohibited by statute orregulation, such description must be sufficiently detailed for a recipient ofordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms ofthat version. You may also choose to use suchCovered Code under the terms of any subsequent version of the Licensepublished by Netscape. No one other than Netscape has the right tomodify the terms applicable to Covered Code created under this License.

6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a)rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or anyconfusingly similar phrase do not appear in your license (except to note that your licensediffers from this License) and (b) otherwise make it clear that Yourversion of the license contains terms which differ from theMozilla Public License and Netscape Public License. (Filling in thename of the Initial Developer, Original Code or Contributorin the notice described in Exhibit A shall not of themselves bedeemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGING. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply withterms herein and fail to cure such breach within 30 days of becomingaware of the breach. All sublicenses to the Covered Code which 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 declatory judgment actions) against Initial Developer or a Contributor (the Initial Developeror Contributor against whom You file such action is referred to as"Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and allrights granted by such Participant to You under Sections 2.1and/or 2.2 of this License shall, upon 60 days notice fromParticipant terminate prospectively, unless if within 60 days after receipt ofnotice You either: (I) agree in writing to pay Participant amutually agreeable reasonable royalty for Your past and future use ofModifications made by such Participant, or (ii) withdraw Yourlitigation claim with respect to the Contributor Version against suchParticipant. If within 60 days of notice, a reasonable royalty andpayment arrangement are not mutually agreed upon in writing by theparties or the litigation claim is not withdrawn, the rights granted byParticipant to You under Sections 2.1 and/or 2.2 automaticallyterminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly orindirectly infringes any patent, then any rights granted to You by suchParticipant under Sections 2.1(b) and 2.2(b) are revoked effective as 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 wheresuch claim is resolved (such as by license or settlement) prior to theinitiation of patent infringement litigation, then thereasonable value of the licenses granted by such Participant underSections 2.1 or 2.2 shall be taken into account in determining the amount orvalue of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU.

10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consistingof "commercial computer software" and "commercialcomputer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept.1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1through 227.7202-4 (June 1995), all U.S. Government End Users acquireCovered Code with only those rights set forth herein.

11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of thisLicense is held to be unenforceable, such provision shall bereformed only to the extent necessary to make it enforceable. ThisLicense shall be governed by California law provisions (except to theextent applicable law, if any, provides otherwise), excluding itsconflict-of-law provisions. With respect to disputes in which atleast one party is a citizen of, or an entity chartered or registered todo business in the United States of America, any litigationrelating to this License shall be subject to the jurisdiction of theFederal Courts of the Northern District of California, with venue lyingin Santa Clara County, California, with the losing partyresponsible for costs, including without limitation, court costs andreasonable attorneys' fees and expenses. The application of the UnitedNations Convention on Contracts for the International Sale ofGoods is expressly excluded. Any law or regulation which provides thatthe language of a contract shall be construed against the draftershall not apply to this License.

12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

EXHIBIT A -Mozilla Public License.

``The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You mayobtain a copy of the License at http://www.mozilla.org/MPL/

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governingrights and limitations under the License.

The Original Code is ____

The Initial Developer of the Original Code is ______ Portions created by ______ are Copyright (C) ______ _____. All Rights Reserved.

Contributor(s): ____

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[___] License"), in which case the provisions of [_____] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [____] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [___] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [___] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

**MD Datejs

Copyright © 2006-2010 Coolite Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to dealin the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE ANDNONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THESOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

**JQuery

Copyright © 2009 John Resig

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE ANDNONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLEFOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OFCONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THESOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

** JCalendar 1.3.2

This product uses JCalendar 1.3.2. JCalendar is distributed pursuant to the terms of the Lesser General Public License. The source code for the JCalendar may be obtained from http://www.toedter.com/en/jcalendar/index.html

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!