

RTView Enterprise Monitor® User's Guide

Version 5.0



RTView®

Copyright © 1998-2019. All rights reserved.

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

RTView Enterprise Monitor®

©2013-2019 Sherrill-Lubinski Corporation. All Rights Reserved.

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

LIMITATIONS ON USE

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

SL, SL-GMS, GMS, RTView, SL Corporation, and the SL logo are trademarks or registered trademarks of Sherrill-Lubinski Corporation in the United States and other countries. Copyright © 1998-2019 Sherrill-Lubinski Corporation. All Rights Reserved.

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

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

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



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

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

Preface	1
About This Guide	1
Audience	1
Document Conventions	1
Additional Resources	2
Release Notes	2
SL Documentation	2
Support Knowledge Base	2
Contacting SL	2
Internet	2
Technical Support	2
 Chapter 1 - Introduction to RTView Enterprise Monitor	 3
Displays	4
RTViewCentral	4
RTView DataServer	5
Solution Packages	6
System Requirements	7
 Chapter 2 - Using the Monitor	 9
Overview	11
Navigation Tabs	12
SERVICE TREE Tab	13
SERVICE VIEWS Tab	18
COMPONENTS Tab	18
ALERTS Tab	21
ADMIN Tab	24
CUSTOM Tab	25
Fundamental Structure of Displays	25
Heatmaps	26
History Heatmaps	28
Tables	29
Multiple Column Sorting	29
Column Visibility	30
Column Filtering	30
Column Locking	32
Column Reordering	32
Saving Settings	32
Row Paging	33
Row Color Code	33
Row Keyboard Selection	34
Trend Graphs	34
Mouse-over	35
Log Scale	35

Time Range	35
Popup Menu	36
Export PDF Report	36
Title Bar	37
Multiple Windows	38
RTViewCentral Displays	39
All Management Areas	39
Area Heatmap.....	40
Area Table.....	41
Multi Area Service Views	43
Group/Service Heatmap	44
Group/Region Heatmap	46
Group / Service Table	47
Services CI Type Summary	49
Services History Heatmap.....	52
Single Area Service Views	54
Single Area: Group/Service Heatmap	55
Single Area: Region/Service Heatmap	56
Single Area: Group / Service Table	58
Single Area: Services CI Type Summary	60
Single Area: Services History Heatmap.....	63
Service Summary Views	65
Service By CI Type	65
Service Summary	69
Service Health Heatmap	73
Key Metrics Views	75
Service KM Heatmap.....	76
Service KM Table.....	80
Service KM History	83
Service KM History (Alt)	87
Available KM Metrics and Alerts	92
Component Views	103
CI / Service Tree View	103
CI / Service Table.....	105
CI / Type Heatmap	107
CI / Type Table	109
Metric Explorer	111
Metric Explorer.....	111
Alert Views	116
RTView Alerts Table	116
Administration	124
Alert Administration	124
Alert Admin Audit	130
Alert Action Audit Trail	132
CMDB Administration.....	133
CMDB Admin.....	133
Architecture	138

System Overview	138
RTView Data Servers	140
Data Server Summary	142
RTView History Table Statistics.....	144
RTView Cache Tables	145
RTView CI Stats Tables	147
RTView CI Type Defs.....	148
RTView KM Defs	150
About	151
Property Views	152
Properties Configuration	152
System Properties	154
Applied Properties	155
All Properties	158
Properties Descriptions.....	160
Diagram Views	162
Steps to Create a Diagram Display.....	164
Create an Object Template Display	165
Node Administration Display	166
Link Admin Display	168
Diagram Properties Admin Display	169
Add Diagrams to your Project.....	171
View Diagram Displays.....	171
Optional Diagram Display Customizations	171
Chapter 3 - RTView DataServer for IBM	175
IBM DB2.....	176
DB2 Instances.....	176
All Instances Heatmap	177
All Instances Table	181
Instance Summary	182
DB2 Members.....	185
All Members Table	185
All Members Grid.....	187
Member Summary	188
All DB2 Databases	191
All Databases Table	191
All Databases Grid	193
Single Database.....	195
Database Summary	195
Partition Table	198
Partition Heatmap	200
Partition Grid	203
Reports.....	205
Trends	206
DB2 Log.....	208

IBM MQ	209
Overview	209
IBM MQ Brokers View	210
All Brokers Grid	211
All Broker Connections	213
Single Broker Summary	214
Channels View	217
All Channels Table	217
Single Channel Summary	219
Single Channel Detail	223
Queue Managers	225
All Queue Managers Table	225
Queues	226
All Queues Table	227
Single Queue Summary	229
IBM WebSphere	232
Overview	232
Key Features	233
Metrics for WebSphere Server	233
End-to-End Context for WebSphere	233
All WebSphere Servers View	233
All Servers Heatmap	234
All Servers Table	236
All Servers Grid	237
All Servers Detail	239
Single Server View	243
Server Summary	244
JDBC Providers	247
Thread Pool	250
All Applications View	251
All Sessions Heatmap	252
Session Charts By App	253
Session Detail By App	255
All Applications Detail	257
Single Application View	258
Application Summary	259
Component Detail	261
Module Totals - Charts	263
Module Totals - Tables	264
Chapter 4 - RTView DataServer for Infrastructure	267
Amazon Web Services	268
EC2 Instances	268
Amazon EC2 Instance Heatmap	269
Amazon EC2 Instance Table	271
Amazon EC2 Instance Summary	273

Docker	275
Engine View	275
Engines Heatmap	276
Engines Table	278
Engine Summary	281
Container View	283
Containers Heatmap	283
Containers Table	286
Container Summary	288
JBoss	291
JBoss Servers View	292
All Servers Heatmap	292
All Servers Table	294
Server Summary	296
JBoss Applications	298
Applications Heatmap	298
Applications Summary	301
MongoDB	303
Mongo Instance View	304
All Instances Heatmap	304
All Instances Table	307
Single Instance Summary	310
Mongo Database View	312
All Databases Heatmap	312
All Databases Table	315
Database Summary	317
Mongo Collection View	320
All Collections Heatmap	320
All Collections Table	323
Collection Summary	325
MySQL	327
All MySQL Databases	327
All Servers Heatmap	327
All Servers Table	330
Single MySQL Database	332
Server Summary	333
Servers Properties	335
Servers Operations	336
User Tables	338
MS SQL Server	339
All Servers View	339
All Servers	339
All Servers Heatmap	342
Single SQL Server View	343
Server Summary	344
Database Details	348

Wait Stats	350
Performance Trends.....	352
Database/Table Sizes.....	355
Node.js	357
Node/Master View.....	357
Node Master Table.....	357
Node Master Summary.....	360
Node Request View	363
Node Requests Table	363
Node Request Summary.....	365
Process View	367
All (Node) Processes Heatmap.....	367
All (Node) Processes Table.....	370
(Node) Process Summary	372
RTView Host Agent.....	375
All Hosts	375
All Hosts Heatmap	375
All Hosts Table	377
All Hosts Grid.....	380
All Processes Table	381
All Network Table	383
All Storage Table	385
Single Host	386
Host Summary.....	387
UX	389
All URLs Table.....	391
All URLs Monitor	394
URL Summary.....	396
All Robots Table	400
All Robots Monitor	403
Robot Summary	405
VMware vCenter	409
Clusters View	410
All Clusters.....	410
Hosts View.....	413
All Hosts	413
Single Host Summary	416
Host Health	419
Host NICs.....	421
Virtual Machines View.....	423
All VMs Heatmap	423
All VMs Table	426
All VMs Disk Table	429
Single VM Summary	431
Datastores View	434
All Datastores Table.....	435

Hosts by Datastore Table.....	437
VMs by Datastore Table.....	438
Single Datastore Summary	440
Networks View	442
All Networks Table.....	442
Events/Alarms View	444
All Events Table	445
All Alarms Table	447
Chapter 5 - RTView DataServer for Kafka.....	449
Apache Kafka	449
Kafka Clusters View.....	450
All Clusters Table.....	451
Cluster Performance	453
Kafka Topics View	455
All Topics Heatmap	455
All Topics Table	457
Single Topic Summary	460
All Topics for Cluster	464
Single Topic for Cluster	466
Kafka Brokers View	468
All Brokers Heatmap	468
All Brokers Table	471
All Broker Metrics Table.....	474
Single Broker Summary	482
Kafka Zookeepers View.....	491
All Zookeepers Heatmap.....	491
All Zookeepers Table.....	494
Zookeepers Summary	497
Kafka Producers View	502
All Producers Heatmap	502
All Producers Table	505
Producer Summary	508
Kafka Consumers View	513
All Consumers Heatmap	513
All Consumers Table	516
Consumers Summary.....	519
Chapter 6 - RTView DataServer for Oracle	525
Oracle Coherence.....	526
Cluster Selector	527
Cluster Views	528
Cluster - Overview.....	529
Caches / Nodes / Alerts	533
Memory/Network Health.....	536
Stability Metrics	538

All Services History	540
All Caches History.....	543
All Nodes History	547
Proxy Services.....	549
Proxy / Extend Overview	550
Proxy / Extend Connections	551
Proxy / Extend Detail	557
Proxy Nodes History	561
Extend Connections History	563
Cache Services	564
Single Service Summary.....	565
Service Metrics Overview.....	569
Service Metric Heatmap.....	570
Single Service History	572
Cache Service Detail	574
Federated Clusters	576
Federated Destination Detail	577
Federated Destination Summary	580
Federated Origin Detail	582
Federated Origin Summary	585
All Caches.....	587
All Caches Heatmap	588
Storage Nodes Cache Map	589
Current Size Chart.....	591
Current Activity Chart	592
Single Cache	594
Single Cache Summary	594
Size Trends	597
Activity Trends.....	600
Cache Detail Tables	602
Storage Manager Detail.....	604
Node/Group Distribution.....	606
Front/Back Analysis	608
All Nodes	609
All Nodes by Type/Host/Memory	610
All Nodes CPU	611
All Nodes Grid View	612
Communication Issues	613
All Nodes - Detail	615
Invocation Service Detail.....	617
Single Node	619
Node Summary	619
Service Trends.....	622
Node Detail	624
JVM Summary.....	626
JVM Memory Trends.....	628
JVM GC Trends.....	630

System Properties	632
Time Range Analysis	633
Service Comparison	633
Cache Comparison	634
OC Administration	636
OC Metrics Administration	636
Cluster MBean Servers	638
Management Settings	640
Node Administration	642
Cache Administration	644
RTView Cache Tables	645
About	647
Oracle Database	647
Database Instances View	648
All Databases Heatmap	648
All Databases Table	649
Database Summary	651
All Instances Heatmap	654
All Instances Table	657
Instance Summary	660
Database Details View	662
Space	662
Contention	665
Processes	667
Sessions	669
Queries	671
Reports	674
Connector for Oracle Enterprise Manager	676
Oracle WebLogic	677
WebLogic Servers View	678
All Servers Heatmap	678
All Servers Table	681
All Servers Grid	684
All Clusters Table	686
Single WebLogic Server View	687
Server Summary	688
WLS JVM Summary	691
WLS Server Detail	694
WLS JDBC Summary	700
WLS ThreadPool Summary	704
Work Manager	707
Persistent Stores	709
Application Views View	711
Cluster Apps Table	711
Cluster App Summary	713
Server Apps Heatmap	716

Server Apps Summary	719
Server Apps Trends	722
App Components Heatmap	724
WebLogic Single Application Summary	726
App Components Summary.....	730
JMS Servers View	732
JMS Servers Grid.....	733
JMS Server Summary.....	734
JMS Metric Trends	738
JMS Server Detail	740
JMS Destinations View	743
Totals by Destination	743
Detail by Destination.....	747
Distribution by Server	750
Destination by Server	753
Detail by Server	756
JMS Bridges View.....	758
All Bridges Table	759
Bridge Summary	761
JMS Connections View	763
JMS Connections - Detail by Server.....	764

Chapter 7 - RTView DataServer for Solace	769
Message Routers.....	770
Message Routers Overview	770
Message Routers Heatmap.....	772
Message Routers Table.....	774
Message Router Summary	782
Message Router Sensors.....	784
Message Router Provisioning	785
Message Router Interface	787
Message Routers Message Spool.....	789
CSPF Neighbors	791
Neighbors Table	791
Neighbors Diagram	792
Neighbors Summary	793
VPNs	795
VPNs Heatmap	796
VPNs Table	800
VPNs Summary	804
Clients.....	807
Clients Table.....	807
Client Summary	812
Bridges.....	815
Bridges Table.....	815
Bridges - Diagram	819

Bridge Summary	820
Endpoints	823
Endpoints Table	823
Endpoint Summary	825
Capacity	828
Capacity Table	828
Capacity - Summary	831
Capacity Trends	833
Syslog	834
All Syslog Events Table	834
Chapter 8 - RTView DataServer for TIBCO	837
TIBCO ActiveMatrix	837
TIBCO ActiveSpaces (2.x)	838
Spaces View	838
All Metaspaces Table	838
Metaspace Summary	841
All Spaces Table	844
All Spaces Heatmap	847
Space Summary	849
All Queries Table	853
Query Summary	855
Members View	857
All Members Table	858
All Members Heatmap	861
Member Summary	864
Member Summary - Process	868
Member Summary -JVM	871
Members by Space Table	874
Members by Space Heatmap	877
Spaces by Member Table	880
Member by Space Summary	883
TIBCO Adapters	886
All Adapters View	886
All Adapters Heatmap	886
All Adapters Table	889
Single Adapter View	891
Adapter Summary	892
TIBCO BusinessEvents	895
Clusters / Nodes View	895
Clusters	895
Cluster Summary	899
Cluster Nodes Table	902
Cluster Nodes Heatmap	905
Inference Node Summary	908

Storage Node Summary	913
Events / Concepts View	917
Agent Events	917
Agent Event Summary	920
Event Cache Hits	923
Event Hit Summary	925
Concept Cache Hits	929
Concept Hit Summary	930
Channels	933
All Inference Agents	935
All RTC Reports	938
TIBCO BusinessWorks	941
BW Applications	941
BW All Applications Heatmap	941
BW All Applications Table	944
BW Single Application Summary	947
BW Containers	950
All Containers Heatmap	950
All Containers Table	953
Single Container Summary	956
BW AppNodes	958
BW All AppNodes Heatmap	958
BW All AppNodes Table	961
BW Single AppNode Summary	963
BW AppSlices	966
BW All AppSlices Heatmap	966
BW All AppSlices Table	969
BW Single AppSlice Summary	971
BW Processes	974
BW All Processes Heatmap	974
BW All Processes Table	977
BW Single Process Summary	980
BW5 Engines	983
All Engines Heatmap	983
All Engines Table	986
All Engines Grid	989
Single Engine Summary	991
BW5 Processes	994
All Processes Heatmap	994
All Processes Table	998
Single Process Summary	1001
BW5 Activities	1003
All Activities Heatmap	1003
All Activities Table	1006
Single Activity Summary	1009
BW5 Servers	1011

All Servers Heatmap	1011
All Servers Table	1014
All Servers Grid	1015
Single Server Summary.....	1017
Server Processes	1019
Single Server Process - Summary	1020
TIBCO Enterprise Message Service	1022
All EMS Servers	1022
All Servers Heatmap	1022
All Servers Table	1025
All Servers Grid	1029
All Servers Topology	1031
Single EMS Server	1035
Single Server Summary.....	1035
Single Server Trends.....	1040
Single Server Tables	1041
EMS Topics.....	1046
All Topics Heatmap	1046
All Topics Table	1051
All Topics Summary	1054
Single Topic Summary	1057
Single EMS Topic-Clients	1061
Single Topic By Server	1064
EMS Queues.....	1066
All Queues Heatmap	1066
All Queues Table	1071
All Queues Summary	1074
Single Queue Summary.....	1078
Single EMS Queue-Clients.....	1082
Single Queue By Server.....	1085
EMS Clients.....	1087
Connections.....	1088
Bridges, Users, Ports.....	1090
Routes	1092
Producers	1096
Producer Summary	1098
Consumers	1101
Consumer Summary	1104
Durables	1107
TIBCO FTL	1109
FTL Servers.....	1110
All Servers Heatmap	1110
All Servers Table	1112
All Group Servers Table.....	1117
All Satellites Table	1118
Server Summary.....	1121

FTL Clients	1124
All Clients Heatmap	1125
All Clients Table	1127
Clients by Group	1130
Client Summary	1132
Client Metrics	1135
FTL Events	1136
Events	1136
TIBCO Hawk	1139
Hawk Agents Table	1139
Hawk Alerts Table	1140
RTView Agent Administration	1142
Chapter 9 - RTView Manager	1145
RTView Manager	1146
Data Servers	1146
Display Servers	1148
Historian Servers	1150
Version Info	1151
JVM	1153
All JVMs	1154
All JVMs Heatmap	1154
All JVMs Table	1155
Single JVM	1158
JVM Summary	1158
JVM System Properties	1162
JVM Memory Pool Trends	1163
JVM GC Trends	1167
Tomcat	1169
Tomcat Servers	1169
All Tomcat Servers	1169
Tomcat Server Summary	1171
Tomcat Applications	1174
Applications Heatmap	1175
Applications Summary	1176
Appendix A - Alert Definitions	1179
Amazon Web Services	1179
Apache Kafka	1181
Docker	1184
Microsoft SQL Server	1185
MongoDB	1186
MySQL Database	1187

Node.js	1188
Oracle Coherence.....	1189
Oracle Database	1196
Oracle WebLogic	1198
RTView Host Agent.....	1202
RTView Manager and RTView Rules	1204
Solace.....	1205
TIBCO ActiveMatrix BusinessWorks.....	1211
TIBCO ActiveSpaces	1216
TIBCO Adapters.....	1218
TIBCO BusinessEvents.....	1219
TIBCO Enterprise Message Service	1221
TIBCO FTL	1226
UX	1228
VMware vCenter	1229
Appendix B - Limitations	1233
iPad Safari Limitations	1233
TIBCO ActiveMatrix BusinessWorks.....	1235
Servers	1235
Business Works 5.7.1 Engine Status.....	1235
BWSE Components	1236
Appendix C - Third Party Notice Requirements	1237
RTView EM.....	1237
RTView Core®	1243

Preface

Welcome to the *RTView Enterprise Monitor® User's Guide*. Read this preface for an overview of the information provided in this guide and the documentation conventions used throughout, additional reading, and contact information. This preface includes the following sections:

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

About This Guide

The *RTView Enterprise Monitor® User's Guide* describes how to use RTView Enterprise Monitor. For information about how to install and configure RTView Enterprise Monitor, see the *RTView Enterprise Monitor Configuration Guide*.

Audience

This guide is written for support teams and managers who are responsible for the performance and health of network and software resources in an organization.

Document Conventions

This guide uses the following standard set of typographical conventions.

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

Additional Resources

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

- [“Release Notes”](#)
- [“SL Documentation”](#)

Release Notes

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

Examine the online release notes before you begin the installation and configuration process. They contain important information about this release of RTView Enterprise Monitor.

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>.

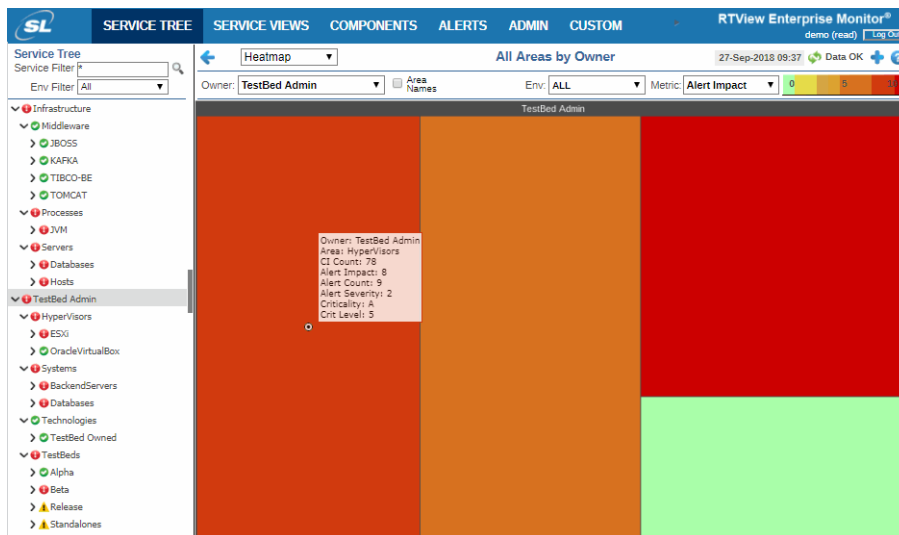
Technical Support

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 RTView Enterprise Monitor

Welcome to RTView Enterprise Monitor®. RTView Enterprise Monitor® is a monitoring system that provides single-pane-of-glass visibility of aggregated real-time and historical information about the performance of complex multi-tier applications, including custom-built applications.



This section contains:

- “Displays”
- “System Requirements”

Note: For details about configuring RTView Enterprise Monitor (also referred to as the *Monitor*), see the *RTView Enterprise Monitor Configuration Guide*.

RTView Enterprise Monitor has the ability to drill-down to the software-component level to help you determine the root cause of issues affecting application performance. RTView Enterprise Monitor (the *Monitor*) enables you to answer questions such as: Are any resources reaching a state of critical health? Do I need to allocate more memory to any resources? Are any having slow response times? Are application deadlocks causing bottlenecks anywhere? Is processing and connection load evenly distributed across resources?

RTView Enterprise Monitor enables application support teams to:


- Provide a single, real-time interface to the end-to-end performance of complex or distributed applications.
- Provide early warning of issues and automate corrective actions tied to alerts, to reduce the number of trouble tickets.
- Leverage historical trends to anticipate possible application degradation and enable preventive care.
- Quickly pinpoint the root cause of issues and initiate repair.
- Reduce costs and minimize lost revenue related to system downtime and degradation.
- Improve performance against SLAs, customer expectations and brand promises.
- Improve business decisions that are tied to application performance.
- Lower the total cost of managing applications.

Displays

RTView Enterprise Monitor® provides multiple sets of displays for monitoring your system. Some displays come with and reside on RTViewCentral. Additional displays can be added on via RTView DataServers.

This section describes “RTViewCentral” and “RTView DataServer” - the functions each provides and the types of displays they provide.

The following figure illustrates “RTViewCentral”, a single “RTView DataServer” and the basic data flow from the monitored data sources.



RTView DataServers collect and store metric data from your data sources. “RTViewCentral” provides the graphic visualization of the metric data collected by RTView DataServers. Performance data collected by RTView DataServers are correlated with the displays that come with RTView Enterprise Monitor.

RTViewCentral

RTViewCentral is where the metric data collected by the RTView DataServer is analyzed, correlated and transformed; historical data is aggregated; alert rules and actions are defined; and where the “master” mapping of everything monitored in your system resides.

RTViewCentral is comprised of a Display Server, the Central Server, the Central Alert Historian and a database. RTViewCentral also has “RTView Manager” displays and RTView RTVRules.

RTViewCentral displays come with RTView Enterprise Monitor and reside on RTViewCentral.

Find these displays organized under the following Views:

- **"All Management Areas"**: These displays show the health of your entire system using aggregated data from all Areas. Use these displays to quickly identify critical conditions across all Areas in your system, then drill-down to investigate in lower-level displays.
- **"Multi Area Service Views"**: These displays show the health of Services for one or more Groups. Use these displays to identify critical conditions across all Areas or a single Area. Drill-down to investigate in lower-level displays.
- **"Single Area Service Views"**: These displays show the health of Services for one or more Groups. Use these displays to identify critical conditions across a single Area. Drill-down to investigate in lower-level displays.
- **"Service Summary Views"**: These displays show the health of CI Types. Use these displays for a closer view of a critical condition, including alert details.
- **"Key Metrics Views"**: These displays show how close a metric is approaching its threshold (rather than your ACTIVE alerts and their impact on the overall application or service), enabling you to anticipate performance problems BEFORE the alert threshold is crossed and analyze the circumstances that led up to error conditions.
- **"Component Views"**: These displays show the lowest level view of CMDB contents--the component level. In these displays, alert states for components are shown by Service and Area in tabular and heatmap formats, while highlighting the most critical alert state for each component.
- **"Metric Explorer"**: The Metric Explorer (MX) is a tool for creating and viewing custom dashboards, referred to as MX Views. An MX View contains a trend graph with up to five traces which you can configure to show numeric metrics from any solution package.
- **"Alert Views"**: These displays show detailed information about all alerts that have occurred in your RTView Enterprise Monitor system (all Owners and all Areas).
- **"Administration"**: These displays enable you to set alert thresholds, and observe how alerts are managed.
- **"CMDB Administration"**: This display allows you to modify your Service Data Model.
- **"Architecture"**: These displays provide a view of RTView Enterprise Monitor component connectivity, mapping between component types, and component level connection and performance information.
- **"Property Views"**: These displays show how your Monitor properties are configured and the values for all connected RTView processes.
- **"Diagram Views"**: These displays are dedicated for the Diagram Generator, a feature that auto-creates diagram displays which mirror your system components and hierarchy.

RTView DataServer

RTView DataServer displays reside on the RTView DataServer. You can add displays to your RTView Enterprise Monitor system by installing one or more RTView DataServers. RTView DataServers contain a series of displays based on the type of technology being monitored. **"Solution Packages"** are bundled into RTView DataServers. For example, the RTView DataServer for TIBCO, which includes the Solution Packages for TIBCO EMS, BusinessWorks, and many other TIBCO applications.

RTView DataServers are also available for IBM, Infrastructure, Kafka, Oracle, Solace and RTView Manager. Performance data collected for these technologies are correlated with the displays that come with RTView Enterprise Monitor.

The following are available:

- **"RTView DataServer for IBM"**: Used to monitor the health and performance across all components for IBM® MQ, IBM DB2 databases and IBM WebSphere servers and applications.
- **"RTView DataServer for Infrastructure"**: Used to monitor the health and performance across all components for Amazon Web Services, Docker, JBoss, MongoDB, MySQL, MS SQL Server, Node.js, VMware, UX (User Experience), and RTView Host Agents.
- **"RTView DataServer for Kafka"**: Used to monitor the health and performance across all components for Apache Kafka.
- **"RTView DataServer for Oracle"**: Used to monitor the health and performance across all components for Oracle Coherence, Oracle Database, Connector for Oracle Enterprise Manager and Oracle WebLogic.
- **"RTView DataServer for Solace"**: Used to monitor the health and performance across all components for Solace routers, bridges, endpoints, clients and Syslog events.
- **"RTView DataServer for TIBCO"**: Used to monitor the health and performance of TIBCO ActiveMatrix, TIBCO ActiveSpaces, TIBCO Adapters, TIBCO BusinessEvents, TIBCO BusinessWorks, TIBCO Enterprise Message Service, TIBCO FTL, and TIBCO Hawk.

Solution Packages

Solution packages gather metrics from infrastructure, middleware, instrumented applications, JVMs, log files, and third party monitoring products. RTView Enterprise Monitor also provides a means for creating custom solution packages to gather most any piece of performance information with a wide array of built-in data adapters. These custom solution packages can be configured without programming. SL Support has many templates for custom solution packages that can be delivered to users or customized as a service.

A solution package provides these main pieces of functionality to RTView Enterprise Monitor:

- **Data Access:** The solution package gathers the performance metrics relevant to the technology being monitored. The data may be gathered by either synchronous or asynchronous direct connections to a technology, or by receiving information from RTView agents deployed on the hosts of the monitored technology.
- **Data Caching:** Performance metrics are stored in in-memory data caches to supply quick access to the most current performance metrics.
- **Data History:** Long-term performance metrics can be stored in a JDBC-enabled relational database. The solution package allows for the configuration of the rules for data compaction and management of long-term data persistence.
- **Alert Event Access:** If the solution package is connecting to another monitoring system, it can gather alert events from that system, bring those events into RTViewCentral and allow alert management to be performed in RTViewCentral. Optionally, the solution package can be configured to synchronize alert states between the two systems.
- **Alert Rules Engine:** The solution packages are configured with alert rule definitions which are processed real-time in the RTView DataServers. Dynamic updates to these alert rule definitions, such as changing alert rule thresholds or policies, can be managed through the RTViewCentral **Alert Administration** interface. When alerts are activated by these alert rule definitions, they are sent to RTViewCentral to be aggregated with other solution package alerts.
- **Data Viewing:** Each solution package comes with designated displays which can be accessed by RTView Enterprise Monitor to show the performance metrics in summary and drill-down views.
- **Data Server:** This Java process is run to begin accessing the data, storing data to internal memory caches, running the alert rules and optionally providing data to the Historian process.
- **Data Historian:** The process manages the storage of information into a relational database and runs the rules relevant to managing this persisted data.

System Requirements

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

CHAPTER 2 Using the Monitor

Welcome to the RTView Enterprise Monitor. This section describes how to access the Monitor, RTView Manager, the user interface and the many displays available.

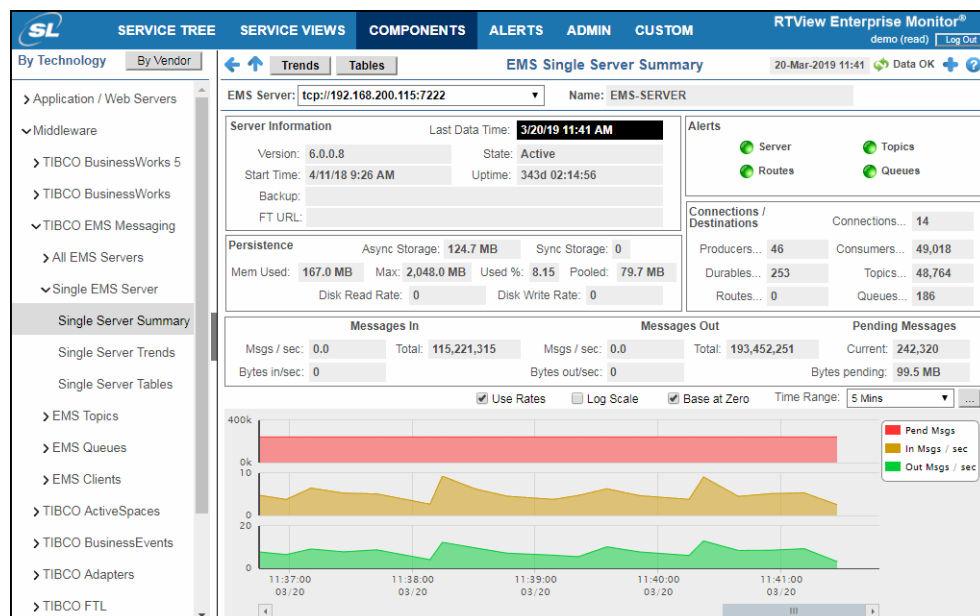
This section contains:

- **"Overview"**: Describes the user interface and navigation, behavior of graphic objects, color codes and icons.
- **"RTViewCentral Displays"**: Describes the displays that come with RTView Enterprise Monitor.

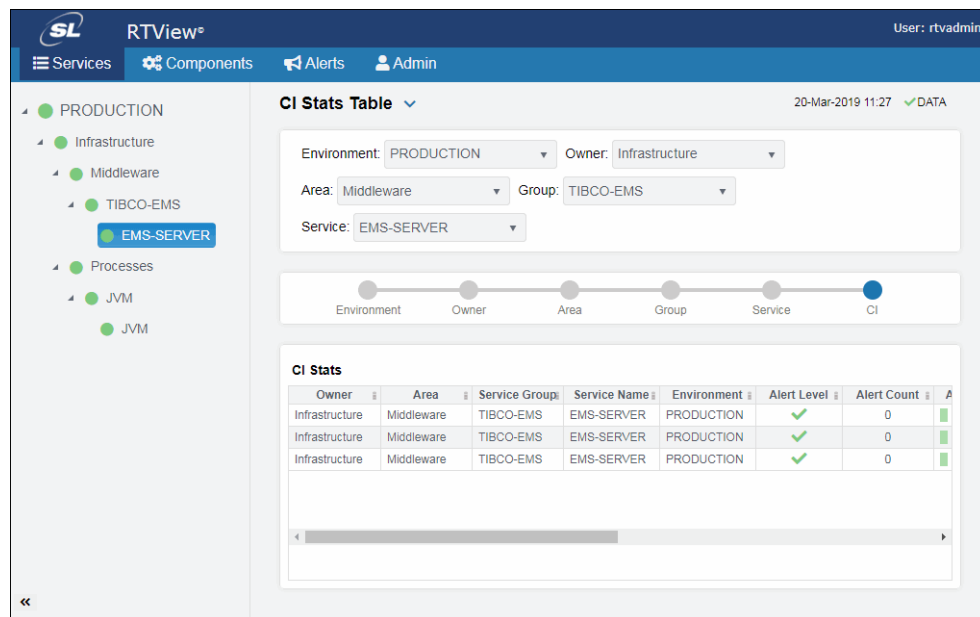
To access RTView Enterprise Monitor:

- Set the `JAVA_HOME` environment variable to the location of your Java installation.
- Execute the **start_servers** script, located in the **RTViewCentral/bin** directory, to start RTView Enterprise Monitor.
- Browse to **http://localhost:10070/rtview-central-classic** and login (use admin/admin)

Note: The first time you open the Monitor displays are visible but not yet populated with monitoring data. After you download and configure one or more RTView DataServers, configure data collection *and also* integrate with RTView Enterprise Monitor, the displays populate with your monitoring data.



Note: Alternatively, you can open the Beta version of the new HTML interface for RTView Enterprise at: **http://localhost:11070/rtview-central** (login as rtvadmin/rtvadmin or rtvuser/rtvuser).



- **"RTView Manager"**: Describes the displays used for monitoring your RTView Enterprise Monitor deployment.

To access RTView Manager:

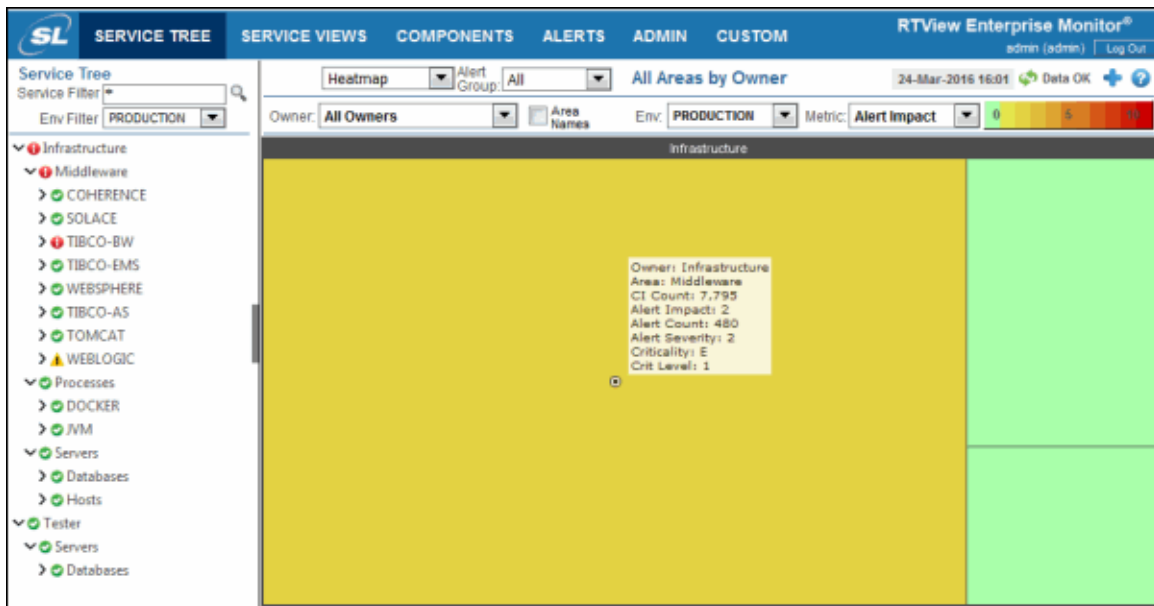
Browse to one of the following and login (use rtvadmin/rtvadmin):

http://<ip_address>:3070/rtview-manager-classic if you are running the monitor remotely.

http://localhost:3070/rtview-manager-classic if you are running the monitor locally.

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

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



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

Overview

RTView Enterprise Monitor uses visual cues (such as color coding, graphic charts and sizing of shapes) to communicate the current state of all elements in your system. This section describes how displays are structured and organized, how to read heatmaps, tables and trend graphs, as well as GUI functionality and navigation.

This section includes:

- [“Navigation Tabs”](#)
- [“Fundamental Structure of Displays”](#)
- [“Heatmaps”](#)
- [“History Heatmaps”](#)
- [“Tables”](#)
- [“Trend Graphs”](#)
- [“Popup Menu”](#)
- [“Title Bar”](#)

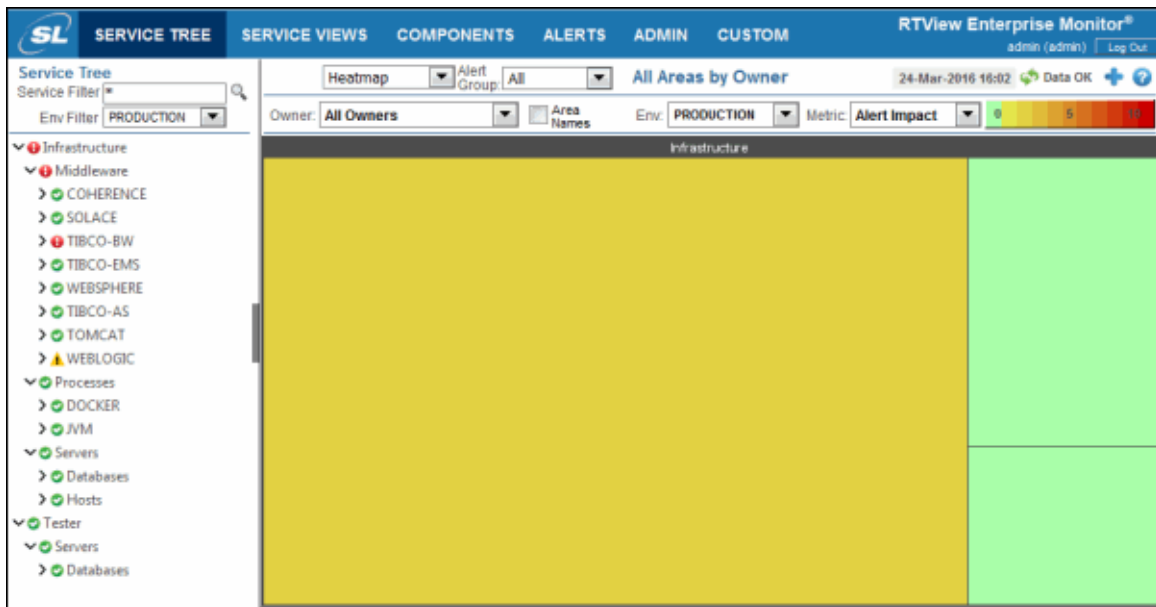
Navigation Tabs



There are six different tabs that run along the upper portion of the window:

- **"SERVICE TREE Tab"**: provides a tree structure view of your defined CMDB with 4 levels of hierarchy: **Owner>Area>Group>Service**. The tree is configurable and should be set up to represent how a support person likes to conceptually think of the vast number of items that are being monitored.
- **"SERVICE VIEWS Tab"**: provides an alternate way of accessing the primary RTView Enterprise Monitor displays also found in the **SERVICE TREE** tab. This tab might be used by power users who are very familiar with their monitoring environment and choose not to visualize the entire service tree.
- **"COMPONENTS Tab"**: provides access to the **JVM Process Views**, the **Tomcat Servers Views**, the **RTView Servers Views**, and any Views included with the solution packages that you have installed. This tab organizes the monitoring information by technology or vendor and allows you to view the health state of your technology footprint without logical or service groupings. Specialists that understand in some depth how the technologies are deployed, configured, load-balanced, and scale based on load will gain benefit on the organization of performance metrics by the important functional sub-components of the technology.
- **"ALERTS Tab"**: provides a view of the current active alerts in the system and allows you to manage those alerts by owning them, acknowledging them, and/or suppressing them.
- **"ADMIN Tab"**: can be accessed by administrators of RTView Enterprise Monitor, who can use this tab during installation to set up proper alert settings, to describe logical and service groupings that drive the construction of the Service Tree, and to "monitor the monitor" view of the current health state of RTView Enterprise Monitor and how it is currently deployed and configured.
- **"CUSTOM Tab"**: provides a location where you can add your own tab and views.


SERVICE TREE Tab



The **SERVICE TREE** tab provides a tree structure view of your defined CMDB with 4 levels of hierarchy: **Owner>Area>Group>Service** (see [“Fundamental Structure of Displays”](#) for more information). This tab is the primary source for understanding the health of your services and for drilling down to analyze issues. The Service Tree, which is configurable, shows user-defined logical groupings of the infrastructure and middleware used to support applications and should be set up to represent how support personnel like to conceptually think of the vast number of items that are being monitored. These groupings could, for example, contain a collection of monitored Configuration Items used to support a specific application or a service, or they could contain Configuration Items relevant to varying technologies located at specific data centers. The Service Tree aggregates the current alert state of any item in a group to indicate which groups have items that need to be investigated, and you can use a variety of visual clues to prioritize and analyze the issues. You can also determine priority using the Alert Impact view in the heatmaps to identify which alert conditions will be the most impactful to your business, and you can then analyze the situation using a variety of tools including:

- **Key Metrics:** allows you to view the cross-correlation of Configuration Items relevant to a grouping or service and how their performance may affect each other and the services they support. For details, see [“Key Metrics Views”](#).
- **Drill Down CI Summary Views:** provides a way to analyze how a particular Configuration Item has been performing over time.
- **Metric Explorer:** allows you to choose specific metrics to chart when analyzing several critical performance metrics over time. For details, see [“Metric Explorer”](#).


This tab allows you to filter the navigation tree content by service and environment (see figure below). The environment you select also sets the **Environment** filter on the main panel. Note that changing the **Environment** filter on the main panel does not set the **Environment** filter in the navigation panel.




Alerts

Each level within the Service Tree has a red, yellow, or green icon next to it, which indicate the highest alert level for that particular Owner, Area, Group, or Service. These icons allow you to instantly recognize problem areas within your system and allow you to drill down to quickly find the source of the issue. A red icon indicates that one or more alerts exceeded their ALARM LEVEL threshold, a yellow icon indicates that one or more alerts exceeded their WARNING LEVEL threshold, and a green icon indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold.

The Owner, Area, and Group automatically display the icon representing the highest level alert for their associated services with red (ALARM LEVEL threshold exceeded) being the most serious, yellow (WARNING LEVEL threshold exceeded) being intermediate, and green meaning everything is functioning normally. For example, if any of the services within a particular **Owner>Area>Group** have one or more alerts that exceeded their ALARM LEVEL threshold and, hence, have a red icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same red icon. In the example below, you can see that the MQ Broker service has one or more alerts that exceeded their ALARM LEVEL threshold and has a red indicator. As a result, the Owner, Area, and Group also have the red indicator



If the highest alert level for the services within a particular **Owner>Area>Group** is a service that has one or more alerts that exceeded their WARNING LEVEL threshold and, hence, has a yellow icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same yellow icon. In the example below, you can see that the DB2 database has one or more alerts that exceeded its WARNING LEVEL threshold and has a yellow indicator. Since none of the other services in this particular tree have alerts that exceeded their ALARM LEVEL threshold, then the associated Owner, Area, and Group also have the yellow indicator since the WARNING LEVEL threshold is the highest alert level threshold exceeded.



Available Displays

The following displays are available in the following levels in this tab:

Owner Level (top level)

To access the following displays, select an Owner Level option (**Infrastructure**, for example) to display an Owner level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

Drop-down Option	Display	Description
Heatmap	"Area Heatmap"	Heatmap of the most critical alerts for all Areas of your system, with the option to filter by Owner, Environment and alert Metric.
Area	"Area Table"	Table of data shown in the "Area Heatmap" with the option to filter by Owner and Environment.

Note: When selecting an Owner Level option, the display that opens by default will be the one that was last selected. For example, if Heatmap was the display that was previously selected, Heatmap will display by default again.

Area Level (second level down)

To access the following displays, select an Area Level option (**Middleware** in the example above) to display an Area Level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

Drop-down Option	Display	Description
By Group	"Group/Service Heatmap"	Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
By Region	"Group/Region Heatmap"	Heatmap as described for the Group / Service Heatmap (above), with the option to filter by Region and no option to show Service Names.
Table	"Group / Service Table"	Table of data shown in the "Group/Service Heatmap".
By CI Type	"Services CI Type Summary"	Table that shows the health state of Services per CI Type.
History	"Services History Heatmap"	Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

Note: When selecting an Area Level option, the display that opens by default will be the one that was last selected. For example, if Group/Service Heatmap was the display that was previously selected, Group/Service Heatmap will display by default again.

Group Level (third level down)

To access the following displays, select a Group Level option (IBM-MQ in the example above) to display a Group Level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

Drop-down Option	Display	Description
By Group	"Single Area: Group/Service Heatmap"	Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
By Region	"Single Area: Region/Service Heatmap"	Heatmap as described for the Group / Service Heatmap (above), with the option to filter by Region and no option to show Service Names.
Table	"Single Area: Region/Service Heatmap"	Table of the data shown in the "Single Area: Group/Service Heatmap".
By CI Type	"Single Area: Services CI Type Summary"	Table that shows the health state of Services per CI Type.
History	"Single Area: Services History Heatmap"	Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

Note: When selecting a Group Level option, the display that opens by default will be the one that was last selected. For example, if Group/Service Heatmap was the display that was previously selected, Group/Service Heatmap will display by default again.


Service Level (fourth level down)

To access the following displays, select a Service Level option (MQ Broker (DEMOSITE) in the example above) to display a Service Level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

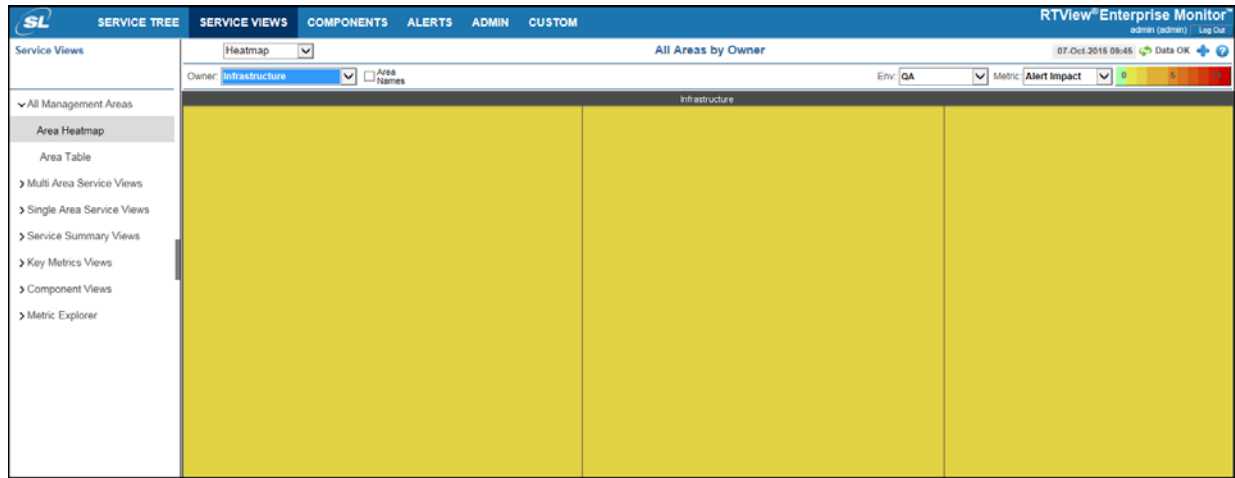
Drop-down Option	Display	Description
By CI Type	"Service By CI Type"	Table of alert states for a Service organized CI Type, with general alert information.
Summary	"Service Summary"	Table of CIs by Service, with detailed alert information.
Health	"Service Health Heatmap"	Heatmap of CIs by Service, with the option to filter by Owner, Area, Group, Environment and alert Metric, and show CI Names.
KM Heatmap	"Service KM Heatmap"	Heatmap of Key Metrics current data for one or more Services in your CMDB hierarchy.
KM Table	"Service KM Table"	Table of Key Metrics current data for one or more Services.
KM History	"Service KM History"	History heatmap of Key Metrics historical data for one or more Services.
KM History (Alt)	"Service KM History (Alt)"	History heatmap of Key Metrics historical data for one or more Services.

Note: When selecting a Service Level option, the display that opens by default will be the one that was last selected. For example, if By CI Type was the display that was previously selected, By CI Type will display by default again.

Select the following button, which is available when you select either **By CI Type (Service By CI Type display)** or **Summary (Service Summary display)** from the drop-down list, to open the associated display:

Button	Display	Description
	"Metric Explorer"	The Metric Explorer (MX) is a tool for creating and viewing custom dashboards, referred to as MX Views.

SERVICE VIEWS Tab



The **SERVICE VIEWS** tab is a simplified version of the **SERVICE TREE** tab that uses drop-down navigation to access displays without the complexity of the service tree. This tab contains the following Views:

- **"All Management Areas"**: Displays in this View show the health of your entire system using aggregated data from all Areas. Use these displays to quickly identify critical conditions across all Areas in your system, then drill-down to investigate in lower-level displays.
- **"Multi Area Service Views"**: Displays in this View show the health of Services for one or more Groups. Use these displays to identify critical conditions across all Areas or a single Area. Drill-down to investigate in lower-level displays.
- **"Single Area Service Views"**: Displays in this View show the health of Services for one or more Groups. Use these displays to identify critical conditions across a single Area. Drill-down to investigate in lower-level displays.
- **"Service Summary Views"**: Displays in this View show the health of CI Types. Use these displays for a closer view of a critical condition, including alert details.
- **"Key Metrics Views"**: The Key Metrics (KM) feature shows how close a metric is approaching its threshold (rather than your ACTIVE alerts and their impact on the overall application or service), enabling you to anticipate performance problems BEFORE the alert threshold is crossed and analyze the circumstances that led up to error conditions.
- **"Metric Explorer"**: The Metric Explorer (MX) is a tool that allows end-users to quickly create custom dashboards for metrics they specifically want to analyze.

COMPONENTS Tab

The **COMPONENTS** tab organizes the monitoring information by technology or vendor and allows you to view the health state of your technology footprint without logical or service groupings. This tab also contains deep summaries and drill-downs to the subcomponents that comprise a particular technology. By default, this tab provides access to the **JVM Process Views**, the **Tomcat Servers Views**, the **RTView Servers Views**, and any Views included with the solution packages that you have installed.

NOTE: The COMPONENTS tab navigation tree (in left panel) has been enhanced to filter the available displays based on the solution packages that are hosted by the RTView DataServer connections enabled by your administrator. If you run a project from a previous release where the navigation tree was customized, this filter is applied in addition to any customization.

The following views are available via this tab:

- **"JVM Processes View/Displays"**: Displays in this View show performance data for monitored Java Virtual Machine (JVM) Processes. Use these displays to monitor performance of your JVMs.
- **"Tomcat Monitor Views/Displays"**: Displays in this View show performance data for monitored Tomcat applications. Use these displays to monitor Tomcat connections and performance of your Web applications and modules.

There are two different ways to view the available displays: **By Technology** and **By Vendor**.

By Technology Button

The **By Technology** button lists the available displays by the type of technology (Application/ Web Servers, Middleware, Databases, Processes, Hosts/VMs, Connectors, Other).

SL

SERVICE TREE

SERVICE VIEWS

COMPONENTS

ALERTS

ADMIN

CUSTOM

By Technology

By Vendor

Heatmap

All JVMs - Table View

23-Sep-2015 17:02

Data OK

> Application / Web Servers

> Middleware

> Databases

> Processes

> JVM Processes

> All JVMs

> All JVMs Heatmap

> All JVMs Table

> Single JVM

> RTView Processes

> Hosts / VMs

> Connectors

> Other

JVM Count: 51

Show Inactive

All JMX Connections

Connection	Expired	Connected	Alert	Host	Port	CPU %	Max Heap	Mem Used %	Display Name	URL	RtvAppTy...
ALERT_SERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10023	18.6	492,896,256	55.7			3 local
ALERTHISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10025	0.6	477,233,152	4.1			1 local
AMXMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3367		0				0 local
AMXMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	6368	2.0	954,466,304	37.8			3 local
AMXMON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	6368	2.0	954,466,304	31.7			3 local
BW6MON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3368	0.9	954,466,304	20.2			3 local
BW6MON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3368	1.0	954,466,304	20.2			3 local
BWMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3367		0				0 local
BWMONITOR-WIN-8	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.138	3368		0				0 local
CONFIG_SERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10013	2.4	477,233,152	34.9			3 local
DISPLAYSERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10024	4.0	477,233,152	62.9			5 local
DISPLAYSERVER_DARK...	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10124	2.5	477,233,152	29.9			5 local
EMSMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3167		0				0 local
EMSMONITOR-WIN-8	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.138	3168	1.3	954,466,304	28.6			3 local
EMSMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3168	1.9	954,466,304	17.1			3 local
EMSMON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3168	1.6	954,466,304	20.4			3 local
local	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			1.8	954,466,304	12.8		local	3 local
MISCMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3967		0				0 local
MISCMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3968	13.0	1,071,316,992	95.4			3 local
MISCMON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3968	5.3	985,661,440	64.4			3 local
MQMON-64-OL7-3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.73	3468	4.2	1,037,959,168	9.4			3 local
MQMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3467		0				0 local
MQMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3468	3.7	954,466,304	35.6			3 local
OCMON-64-OL7-1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.71	9911		0				0 local
OCMON-64-OL7-4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.74	9911	0.4	954,728,448	1.6			3 local
OCMONITOR-WIN-8	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.138	9911		0				0 local
OCMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	9911	3.8	954,466,304	27.6			3 local
OCMON-SLHOST-WIN7	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.137	9911		0				0 local
RTVMGR-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3067		0				0 local
RTVMGR-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3068	1.7	954,466,304	10.9			3 local
RTVMGR-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3068	1.8	954,466,304	12.8			3 local
RTVRULES	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3868	0.5	715,849,728	10.2			3 local
RTVRULES-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3868	0.5	715,849,728	18.1			3 local
SOLMON-64-OL7-6	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.76	4168	1.5	954,728,448	47.7			3 local

By Vendor Button

The **By Vendor** button lists the displays by vendor name (for example: TIBCO, Oracle, and IBM).

[illegible]

ALERTS Tab

The screenshot shows the RTView Enterprise Monitor Alerts tab. The top navigation bar includes 'SERVICE TREE', 'SERVICE VIEWS', 'COMPONENTS', 'ALERTS', 'ADMIN', and 'CUSTOM'. The 'ALERTS' tab is active. On the left, the 'Service Tree' shows a hierarchy: Infrastructure (red icon), Middleware (red icon), COHERENCE (green icon), SOLACE (green icon), TIBCO-AS (green icon), TIBCO-BW (red icon), TIBCO-EMS (green icon), Processes (green icon), JVM (green icon), Servers (green icon), and Hosts (green icon). The main panel displays the 'Alerts Table' for the 'Current' alert group. It includes filters for 'Field Filter', 'Search Text', 'CMB Filter', and 'Alert Group'. The table shows a list of alerts with columns: Total (197 / 197), Critical (197 / 197), Warning (0 / 0), Suppressed (0 / 0), First Occ, Last Occ, Count, Sup, Owner, Alert Name, Primary Service, and CI. The table lists multiple alerts for 'BwActivityExecutionTimeHi' and 'BwProcessAbortRateHigh'.

The **ALERTS** tab provides a view of the current active alerts in the system and allows you to manage those alerts by owning them, acknowledging them, and/or suppressing them. You can navigate and filter the alert list by using the service tree to focus on alerts by logical or service groupings. This tab is customizable and can be interfaced with an existing trouble ticket system so that alerts that require an action can be tracked and managed by those systems.


This tab allows you to filter the navigation tree content by service and environment (see figure below). The environment you select also sets the **Environment** filter on the main panel. Note that changing the **Environment** filter on the main panel does not set the **Environment** filter in the navigation panel.




Alert Icons

Each level within the Alerts tab service tree has a red, yellow, or green icon next to it, which indicate the highest alert level for that particular Owner, Area, Group, or Service. These icons allow you to instantly recognize problem areas within your system and allow you to drill down to quickly find the source of the issue. A red icon indicates that one or more alerts exceeded their ALARM LEVEL threshold, a yellow icon indicates that one or more alerts exceeded their WARNING LEVEL threshold, and a green icon indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold.

The Owner, Area, and Group automatically display the icon representing the highest level alert for their associated services with red (ALARM LEVEL threshold exceeded) being the most serious, yellow (WARNING LEVEL threshold exceeded) being intermediate, and green meaning everything is functioning normally. For example, if any of the services within a particular **Owner>Area>Group** have one or more alerts that exceeded their ALARM LEVEL threshold and, hence, have a red icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same red icon. In the example below, you can see that the MQ Broker service has one or more alerts that exceeded their ALARM LEVEL threshold and has a red indicator. As a result, the Owner, Area, and Group also have the red indicator



If the highest alert level for the services within a particular **Owner>Area>Group** is a service that has one or more alerts that exceeded their WARNING LEVEL threshold and, hence, has a yellow icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same yellow icon. In the example below, you can see that the DB2 database has one or more alerts that exceeded its WARNING LEVEL threshold and has a yellow indicator. Since none of the other services in this particular tree have alerts that exceeded their ALARM LEVEL threshold, then the associated Owner, Area, and Group also have the yellow indicator since the WARNING LEVEL threshold is the highest alert level threshold exceeded.



Available Displays

To access the following displays, select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

Drop-down Option	Display	Description
Current	"RTView Alerts Table"	This display allows you to track and manage all alerts that have occurred in the system, as well as to add comments, acknowledge, or assign Owners to alerts.
History	"Alert History Table"	This display allows you to track the history of any alert that has occurred in your RTView Enterprise Monitor system.

Note: When selecting an option at any level, the display that opens by default will be the one that was last selected. For example, if History was the display that was previously selected, History will display by default again.

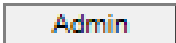
When you select an option at any of the **Owner>Area>Group>Services** levels in the **RTView Alerts Table** display, the display automatically filters the list of alerts based on the level you selected. For example, if you were to select the IBM-MQ option at the Group level, then the filter will be set to Owner=Infrastructure, Area=Middleware, Group=IBM-MQ (the option at the level you selected), and Service and Environment will be set to * (or all services and environments for that particular Group).

If you were to select the Middleware option at the Area level, then the filter will be set to Owner=Infrastructure, Area=Middleware (the option at the level you selected), and Group and Service and Environment will be set to * (or all groups, services, and environments for that particular Area).

If you were to select the Middleware option at the Area level, then the filter will be set to Owner=Infrastructure, Area=Middleware, Group=IBM-MQ, Service=MQ-BROKER, and Environment=DEMOSITE (the option at the level you selected).

Available Display via a Button

If you select the **Current** option from the drop-down list, the following button is available on the **RTView Alerts Table** display. Select the following button to open the associated display:

Button	Display	Description
	"Alert Administration"	This display allows you to set global or override alert thresholds.

ADMIN Tab

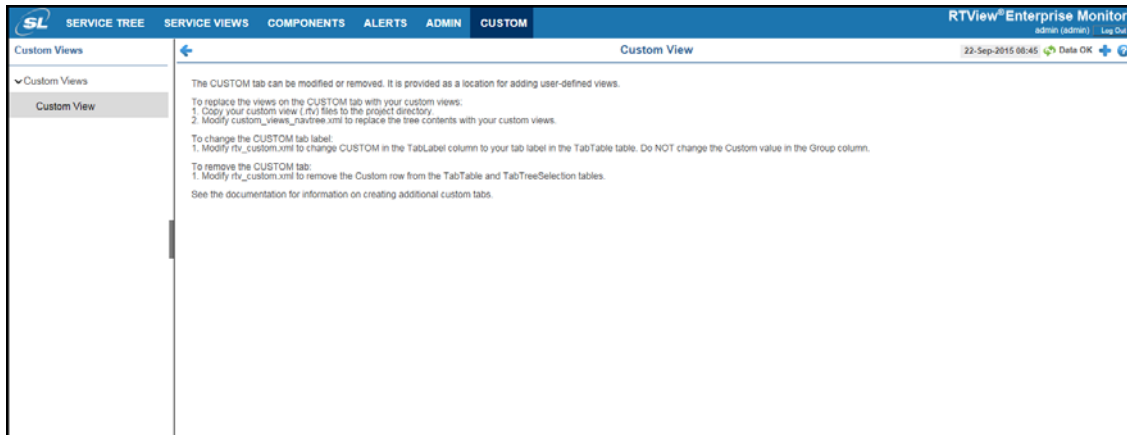
The screenshot displays the **ADMIN** tab in the RTView Enterprise Monitor interface. The main area is titled **Alert Administration** and shows a table of configured alerts. The table has the following columns: **Alert**, **Warning Level**, **Alarm Level**, **Duration**, **Alert Enabled**, and **Override Count**. The table lists various system alerts such as **AcwInstanceCpuHigh**, **AcwInstanceDiskReadBytesHigh**, and **AmxServiceHitRateHigh**, each with specific threshold values and a duration of 30 seconds. Most alerts are marked as **Alert Enabled** (checked box) and have an **Override Count** of 0.

Below the table, the **Settings for Selected Alert** section is visible, containing input fields for **Name** (with a dropdown to select an alert from the table), **Description**, **Warning Level**, **Alarm Level**, **Duration (Secs.)**, and an **Enabled** checkbox. A **Save Settings** button is located at the bottom right of this section.

The **ADMIN** tab can only be accessed by administrators of RTView Enterprise Monitor, who can use this tab during installation to set up proper alert settings, to describe logical and service groupings that drive the construction of the Service Tree, and to “monitor the monitor” view of the current health state of RTView Enterprise Monitor and how it is currently deployed. This tab provides access to the **Alert Administration**, **CMDB Administration**, and **Architecture** Views. See the following sections for more information:

- **“Administration”**: Displays in this View allow you to set alert thresholds, track alert management, and modify your Service Data Model.
- **“CMDB Administration”**: Use this display to setup, view, or modify your Service Data Model (CMDB), including: adding, renaming, deleting or merging your CMDB hierarchical elements (Owners, Areas, Groups or Services), associating CIs with Services and assigning or modifying CI attributes (such as Criticality).
- **“Architecture”**: Displays in this View show RTView Enterprise Monitor system information such as a topological view of your components and their connection state, configuration definitions and mapping, and performance metrics for your Cache Tables and Data Servers.
- **“Property Views”**: Use this display to see how your properties are configured and the values for all connected RTView processes.

CUSTOM Tab



The **CUSTOM** tab provides a location where you can add your own custom tab and views, and create diagram displays. See the following sections for more information:

- **Modify the CUSTOM Tab**
- **“Diagram Views”**: Provides the Diagram Generator, which enables you to auto-generate a topology view of your system components.

Fundamental Structure of Displays

To interpret RTView Enterprise Monitor displays it is helpful to understand the Service Data Model. The Service Data Model, also referred to as the CMDB, is a database that forms the fundamental structure of all RTView Enterprise Monitor displays, and enables data aggregation and filtering.

The Service Data Model has a four level hierarchy which is, from the highest level (Owner) to the lowest level (Service):

- Owner
- Area
- Group
- Service

The Service Data Model maps all the Configuration Items (CIs) in your RTView Enterprise Monitor system to one or more Services (CIs are items being monitored by RTView Enterprise Monitor--servers, processes and so forth--anything that can be configured). Each Service is mapped to a Group, each Group to an Area and each Area to an Owner. Displays are organized and populated with data according to this hierarchy. This mapping enables RTView Enterprise Monitor to aggregate data for several hundreds of CIs, and allows objects (heatmaps, tables and so forth) to filter data shown according to user selections.

For details about the configuring the Service Data Model, see the Configure Service Data Model section.

Heatmaps

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

Heatmaps scale color for a given metric according to the following rules and are applied in the following order:

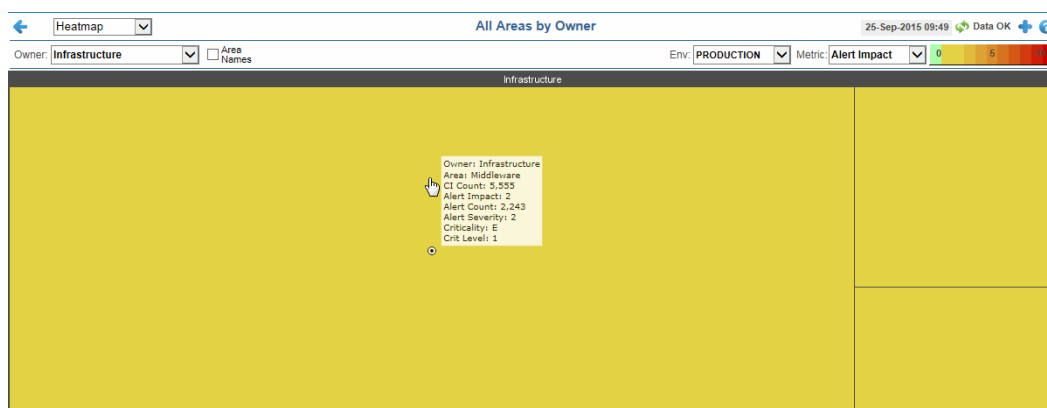
- a) If the metric is associated with an alert, then the color range is scaled from zero to the metric's high Alarm Level threshold, and the color will be red for values near the alarm threshold.
- b) If the metric is not associated with an alert, but the metric is bounded (for example, the **CPU %** utilization value must be in the **0 to 100%** range), then the color is scaled using the user-specified maximum value for the metric.
- c) Otherwise, the metric is autoscaled into a color range from white (minimum) to green (high) using the current highest metric value observed over the monitored entities.


By default, the metric is linearly scaled to an appropriate color. If the **Log** checkbox is checked, then the selected color reflects the logarithm of the current metric value.

Heatmaps include drop-down menus to filter data by Owner, Area, Group, Service, Region and Environment. The filtering options vary among heatmaps.

For example, the **All Management Areas - "Area Heatmap"** (shown in the following figure) illustrates a typical RTView Enterprise Monitor heatmap. The heatmap contains a **Metric** drop-down menu with options to show **Alert Impact**, **Alert Severity**, **Alert Count** and **Criticality** (menu options vary according to the data populating the heatmap). **Alert Impact** is selected and its corresponding color gradient bar is shown. Each rectangle represents all CIs in an Area. The red rectangle in the heatmap indicates that one or more CIs in that Area currently has an alert in an alarm state. The yellow rectangles in the heatmap indicate that one or more CIs in those Areas currently have an alert in a warning state. A green rectangle would indicate that no alert is in a warning or alarm state in an Area.

Continuing with our example, there are two filtering options. You can choose to show all Owners or a single Owner, and all Environments or a single Environment. Each rectangle represents an Area. The rectangle size represents the number of CIs in the rectangle; a larger size is a larger value. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. The following figure illustrates the mouse-over feature in which we see all the **Metric** drop-down values.



In most heatmaps, you can also drill-down to more detail by clicking a rectangle in the heatmap. Or, click Open New Window  and then drill-down. The drill-down opens a display that contains relevant and more detailed data.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.








Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

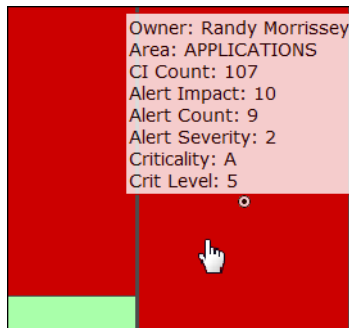
Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

Alert Impact	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from 0 - 10 , as indicated in the color gradient  bar, where 10 is the highest Alert Impact.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	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.
Criticality	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from 1 to 5, as indicated in the color gradient  bar, where 5 is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - "CI / Service Table" display, which range from A to E, where A is the highest Criticality (level 5 maps to a Criticality of A and level 1 maps to a Criticality of E with equally spaced intermediate values).</p>


Mouse-over

The mouse-over functionality provides additional detailed data in an over imposed pop-up window when you mouse-over a heatmap. The following figure illustrates mouse-over functionality in a heatmap object. In this example, when you mouse-over a host, details are shown such as **CI Count**, **Alert Impact**, **Alert Severity**, and **Criticality**.



History Heatmaps

History heatmaps, such as the following cache heatmap, show you utilization trends, over time, for an entire Oracle Coherence 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 time-stamp, 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.

Tables

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

All Properties			
Order	File Name	Property Name	Property Value
2		sl.rtvview.013420134501337013390134401302	01350013480135101335
2	Sort Ascending	sl.rtvview.cache.config	adaptris_cache_source.rtv \$adap
2	Sort Descending	sl.rtvview.jmx.jmxconn	demoAdapter - - URL:service:jmx
2		rtvapm_package	adaptrismon
2	Columns	rtv_proctag	ADAPTRISMON
2		rtv_title	ADAPTRISMON Package
2	Filter	sl.rtvview.sql.sqlldb	ALERTDEFS sa - jdbc:hsqldb:hs
2		sl.rtvview.cp	C:/rtvdemos/mysql-connector-jav
2	Settings	historian.sl.rtvview.historian.driver	com.mysql.jdbc.Driver
213	rtvview	historian.sl.rtvview.historian.url	jdbc:mysql://192.168.200.42:3306
215	rtvview	historian.sl.rtvview.historian.password	my-secret-pw

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

Additional features are:

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

Multiple Column Sorting

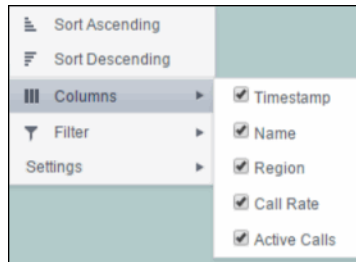
Click on a column header to sort the table by that column. On the first click, the column is sorted in ascending order (smallest value at the top), on the second click the sort is in descending order, and on the third click, the column is returned to its original unsorted state. A sort on a string column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

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

The grid's row selection is cleared if the sort is changed or if columns are resized or reordered. Column sorting is reflected in an export to HTML and Excel.

Column Visibility

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



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


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

Column Filtering

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

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

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



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

- **String columns:** You can enter a filter string such as "abc" and, from the dropdown list, select the operator (equal to, not equal to, starts with, contains, etc) to be used when comparing the filter string to each string in the column. All of the filter comparisons on strings are case-insensitive. You can optionally enter a second filter string (e.g. "xyz") and specify if an AND or OR combination should be used to combine the first and second filter results on the column.

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

The numeric and boolean filter dialogs are shown below.

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


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

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

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

Column Locking

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



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

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


Column Reordering

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

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

Saving Settings

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



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

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

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

Revert Table Settings

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

Row Paging

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

217	emreference	sl.rtvew.sql.sqldb	RTVMGR_FORT_1001-my-secret-pw-joe.mysql/102
229	emreference	sl.rtvew.sub	\$rtvConfigDataServer.CONFIG_SERVER
216	emreference	sl.rtvew.properties.queryTimeOut	10
		sl.rtvew.sql.sqldb	ALERTDEFS --- _none ---

Page 1 of 2 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	10	10
JVM	localhostMYDEMO_DATASERVER	8	8
JVM	localhostMYDEMO_DISPLAYSERVER	8	8
JVM	slidemos.com213415_RTVDB	10	10
JVM	localhostRTVMGR_DB-1	5	5
WAS	SLHOST12Node01Cell SLHOST12Node01.server1	5	5
JVM	localhostRTVMGR_DATABASE	5	5
JVM	localhostRTVMGR_DATASERVER	0	0
JVM	localhostWLM_DATABASE	0	0
EMS	tcp:SLHOST10.7021	0	0
EMS	tcp:SLHOST10.7020	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.
- 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.rtvview.sql.dbretry
9	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.global
10	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.global
11	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.xml.xmlsource
12	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.jmx.jmxconn
13	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.dsenable

Trend Graphs

Trend graphs enable you to view and compare various important metrics over time, such as server memory utilization, server throughput, the number of clients being served by the server, or the total amount of data sent to clients. You can use trend graphs to assess utilization and performance trends.

For example, the following figure illustrates a typical trend graph. In this example, metrics for **Pending Messages**, **Incoming Messages** and **Outgoing Messages** are traced.



By default, the time range end point is the current time. To change the time range for the trend graph click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, **Apr 26, 2012 5:01 PM**. Click **Apply**. Use the Navigation Arrows to move forward or backward one time period (the time period selected from the **Time Range** drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.


Mouse-over


The mouse-over functionality provides additional detailed data in an over imposed pop-up window when you mouse-over trend graphs. The above figure illustrates mouse-over functionality. In this example, when you mouse-over a single dot, or data point, in the **Out Msgs / sec** trend graph, a pop-up window shows data for that data point. In this case, the X-axis value is **13:15:29 hours on September 6th**, and the Y-axis value is **22 Outbound messages per second**.




Log Scale

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

Time Range


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



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

Popup Menu

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

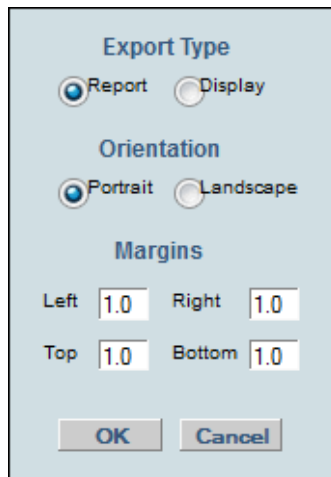


Export PDF Report

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

To generate a report for a display:

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




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

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

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

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



Title Bar

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




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



Opens the previously open display.



Opens the display that is up one level.




Navigates to a display that is most commonly accessed from the current display. The target display differs among displays.




Navigates to displays that are most commonly accessed from the current display. The drop-down menu options differ among displays.



Opens the Alerts Table display in a new window.



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



The data connection state. Red indicates the data source is disconnected (for example, if the Data Server is not receiving data, or if the Display Server does not receive data from the Data Server, this will be red). Green indicates the data source is connected. When the date and time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.



Opens an instance of the same display in a new window. Each window operates independently, allowing you to switch views, navigate to other displays in RTView EM, and compare server performance data. For illustration, see **Multiple Windows**.




Opens the online help page for the current display.

6,047

The number of items (for example, CIs or Areas) in the display.

Area Count: 9

Multiple Windows

The following illustrates the usage of the Open New Window  to open multiple windows.

The screenshot illustrates the usage of the Open New Window button (plus icon) to open multiple windows. The main window displays the 'All Areas by Owner' and 'All Services Status History' views. A secondary window, titled 'Single Service Summary', is overlaid on the main window, showing details for the service 'IRIS-SCAN-DFW'. This window includes a table of CI types and their status, along with a 'Selected CI' section.

CIType	CName	Quality	Severity	Alerts	Criticality	Impact	Region
EMS-QUEUE	tcp://MIRIS1034:7222:SCAN-QUEUE	0	0	0	C	0	AMER
EMS-QUEUE	tcp://MIRIS1034:7222:WORK-QUEUE	0	0	0	C	0	AMER
EMS-QUEUE	tcp://MIRIS1075:7222:REPORT-QUEUE	0	0	0	C	0	AMER
EMS-SERVER	tcp://MIRIS1034:7222	0	0	0	C	0	AMER
JVM	localhost:VMIRIS1034-SCAN-JVM-DFW	0	0	0	C	0	AMER
VMWARE-HOST	vSphereS:esxi-1.south	0	0	0	C	0	AMER
VMWARE-VM	vSphereS:VMIRIS1034	0	0	0	C	0	AMER
WLS	Domain-SOUTH:WLS-SERVER-DFW	0	0	0	C	0	AMER
WLS-APP	Domain-SOUTH:WLS-SERVER-DFW:control	0	0	0	C	0	AMER
WLS-APP	Domain-SOUTH:WLS-SERVER-DFW:scanner	0	0	0	C	0	AMER

RTViewCentral Displays

This section describes the Views and displays that come with RTView Enterprise Monitor and reside on RTViewCentral. This section describes the following Views:

- [“All Management Areas”](#)
- [“Multi Area Service Views”](#)
- [“Single Area Service Views”](#)
- [“Service Summary Views”](#)
- [“Key Metrics Views”](#)
- [“Component Views”](#)
- [“Metric Explorer”](#)
- [“Alert Views”](#)
- [“Administration”](#)
- [“CMDB Administration”](#)
- [“Architecture”](#)
- [“Property Views”](#)
- [“Diagram Views”](#)

For details about add-on displays, see:

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

All Management Areas

These displays present the highest-level summary views of alert states for your entire system. Aggregated data is organized by Owners and shows all Areas, while highlighting the most critical alert states using color. Data can be filtered by Owner, Area, Environment and alert Metric. Data is filtered by the \$rtvOwnerMask and \$rtvAreaMask values for the logged in user. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.

Use these displays to monitor critical alerts anywhere in your system, and investigate those alerts in lower-level displays. Because these displays immediately show you any critical alert in your system, users typically keep one of these displays open for quick monitoring. Click an Area in the display to drill-down and view the selected Area in the **Multi Area Service Views** displays.

The **All Management Areas** displays present the same aggregated data in tabular and heatmap formats. Displays in this View are:

- [“Area Heatmap”](#): Heatmap of the most critical alerts for all Areas of your system, with the option to filter by Owner, Environment and alert Metric.
- [“Area Table”](#): Table of data shown in the **All Management Areas - “Area Heatmap”** with the option to filter by Owner and Environment.

Area Heatmap

View the most critical alert state for all monitored instances throughout your system. Consider keeping this display open to monitor conditions in your system. The heatmap organizes monitored instances by one or all Owners for all Areas, and uses color to show the most critical alert state in each. Each rectangle in the heatmap represents a management Area (for example, Applications, Demo Systems and so forth), which are also grouped by Owner. The rectangle size represents the number of CIs in the rectangle; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. By default, this display shows all Owners, all Environments and the Alert Impact.

Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Area in the display that was last selected under **Multi Area Service Views**. For example, if the last selected display under **Multi Area Service Views** was “[Group / Service Table](#)”, then clicking an Area in the heatmap results in displaying details in the **Group/Service 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.
- 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.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.







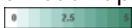
Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.


Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

Alert Impact	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from 0 - 10 , as indicated in the color gradient  bar, where 10 is the highest Alert Impact.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	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.
Criticality	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from 1 to 5, as indicated in the color gradient  bar, where 5 is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - "CI / Service Table" display, which range from A to E, where A is the highest Criticality (level 5 maps to a Criticality of A and level 1 maps to a Criticality of E with equally spaced intermediate values).</p>

Area Table

View data from the **All Management Areas** - "[Area Heatmap](#)" in a tabular format: all alert states (alert Impact, Severity, Count, Criticality and CI Count) for all Areas, Owners and Environments. Each row in the table is a different Area (for example, **Applications**, **Demo Systems** and so forth). Use this display to check the status of your systems by Area, Owner and Environment, and to compare detailed metrics across all Areas in your organization.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Drill-down and investigate by clicking a row in the table to view details for the selected Area in the display that was last selected under **Multi Area Service Views**. For example, if the last selected display under **Multi Area Service Views** was “[Group / Service Table](#)”, then clicking an Area in the heatmap results in displaying details in the **Group/Service Table** display.

←

Table

All Areas by Owner

24-Sep-2015 11:38

Data OK






+

?

Owner: All Owners

Env: QA

Area Count: 5

Owner	Area	Severity	Alert Count	Max Alert Impact	Criticality	CI Count
Infrastructure	Middleware		17	2	E	12
Infrastructure	Processes		59	2	E	36
Infrastructure	Servers		18	2	E	16
Jerelyn Parker	Backends		27	10		29
Jerelyn Parker	Systems		111	10		112

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.

Row Color Code:

Tables with colored rows indicate the following:

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

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.




Fields and Data

This display includes:

Area Count The current number of Areas shown in the table.

(Table)

Each row in the table is a different Area.

Owner	The name of the person or Group the Area is designated to.
Area	The name of the Area where the alert data originated.
Severity	The maximum level of alerts in the Area. Values range from 0 to 2 , where 2 is the greatest Severity:  One or more alerts exceeded their ALARM LEVEL threshold in the Area.  One or more alerts exceeded their WARNING LEVEL threshold in the Area.  No alert thresholds have been exceeded in the Area.
Criticality	The Criticality (rank of importance) specified in the Service Data Model (CMDDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E, where A is the highest Criticality. This value is used to determine the value for Alert Impact.
Max Alert Impact	The highest value that Alert Impact has had for the Area.
Alert Count	The total number of critical and warning alerts for the Area.
CI Count	The total number of configurable items associated with the Area.

Multi Area Service Views

These displays present aggregated data of alert states for all Services for all Areas. Data can be filtered by Area, Group, Environment, and alert Metric. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask and \$rtvGroupMask values for the logged in user. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.

Use these displays, for example, to isolate the Area and Environment in which a critical alert is occurring. If you see a critical alert, get information by comparing alert metrics (such as how many other items are potentially affected).

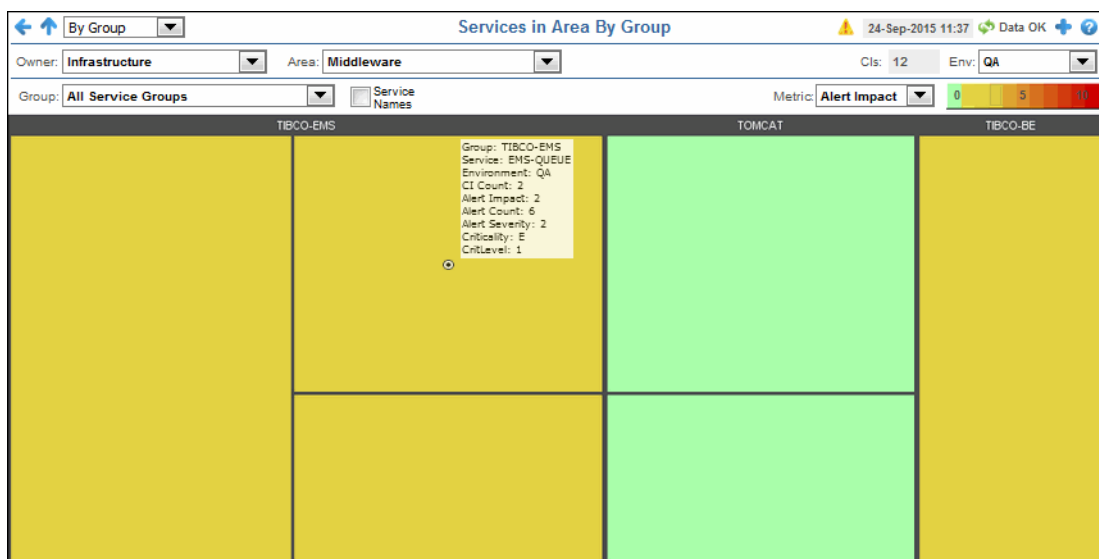
These displays drill-down to the **Service Summary Views** - “[Service By CI Type](#)” display. The **Multi Area Service Views** displays present data in tabular and heatmap formats. Displays in this View are:


- “[Group/Service Heatmap](#)”: Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
- “[Group/Region Heatmap](#)”: Heatmap as described for the **Group / Service Heatmap** (above), with the option to filter by Region and no option to show Service Names.
- “[Group / Service Table](#)”: Table of **Group/Service Heatmap** data.
- “[Services CI Type Summary](#)”: Table that shows the health state of Services per CI Type.
- “[Services History Heatmap](#)”: Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

Group/Service Heatmap

View heatmap of alert states for Services in one or all Areas, filter by Group or Environment, and optionally show Service Names. The heatmap organizes Services by one or all Areas. Each rectangle in the heatmap represents a Service (for example, Applications, Demo Systems and so forth), which are grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the “[Service Summary](#)” display under “[Service Summary Views](#)” and you clicked on a rectangle in the **Group / Service Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the “[Service KM Table](#)” display under “[Key Metrics Views](#)”, then clicking a rectangle in the **Group / Service Heatmap** displays the details in the **Service KM Table**.





Note: The “Up” Arrow (↑) opens the most recently viewed display under “All Management Areas”. For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking ↑ opens the “Area Table” display.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.




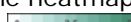
Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Metric:


Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

Alert Impact	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from 0 - 10 , as indicated in the color gradient  bar, where 10 is the highest Alert Impact.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none"> ● Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2. ● Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1. ● Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	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.
Criticality	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from 1 to 5, as indicated in the color gradient  bar, where 5 is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - “C1 / Service Table” display, which range from A to E, where A is the highest Criticality (level 5 maps to a Criticality of A and level 1 maps to a Criticality of E with equally spaced intermediate values).</p>

Group/Region Heatmap

View heatmap of alert states for one or all Services, Areas, Environment or Regions, and optionally show Service Names. The heatmap organizes CIs by one or all Groups. Each rectangle in the heatmap represents a Group, which are grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a rectangle in the **Group / Region Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a rectangle in the **Group / Region Heatmap** displays the details in the **Service KM Table**.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Note: The “Up” Arrow () opens the most recently viewed display under “All Management Areas”. For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking opens the “Area Table” display.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.







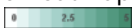
Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

Alert Impact	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from 0 - 10 , as indicated in the color gradient  bar, where 10 is the highest Alert Impact.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	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.
Criticality	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from 1 to 5, as indicated in the color gradient  bar, where 5 is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - "CI / Service Table" display, which range from A to E, where A is the highest Criticality (level 5 maps to a Criticality of A and level 1 maps to a Criticality of E with equally spaced intermediate values).</p>

Group / Service Table

This table displays data shown in the **Group/Service** and **Group/Region** heatmaps. View Service metrics (Impact, Severity, Count and Criticality, and CI Count) for one or all Areas, Owners, Groups and Environments, and compare detailed metrics across all Areas in your organization. The table lists Services by Owner and Area. Each row in the table is a different Service. The color of the circle in the **Alert Severity** column represents the most critical alert state for that Service.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.




Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Fields and Data

This display includes:

Service/ Region Count	The total number of Services listed in the table. This value is determined by the selections made from display drop-down menus.
Area	The name of the Area where the alert data originated.
Service	The name of the Service where the alert data originated.
Region	The name of the Region to which the Service applies.
Severity	The maximum level of alerts in the row. Values range from 0 to 2 , where 2 is the greatest Severity:  One or more alerts exceeded their ALARM LEVEL threshold in the Service.  One or more alerts exceeded their WARNING LEVEL threshold in the Service.  No alert thresholds have been exceeded in the Service.
Alert Count	The total number of critical and warning alerts for the Service.
Alert Impact	The maximum of the products of maximum Alert Severity multiplied by the Criticality of all CIs for the Service. Values range from 0 - 10 , where 10 is the highest Alert Impact.
Service Criticality	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E, where A is the highest Criticality.
CIs	The total number of configurable items in the display.
Environment	The name of the Environment to which the Service applies.
Group	The name of the Environment to which the Service applies.
CI Count	The total number of configurable items associated with the Service.

Services CI Type Summary

This display lists the health state of Services by CI Type and allows you to manage alerts. In the upper table, each column is a CI Type and each row is a Service. Select a row in the table to view details in the lower table.

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

Service Health By CI Type

Owner: Infrastructure Area: All Areas Cls: 80 Env: QA

Group: All Service Groups ☐ Valid CI Types Only Service/Region Count: 12

Service	All	JVM	AMX	AMX	AMX	BW	BW	BW	EMS	EMS	EMS	Active	Tomcat	Tomcat	Oracle
ORACLE		0	0	0	0	0	0	0	0	0	0	0	0	0	
HOST		0	0	0	0	0	0	0	0	0	0	0	0	0	0
VMWARE-HOST		0	0	0	0	0	0	0	0	0	0	0	0	0	0
VMWARE-VM		0	0	0	0	0	0	0	0	0	0	0	0	0	0
JVM			0	0	0	0	0	0	0	0	0	0	0	0	0
SOLACE-BRIDGE		0	0	0	0	0	0	0	0	0	0	0	0	0	0
TBE-CLUSTER		0	0	0	0	0	0	0	0	0	0	0	0	0	0
EMS-QUEUE		0	0	0	0	0	0	0	0	0		0	0	0	0
EMS-SERVER		0	0	0	0	0	0	0			0	0	0	0	0
EMS-TOPIC		0	0	0	0	0	0	0			0	0	0	0	0

Service: ORACLE CI Type: *

First Occ	Last Occ	Count	Sup	Owner	Alert Name	Primary Service	CI	Alert Text
10/06/15 11:05:16	10/07/15 06:18:...	580	<input type="checkbox"/>		OrainstanceNumActive...	Oracle	testBedOracle11g	High Alert Limit exceeded, current value: 25.0 limit: 15.0
10/06/15 11:05:16	10/07/15 06:19:...	2	<input type="checkbox"/>		OraDatabaseTablesac...	Oracle	testBedOracle11g	High Alert Limit exceeded, current value: 93.38 limit: 90.0
10/06/15 11:05:16	10/07/15 06:19:...	9	<input type="checkbox"/>		OraDatabaseTablesac...	Oracle	testBedOracle11g	High Alert Limit exceeded, current value: 93.74 limit: 90.0
10/06/15 11:05:16	10/07/15 06:18:...	580	<input type="checkbox"/>		OrainstanceNumCurren...	Oracle	testBedOracle11g	High Alert Limit exceeded, current value: 18.0 limit: 15.0
09/21/15 10:28:05	10/07/15 06:18:...	5573	<input type="checkbox"/>		OrainstanceMaxQueryT...	Oracle	testBedOracle11g	High Alert Limit exceeded, current value: 19754.0275 limit: 15000.0

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.

Note: The "Up" Arrow () opens the most recently viewed display under "All Management Areas". For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking opens the "Area Table" display.

For each Service in a selected Group, the round indicator shows the current maximum Alert Severity of all the CIs associated with each CI Type.

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

The cell background color indicates the current maximum Alert Impact of all the CIs associated with the Service and CI Type. The Alert Impact is calculated for each CI, which is the product of the CI Criticality times the current maximum Alert Severity. Background colors range from green to red, green being the lowest possible alert impact and red the highest possible value.

For example, in the following figure the first five Services in the list have an alert condition due to a BW Engine problem, and additionally the **INVENTORY MANAGER** Service has a TIBCO EMS Server problem. The **All CI Types** column shows the global highest level for all CI Types.

Service Name	All CI Types	User Experience	JVM	BW Server	BW Engine	TibcoEMS Server	TibcoEMS Topic	Tomcat
ACCOUNTING		0	0					
COMPLIANCE		0						
INVENTORY MANAGER		0	0				0	
ORDER PROCESSING		0	0				0	
REPORTING		0	0				0	
TUCON-EXCHANGE		0	0	0	0			0

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.


Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.


Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.






Fields and Data

This display includes:

Valid CI Types Only	Check to only show CI Type columns that contain data in the table, uncheck to include columns that are empty. Including empty table columns can be helpful when you are comparing Services (using the Group drop-down menu) because the table columns retain their order.
Service/Region Count	The total number of Services currently listed in the table.
Service Name	The name of the Service.
All CI Types	The circular indicator  shows the current maximum Alert Severity of all the CIs associated with the CI Type, and the cell background color shows the current maximum Alert Impact of all the CIs--across all CI Types-- associated with the Service.
Service	Shows the Service selected in the upper table.
CI Type	Shows the CI Type selected in the upper table.

Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details**, **Own**, **Suppress**, **Close**, **Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** buttons are disabled.
-  Opens the **Alerts Table** display in a new window.

Own	Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
Suppress	Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
Close	Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
Details	Select an alert, right-click and choose Alert/Details to open the Alert Detail window and view alert details. Or, double-click an alert to open the Alert Detail window.


Annotate	Select one or more alerts, right-click and choose Alert/Annotate to open the Set Owner and Comments dialog and enter comments or change alert owner.
Options	Select an alert, right-click and choose Alert/Options to open the Alert Options dialog. This dialog is provided for customizing your own alert options.
First Occ	The date and time the alert first occurred.
Last Occ	The date and time the alert last occurred.
Count	The number of times the alert was generated.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Primary Service	The name of the Service with which the alert is associated.
CI	The CI alert source.
Alert Text	Description of the alert.
AlertClass	An optional alert field which can be used when integrating with other alerting systems.
CompID	An optional alert field which can be used when integrating with other alerting systems.
TicketID	An optional alert field which can be used when integrating with other alerting systems.
TicketGroup	An optional alert field which can be used when integrating with other alerting systems.

Services History Heatmap

View history heatmap of alert states, over time, for Services in one Area, filtered by Group and Environment.

The history heatmap displays Services from one or more Groups and Environments of a given Owner and Area. Each row in the heatmap represents a different Service. The row color shows the Alert Impact or Alert Severity of a Service across time.

Use the available drop-down menus or right-click to filter data shown in the display. Mouse-over each row to see the time of alert state changes for particular Service occurred. For example, you can see at what time an alert state changed from green to red. Use the checkboxes ☒ to include or exclude labels in the heatmap. Drill-down and investigate by clicking a row in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metric Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a row in the **Services History Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a row in the **Services History Heatmap** displays the details in the **Service KM Table**.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

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.

Note: The “Up” Arrow (↑) opens the most recently viewed display under “All Management Areas”. For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking ↑ opens the “Area Table” display.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

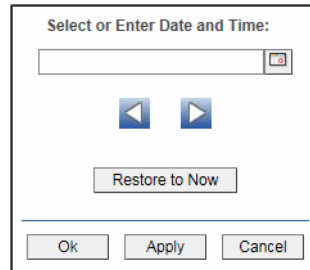
Color Code:


Row color indicates the following:



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

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 for the graph, click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the Time Range drop-down menu. Click Restore to Now to reset the time range end point to the current time.

Single Area Service Views

These displays present aggregated data of alert states for all Services for a specific Area. Data can be filtered by Area, Group, Environment, and alert Metric. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask and \$rtvGroupMask values for the logged in user. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.

Use these displays, for example, to isolate the Area and Environment in which a critical alert is occurring. If you see a critical alert, get information by comparing alert metrics (such as how many other items are potentially affected).

These displays drill-down to the **Service Summary Views** - "[Service By CI Type](#)". The **Single Area Service Views** displays present data in tabular and heatmap formats. Displays in this View are:

- "[Single Area: Group/Service Heatmap](#)": Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
- "[Single Area: Region/Service Heatmap](#)": Heatmap as described for the **Group / Service Heatmap** (above), with the option to filter by Region and no option to show Service Names.
- "[Single Area: Group / Service Table](#)": Table view of **Group/Service Heatmap** data.
- "[Single Area: Services CI Type Summary](#)": Table that shows the health state of Services per CI Type.
- "[Single Area: Services History Heatmap](#)": Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

Single Area: Group/Service Heatmap

View heatmap of alert states for Services in one Area, filter by Group or Environment, and optionally show Service Names. Each rectangle in the heatmap represents a Service (for example, Applications, Demo Systems and so forth), which are grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a rectangle in the **Group / Service Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a rectangle in the **Group / Service Heatmap** displays the details in the **Service KM Table**.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Note: The “Up” Arrow () opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking opens the “Group/Region Heatmap” display.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.





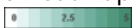
Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Metric:

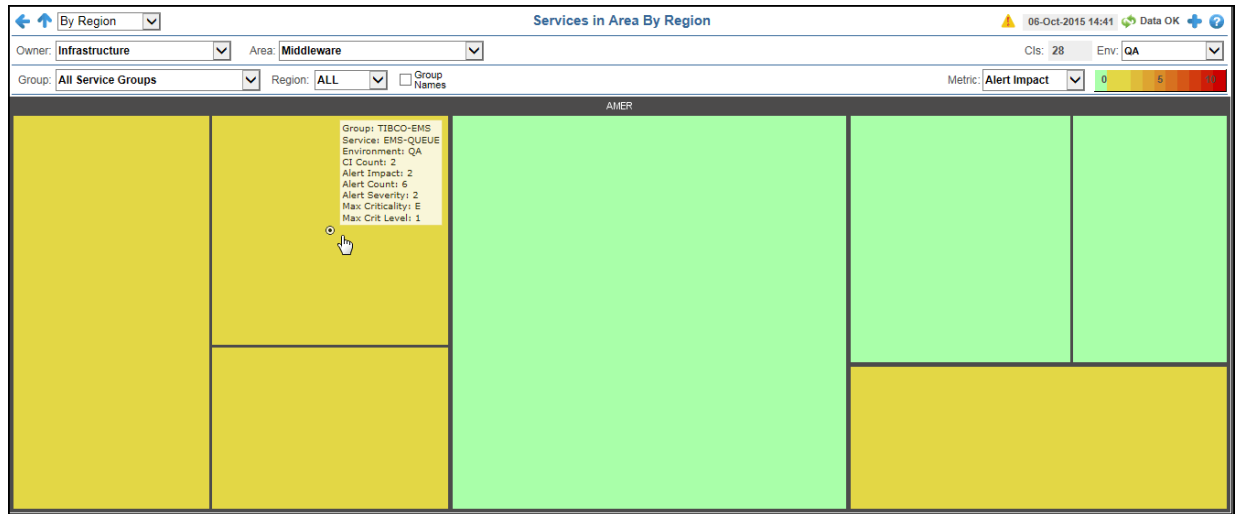
Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

Alert Impact	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from 0 - 10 , as indicated in the color gradient  bar, where 10 is the highest Alert Impact.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	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.
Criticality	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from 1 to 5, as indicated in the color gradient  bar, where 5 is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - "CI / Service Table" display, which range from A to E, where A is the highest Criticality (level 5 maps to a Criticality of A and level 1 maps to a Criticality of E with equally spaced intermediate values).</p>

Single Area: Region/Service Heatmap

View heatmap of alert states for one Owner, one specific Area, one or all Service Groups, and one or all Regions. You can also optionally show Service Group Names. The heatmap organizes CIs by one or all Groups. Each rectangle in the heatmap represents a Group, which is grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a rectangle in the **Group / Region Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a rectangle in the **Group / Region Heatmap** displays the details in the **Service KM Table**.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

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.

Note: The “Up” Arrow (↑) opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under Multi Area Service Views was Group/Region Heatmap, then clicking ↑ opens the “Group/Region Heatmap” display.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.








Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.


Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

Alert Impact	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from 0 - 10 , as indicated in the color gradient  bar, where 10 is the highest Alert Impact.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	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.
Criticality	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from 1 to 5, as indicated in the color gradient  bar, where 5 is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - "CI / Service Table" display, which range from A to E, where A is the highest Criticality (level 5 maps to a Criticality of A and level 1 maps to a Criticality of E with equally spaced intermediate values).</p>

Single Area: Group / Service Table

This table displays data shown in the **Group/Service** and **Region/Service** heatmaps. View Service metrics (Impact, Severity, Count and Criticality, and CI Count) for a specific Area, for one or all Owners, Groups, and Environments. The table lists Services by Owner and Area. Each row in the table is a different Service. The color of the circle in the **Alert Severity** column represents the most critical alert state for that Service.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Drill-down and investigate by clicking a row in the table to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metric Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a row in the table, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a row in the table displays the details in the **Service KM Table**.

Table

Services in Area By Group

Owner

Infrastructure

Area

Middleware

Cls

28

Env

QA

Group







All Service Groups


Service/Region Count


7



Service	Region	Alert Severity	Alert Count	Alert Impact	Service Criticality	CI Count	Environment	Group	
SOLACE-BRIDGE	AMER		0	0	E	16	QA	SOLACE	
TBE-CLUSTER	AMER		1	1	E	2	QA	TIBCO-BE	
EMS-QUEUE	AMER		6	2	E	2	QA	TIBCO-EMS	
EMS-SERVER	AMER		2	2	E	1	QA	TIBCO-EMS	
EMS-TOPIC	AMER		8	2	E	4	QA	TIBCO-EMS	
TOMCAT	AMER		0	0	E	1	QA	TOMCAT	
TOMCAT-APP	AMER		0	0	E	2	QA	TOMCAT	

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.
-   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.

Note: The “Up” Arrow () opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking  opens the “Group/Region Heatmap” display.

Row Color Code:

Tables with colored rows indicate the following:

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

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.




Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.


Fields and Data

This display includes:

Service/ Region Count	The total number of Services listed in the table. This value is determined by the selections made from display drop-down menus.
Area	The name of the Area where the alert data originated.
Service	The name of the Service where the alert data originated.
Region	The name of the Region to which the Service applies.
Severity	The maximum level of alerts in the row. Values range from 0 to 2 , where 2 is the greatest Severity:  One or more alerts exceeded their ALARM LEVEL threshold in the Service.  One or more alerts exceeded their WARNING LEVEL threshold in the Service.  No alert thresholds have been exceeded in the Service.
Alert Count	The total number of critical and warning alerts for the Service.
Alert Impact	The maximum of the products of maximum Alert Severity multiplied by the Criticality of all CIs for the Service. Values range from 0 - 10 , where 10 is the highest Alert Impact.
Service Criticality	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E, where A is the highest Criticality.
CIs	The total number of configurable items in the display.
Environment	The name of the Environment to which the Service applies.
Group	The name of the Environment to which the Service applies.
CI Count	The total number of configurable items associated with the Area.

Single Area: Services CI Type Summary

This display lists the health state of Services by CI Type and allows you to manage alerts. In the upper table, each column is a CI Type and each row is a Service. Select a row in the table to view details in the lower table.

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

Service Health By CI Type

Owner: Infrastructure Area: Middleware CIs: 28 Env: QA

Group: All Service Groups ☐ Valid CI Types Only Service/Region Count: 7

Service Name	All CI Types	AMX Node	AMX Service	AMX ServiceNod	BW Server	BW Engine	BW Proc	EMS Server	EMS Topic	EMS Queue	Active Spaces	Tomcat	Tomcat App	ServiceGroup
SOLACE-BRIDGE		0	0	0	0	0	0	0	0	0	0	0	0	SOLACE
TBE-CLUSTER		0	0	0	0	0	0	0	0	0	0	0	0	TIBCO-BE
EMS-QUEUE		0	0	0	0	0	0	0	0		0	0	0	TIBCO-EMS
EMS-SERVER		0	0	0	0	0	0		0	0	0	0	0	TIBCO-EMS
EMS-TOPIC		0	0	0	0	0	0	0		0	0	0	0	TIBCO-EMS
TOMCAT		0	0	0	0	0	0	0	0	0	0		0	TOMCAT
TOMCAT-APP		0	0	0	0	0	0	0	0	0	0		0	TOMCAT

Service: EMS-QUEUE CI Type: *


First Occ	Last Occ	Count	Sup	Owner	Alert Name	Primary Service	CI	Alert Text
10/06/15 06:18:02	10/06/15 06:18:02	1	<input type="checkbox"/>		EmsQueueProviderId...	EMS	tcp://192.168.200.13	High Alert Limit exceeded, current value: 97.0 limit: 80.0
10/06/15 06:18:02	10/06/15 06:18:02	1	<input type="checkbox"/>		EmsQueueProviderId...	EMS	tcp://192.168.200.13	High Alert Limit exceeded, current value: 97.0 limit: 80.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesProducerC...	EMS	tcp://192.168.200.13	Low Alert Limit exceeded, current value: 0.0 limit: 5.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesProducerC...	EMS	tcp://192.168.200.13	Low Alert Limit exceeded, current value: 0.0 limit: 5.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesConsumer...	EMS	tcp://192.168.200.13	Low Alert Limit exceeded, current value: 0.0 limit: 5.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesConsumer...	EMS	tcp://192.168.200.13	Low Alert Limit exceeded, current value: 0.0 limit: 5.0


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.

Note: The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking  opens the "Group/Region Heatmap" display.

For each Service in a selected Group, the round indicator  shows the current maximum Alert Severity of all the CIs associated with each CI Type.

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

The cell background color indicates the current maximum Alert Impact of all the CIs associated with the Service and CI Type. The Alert Impact is calculated for each CI, which is the product of the CI Criticality times the current maximum Alert Severity. Background colors range from green to red, green being the lowest possible alert impact and red the highest possible value.

For example, in the following figure the first five Services in the list have an alert condition due to a BW Engine problem, and additionally the **INVENTORY MANAGER** Service has a TIBCO EMS Server problem. The **All CI Types** column shows the global highest level for all CI Types.

Service Name	All CI Types	User Experience	JVM	BW Server	BW Engine	TibcoEMS Server	TibcoEMS Topic	Tomcat
ACCOUNTING		0	0					
COMPLIANCE		0						
INVENTORY MANAGER		0	0					
ORDER PROCESSING		0	0					
REPORTING			0					
TUCON-EXCHANGE		0	0	0	0			0

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.


Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.


Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.






Fields and Data

This display includes:

- Valid CI Types Only** Check to only show CI Type columns that contain data in the table, uncheck to include columns that are empty. Including empty table columns can be helpful when you are comparing Services (using the **Group** drop-down menu) because the table columns retain their order.
- Service / Region Count** The total number of Services currently listed in the table.
- Service Name** The name of the Service.
- All CI Types** The circular indicator  shows the current maximum Alert Severity of all the CIs associated with the CI Type, and the cell background color shows the current maximum Alert Impact of all the CIs--across all CI Types-- associated with the Service.
- Service** Shows the Service selected in the upper table.
- CI Type** Shows the CI Type selected in the upper table.

Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details**, **Own**, **Suppress**, **Close**, **Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** buttons are disabled.
-  Opens the **Alerts Table** display in a new window.

- Own** Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Suppress** Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Close** Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.

Details	Select an alert, right-click and choose Alert/Details to open the Alert Detail window and view alert details. Or, double-click an alert to open the Alert Detail window.
Annotate	Select one or more alerts, right-click and choose Alert/Annotate to open the Set Owner and Comments dialog and enter comments or change alert owner.
Options	Select an alert, right-click and choose Alert/Options to open the Alert Options dialog. This dialog is provided for customizing your own alert options.
First Occ	The date and time the alert first occurred.
Last Occ	The date and time the alert last occurred.
Count	The number of times the alert was generated.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Primary Service	The name of the Service with which the alert is associated.
CI	The CI alert source.
Alert Text	Description of the alert.
AlertClass	An optional alert field which can be used when integrating with other alerting systems.
CompID	An optional alert field which can be used when integrating with other alerting systems.
TicketID	An optional alert field which can be used when integrating with other alerting systems.
TicketGroup	An optional alert field which can be used when integrating with other alerting systems.


Single Area: Services History Heatmap

View history heatmap of alert states, over time, for Services in one Area, filtered by Group and Environment.

The history heatmap displays Services from one or more Groups and Environments of a given Owner and Area. Each row in the heatmap represents a different Service. The row color shows the Alert Impact or Alert Severity of a Service across time.

Use the available drop-down menus or right-click to filter data shown in the display. Mouse-over each row to see the time of alert state changes for particular Service occurred. For example, you can see at what time an alert state changed from green to red. Use the checkboxes ☒ to include or exclude labels in the heatmap. Drill-down and investigate by clicking a row in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metric Views**.

For example, if the last selected display was the “Service KM Table” display under “Key Metrics Views” and you clicked on a row in the table, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a row in the table displays the details in the **Service KM Table**.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

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.

Note: The “Up” Arrow (↑) opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking ↑ opens the “Group/Region Heatmap” display.

Filter By:

The following filtering options are typically included:

- Owner:** Choose an Owner to see metrics for Areas associated with that Owner.
- Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.
- Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.
- Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.
- Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

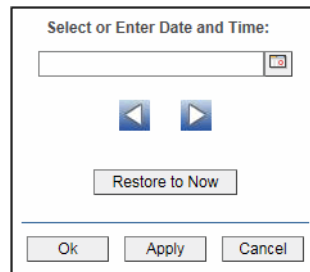
Color Code:


Row color indicates the following:



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

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 for the graph, click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Service Summary Views

These displays present alert states at the component-level by Service in tabular and heatmap formats, while highlighting the most critical alert state. Data can be filtered by Owner, Area, Group, Service or Environment. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.


Use these displays to get alert details and detailed status information for a particular Service, such as a list of all the CI Types relevant to a Service and the quality of the performance metrics for each CI. Displays in this View are:

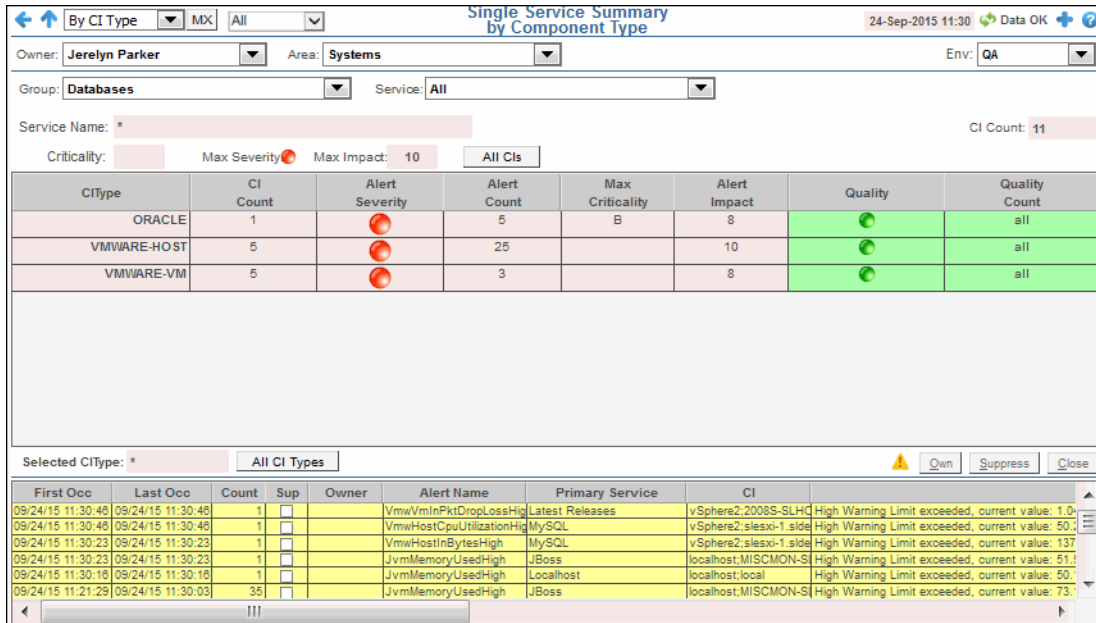
- **"Service By CI Type"**: Table of alert states for a Service organized CI Type, with general alert information.
- **"Service Summary"**: Table of CIs by Service, with detailed alert information.
- **"Service Health Heatmap"**: Heatmap of CIs by Service, with the option to filter by Owner, Area, Group, Environment and alert Metric, and show CI Names.

Service By CI Type

View alert states for a Service organized CI Type and manage alerts. See the CI Count for a Service and obtain alert statistics for CI Types such as Alert Severity and Alert Count. Use this display to summarize alerts occurring for a Service and determine which component types are malfunctioning. View a list of all active alerts associated with the CI Type.

The upper table lists all CI Types for the selected Service with alert details such as the highest Alert Severity. Each row is a CI Type. The color of each row represents the maximum Alert Impact for the row. Select a row that has an active alert (the Alert Severity is red or yellow) to view the active alerts in the lower table. Double-click a row to view a detailed list of CIs associated with the CI Type in the **Service Summary** display. In the lower table, each row is a different alert for a CI that is associated with the CI Type selected from the upper table.

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



CType	CI Count	Alert Severity	Alert Count	Max Criticality	Alert Impact	Quality	Quality Count
ORACLE	1		5	B	8		all
VMWARE-HOST	5		25		10		all
VMWARE-VM	5		3		8		all

First Occ	Last Occ	Count	Sup	Owner	Alert Name	Primary Service	CI	
09/24/15 11:30:46	09/24/15 11:30:46	1	<input type="checkbox"/>		VmwVmInPktDropLossHig	Latest Releases	vSphere2:2008S-SLHC	High Warning Limit exceeded, current value: 1.0
09/24/15 11:30:46	09/24/15 11:30:46	1	<input type="checkbox"/>		VmwHostCpuUtilizationHig	MySQL	vSphere2:slesxi-1.sle	High Warning Limit exceeded, current value: 50.0
09/24/15 11:30:23	09/24/15 11:30:23	1	<input type="checkbox"/>		VmwHostInBytesHigh	MySQL	vSphere2:slesxi-1.sle	High Warning Limit exceeded, current value: 137
09/24/15 11:30:23	09/24/15 11:30:23	1	<input type="checkbox"/>		JvmMemoryUsedHigh	JBoss	localhost:MISCMON-S	High Warning Limit exceeded, current value: 51.1
09/24/15 11:30:16	09/24/15 11:30:16	1	<input type="checkbox"/>		JvmMemoryUsedHigh	Localhost	localhost:local	High Warning Limit exceeded, current value: 50.0
09/24/15 11:21:29	09/24/15 11:30:03	35	<input type="checkbox"/>		JvmMemoryUsedHigh	JBoss	localhost:MISCMON-S	High Warning Limit exceeded, current value: 73.0

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.
- 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.

Note: The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the "Services CI Type Summary" display.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.




Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Fields and Data

This display includes:








Service Name The name of the selected Service.

CI Count The total number of configurable items in the display.


Criticality	The Criticality (rank of importance) specified in the Service Data Model (CMDDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E , where A is the highest Criticality. This value is used to determine the value for Alert Impact.
Max Severity	<p>The highest Alert Severity value of any CI associated with the selected Service. Values range from 0 to 2, where 2 is the greatest Severity:</p> <ul style="list-style-type: none">  One or more alerts exceeded their ALARM LEVEL threshold in the Service.  One or more alerts exceeded their WARNING LEVEL threshold in the Service.  No alert thresholds have been exceeded in the Service.
Max Impact	The highest Alert Impact value of any CI associated with the selected Service.
All CIs	Opens the Service Summary display.





(CI Type Table)

This table lists all CI Types for the selected Service. Each row in the table is a CI Type. Click a row to view details in the lower table about alerts associated with the CI Type. Double-click a row to drill-down to Service Summary display describing alert details relevant to this CI Type.

CIType	The type of CI.
CI Count	The total number of configurable items associated with the CI Type.
Alert Severity	<p>The highest Alert Severity value of any CI associated with the selected Service. Values range from 0 to 2, where 2 is the greatest Severity:</p> <ul style="list-style-type: none">  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
Alert Count	The total number of active alerts for the CIs associated with the CI Type.
Quality	<p>Shows whether performance metrics are being received from the CIs associated with the CI Type.</p> <ul style="list-style-type: none">  One or more performance metrics are not being received from the CIs associated with the CI Type.  All performance metrics are being received from the CIs associated with the CI Type.
Quality Count	<p>Shows the number of CIs for that CI Type that have a known state. It displays all when that number is the total count of CI's.</p> <ul style="list-style-type: none">  One or more performance metrics are not being received from the CIs associated with the CI Type.  All performance metrics are being received from the CIs associated with the CI Type.
Selected CI Type	Shows the CI Type selected in the upper table.
All CI Types	Shows all active alerts for all CIs associated with the CI Type selected.

Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details**, **Own**, **Suppress**, **Close**, **Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** buttons are disabled.

 Opens the **Alerts Table** display in a new window.

- Own** Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Suppress** Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Close** Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Details** Select an alert, right-click and choose **Alert/Details** to open the **Alert Detail** window and view alert details. Or, double-click an alert to open the **Alert Detail** window.
- Annotate** Select one or more alerts, right-click and choose **Alert/Annotate** to open the **Set Owner and Comments** dialog and enter comments or change alert owner.
- Options** Select an alert, right-click and choose **Alert/Options** to open the **Alert Options** dialog. This dialog is provided for customizing your own alert options.


First Occ	The date and time the alert first occurred.
Last Occ	The date and time the alert last occurred.
Count	The number of times the alert was generated.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Primary Service	The name of the Service with which the alert is associated.
CI	The CI alert source.
Alert Text	Description of the alert.
AlertClass	An optional alert field which can be used when integrating with other alerting systems.
CompID	An optional alert field which can be used when integrating with other alerting systems.
TicketID	An optional alert field which can be used when integrating with other alerting systems.
TicketGroup	An optional alert field which can be used when integrating with other alerting systems.

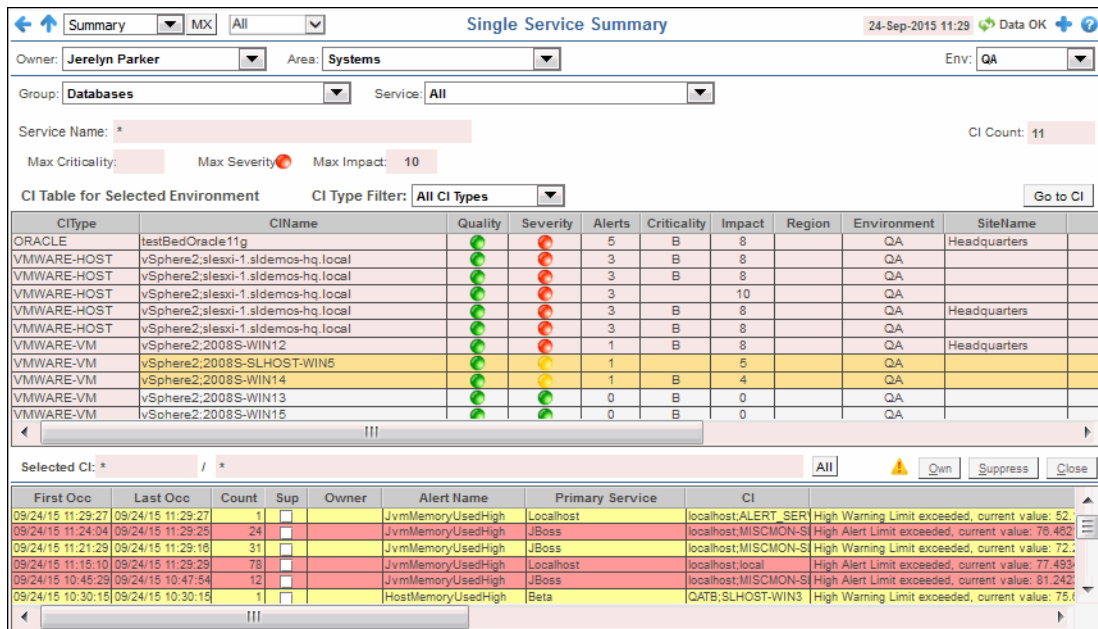
Service Summary

View alert states at the component-level per Service, manage alerts, obtain component details such as the number of active alerts for the component, which operating system the component uses and the Data Server associated with the component.

Use this display to monitor a Service in a specific Group or Environment anywhere in your organization, and determine whether a component is malfunctioning.

The table lists all components for a selected Service. Each row in the table is a different CI (configurable item or component). Each CI can have multiple alerts. Click a row to view details in the lower table about any alerts associated with the CI.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Double-click a row to drill-down to a summary page describing information relevant to this CI. This action can also be performed by selecting (a single click) on a row and selecting the **Go to CI** button.



Single Service Summary 24-Sep-2015 11:29 Data OK

Owner: Jerelyn Parker Area: Systems Env: QA

Group: Databases Service: All

Service Name: * CI Count: 11

Max Criticality: Max Severity: Max Impact: 10

CI Table for Selected Environment CI Type Filter: All CI Types Go to CI

CIType	CIName	Quality	Severity	Alerts	Criticality	Impact	Region	Environment	SiteName
ORACLE	testBedOracle11g			5	B	8		QA	Headquarters
VMWARE-HOST	vSphere2.slesxi-1.sldemos-hq.local			3	B	8		QA	
VMWARE-HOST	vSphere2.slesxi-1.sldemos-hq.local			3	B	8		QA	
VMWARE-HOST	vSphere2.slesxi-1.sldemos-hq.local			3	B	10		QA	
VMWARE-HOST	vSphere2.slesxi-1.sldemos-hq.local			3	B	8		QA	Headquarters
VMWARE-HOST	vSphere2.slesxi-1.sldemos-hq.local			3	B	8		QA	
VMWARE-VM	vSphere2.2008S-WIN12			1	B	8		QA	Headquarters
VMWARE-VM	vSphere2.2008S-SLHOST-WIN5			1	B	5		QA	
VMWARE-VM	vSphere2.2008S-WIN14			1	B	4		QA	
VMWARE-VM	vSphere2.2008S-WIN13			0	B	0		QA	
VMWARE-VM	vSphere2.2008S-WIN15			0	B	0		QA	

Selected CI: * / * All Own Suppress Close

First Occ	Last Occ	Count	Sup	Owner	Alert Name	Primary Service	CI	
09/24/15 11:29:27	09/24/15 11:29:27	1	<input type="checkbox"/>		JvmMemoryUsedHigh	Localhost	localhost:ALERT_SER	High Warning Limit exceeded, current value: 52.
09/24/15 11:24:04	09/24/15 11:29:25	24	<input type="checkbox"/>		JvmMemoryUsedHigh	JBoss	localhost:MISCMON-S	High Alert Limit exceeded, current value: 78.462
09/24/15 11:21:29	09/24/15 11:29:18	31	<input type="checkbox"/>		JvmMemoryUsedHigh	JBoss	localhost:MISCMON-S	High Warning Limit exceeded, current value: 72.
09/24/15 11:15:10	09/24/15 11:29:29	78	<input type="checkbox"/>		JvmMemoryUsedHigh	Localhost	localhost:local	High Alert Limit exceeded, current value: 77.459
09/24/15 10:45:29	09/24/15 10:47:54	12	<input type="checkbox"/>		JvmMemoryUsedHigh	JBoss	localhost:MISCMON-S	High Alert Limit exceeded, current value: 81.242
09/24/15 10:30:15	09/24/15 10:30:15	1	<input type="checkbox"/>		HostMemoryUsedHigh	Beta	QATB:SLHOST-WIN3	High Warning Limit exceeded, current value: 75.1

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 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.

Note: The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the "Services CI Type Summary" display.

Row Color Code:

Tables with colored rows indicate the following:

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

Filter By:

Use the available drop-down menus to filter data shown in the table. The display might include these filtering options:

Owner: Choose an Owner to see metrics in the heatmap for Areas associated with that Owner.

Area: Choose an Area to see metrics in the heatmap for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics in the heatmap for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics in the heatmap for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics in the heatmap for Environments associated with that Service, Group, Area and Owner.






Fields and Data

This display includes:

Service Name	The name of the selected Service.
Criticality	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E , where A is the highest Criticality. This value is used to determine the value for Alert Impact.
CI Count	The total number of configurable items in the display.


CI Table for Selected Environment






This table lists all CIs for the selected Environment. Each row in the table is a CI. Each CI can have multiple alerts. Click a row to view details about any alerts associated with the CI in the lower table. Double-click a row to drill-down to a summary page describing information relevant to this CI. This action can also be performed by selecting (a single click) on a row and selecting the **Go to CI** button.

CI Type Filter	Select a CI Type to display in the table or select All CI Types.
Go to CI	Drill-down to a summary page describing information relevant to this CI.
CIType	The type of CI.
Quality	Shows whether performance metrics are being received from the CI:  Performance metrics are not being received from the CI.  Performance metrics are being received from the CI.
Severity	Shows the most critical alert state for the selected CI:  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
Alerts	The number of currently active alerts for the selected CI.
Region	The name of the Region for the CI.
SiteName	The name of the Site for the CI.
OSType	The operating system currently running on the CI.
City	The name of the City for the CI.
Country	The name of the Country for the CI.
Data Server	The name of the Data Server with which the CI is associated.

- Selected CI** Shows the CI Type selected in the upper table.
- All** Shows all active alerts for all CIs associated with the CI Type selected.

Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details**, **Own**, **Suppress**, **Close**, **Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** buttons are disabled.
-  Opens the **Alerts Table** display in a new window.

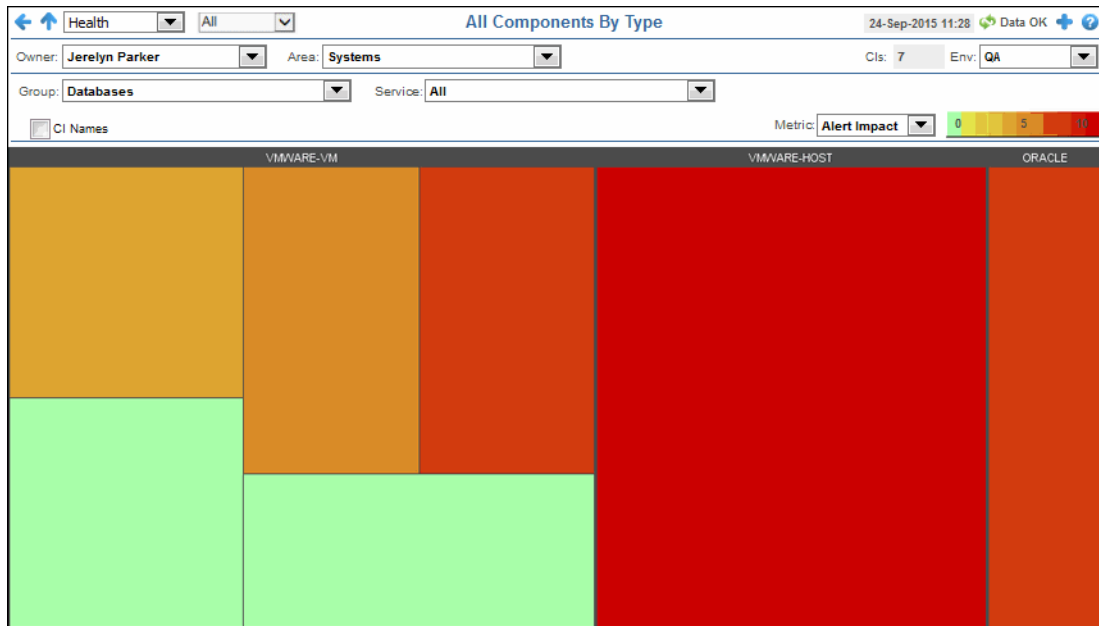
- Own** Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Suppress** Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Close** Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Details** Select an alert, right-click and choose **Alert/Details** to open the **Alert Detail** window and view alert details. Or, double-click an alert to open the **Alert Detail** window.
- Annotate** Select one or more alerts, right-click and choose **Alert/Annotate** to open the **Set Owner and Comments** dialog and enter comments or change alert owner.
- Options** Select an alert, right-click and choose **Alert/Options** to open the **Alert Options** dialog. This dialog is provided for customizing your own alert options.

First Occ	The date and time the alert first occurred.
Last Occ	The date and time the alert last occurred.
Count	The number of times the alert was generated.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Primary Service	The name of the Service with which the alert is associated.
CI	The CI alert source.
Alert Text	Description of the alert.
AlertClass	An optional alert field which can be used when integrating with other alerting systems.
CompID	An optional alert field which can be used when integrating with other alerting systems.
TicketID	An optional alert field which can be used when integrating with other alerting systems.
TicketGroup	An optional alert field which can be used when integrating with other alerting systems.

Service Health Heatmap

View heatmap of alert states for CIs associated with a Service. The heatmap organizes CIs by the Service selected. Each rectangle in the heatmap represents a CI (for example, **localhost-14**). Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. By default, this display shows Alert Impact.



Title Bar (possible features are):

- Open the previous display.
 - Open the 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.

Note: The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the "Services CI Type Summary" display.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.








Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

Alert Impact	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from 0 - 10 , as indicated in the color gradient  bar, where 10 is the highest Alert Impact.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	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.
Criticality	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from 1 to 5, as indicated in the color gradient  bar, where 5 is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - “CI / Service Table” display, which range from A to E, where A is the highest Criticality (level 5 maps to a Criticality of A and level 1 maps to a Criticality of E with equally spaced intermediate values).</p>

Key Metrics Views

The Key Metrics (KM) feature is an entirely new way of looking at and interpreting application health and performance data.

In contrast to the traditional Alert Impact view showing your ACTIVE alerts and their impact on the overall application or service, the Key Metrics view shows how close a metric is approaching its threshold over a period of time – both before and after the alert threshold is reached.

This allows you to both proactively anticipate performance problems BEFORE the alert threshold is crossed as well analyze the circumstances that led up to error conditions AFTER you got an alert. Armed with this knowledge, you can avert disasters before they happen and resolve problems faster after they happen.

RTView does this by correlating the most valuable key metrics over multiple components within a service and displaying them in context with both real-time and historical data. This is valuable because health problems in one component may be caused by performance problems in another and only by viewing each of these metrics in context with one another over a period of time are you able to visually link the relationship between troubled components.

It is important to note that your Alert Impact heatmaps may look very different from your Key Metrics heatmaps given that KM will indicate potential threats BEFORE they show up as alerts.

Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.

For Key Metrics definitions by technology, see [“Available KM Metrics and Alerts”](#).

Dependencies

The KM package is dependent on the Metric Explorer package. Both must be included in your project in order for KM to work. If you are upgrading from a version previous to 1.5.0 and have not added Metric Explorer to your project, see the *RTView Enterprise Monitor® User's Guide Upgrade Notes* section for information about including it.

Displays in this View are:

- [“Service KM Heatmap”](#): Heatmap of Key Metrics current data for one or more Services in your CMDB hierarchy.
- [“Service KM Table”](#): Table of Key Metrics current data for one or more Services.
- [“Service KM History”](#): History heatmap of Key Metrics historical data for one or more Services.
- [“Service KM History \(Alt\)”](#): History heatmap of Key Metrics historical data for one or more Services.


This section also includes:






- [“Available KM Metrics and Alerts”](#): List and descriptions of available key metrics.

Service KM Heatmap

View Key Metrics current data for one or more Services in your CMDB hierarchy in a heatmap. The **Service KM Heatmap** provides one view of all your Services and whether they are approaching an alert condition.

The most important overview of your Services is the Alert Impact View. The Alert Impact View lets you know what is a problem NOW. The **Service KM Heatmap** gives you a proactive view of which Services might be approaching a serious problem so that you can take action before they become critical. First look at the Alert Impact View to address current issues, then move to the **Service KM Heatmap** for proactive analysis.

The colors in the display are determined by the **Threshold %** and **Quality** values. As shown in the color gradient bar , a rectangle is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

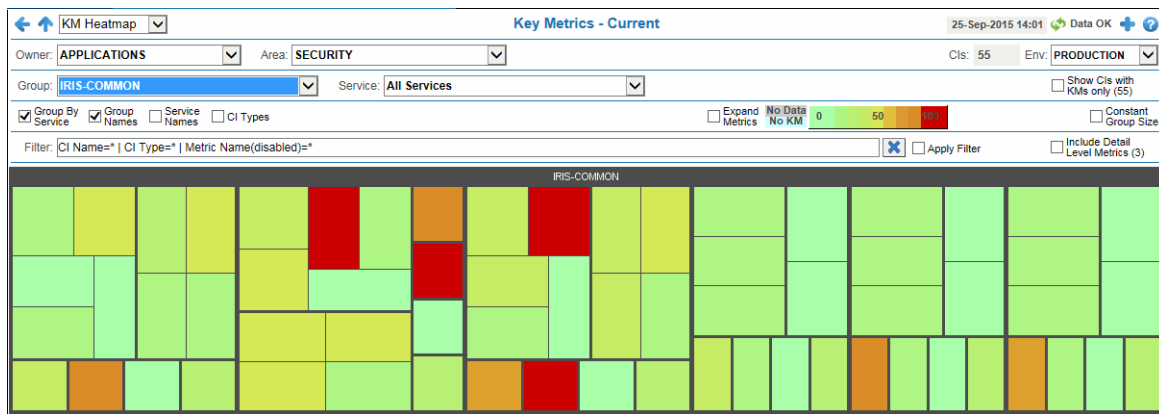
-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

Select **Group By Service** to include the **Group** and **Service** labels in the display. Select **Expand Metrics** to include the **Metric Name**, **Metric Value** and **Threshold** labels in the display.

For an overview about the Key Metrics feature, see [“Key Metrics Views”](#).

For Key Metrics definitions by technology, see [“Available KM Metrics and Alerts”](#).

Use the available drop-down menus or right-click to filter data shown in the display. Drill-down and investigate by double-clicking a rectangle in the display to view details in the corresponding 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.

Note: The “Up” Arrow (↑) opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking ↑ opens the “Services CI Type Summary” display.

Filter By:

The following filtering options are typically included:



Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.







Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.



You can also filter KMs by **CI Name**, **CI Type** and **Metric Name**. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item in the display and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click Clear  to clear the filter.


Note: The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.


Fields and Data

This display includes:

Show CIs with KMs only (x)	When selected, hides any CIs that do not have KMs defined. The number following the label (x) is the number of CIs with KMs defined.
Group By Service	When selected, includes the Service Group and Service Name in the KM data. CIs that are included in multiple Services will appear multiple times, once for each Service they are associated with.
Group Names	When selected, includes the Group Name in the display. Only available if Group By Service is selected.
Service Names	When selected, includes both the Group Name and Service Name in the display. Only available if Group By Service is selected.
CI Types	When selected, includes the CI Type in the display. If Group By Service is selected, this is shown in addition to the Group Name and Service Name .
Expand Metrics	When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum Threshold % and the minimum Quality value for the CI.
	The No Data No KM indicates the Quality value for the data. If no KMs are defined for the CI Type, the Quality is set to 0 and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the Quality is set to -1 and the color is shown as gray. If data was received for the metric, the Quality is set to 1 and the color is set based on the Threshold % value as described above. If the Expand Metrics checkbox is selected, this is the Quality of a single KM. If the Expand Metrics checkbox is not selected, this is the lowest Quality for all of the KMs on the CI.
	<p>The gradient bar is the legend for the display colors, which are determined by the Threshold % and Quality values. A row is green when the value is close to 0 changing to yellow, orange and red as the value gets closer to 100. Values at or over 100 are shown as red.</p> <ul style="list-style-type: none">  Red indicates the value is at or over 100.  Yellow indicates the value is between 0 and 100.  Green indicates the value is close to 0.  Teal indicates no KMs are defined for the CI Type.  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.
Constant Group Size	When selected, Groups are equally sized in the display. When not selected, Groups are sized according to the number of elements in the Group (a Group containing the most elements is rendered with the largest rectangle).

Filter: Shows the current filter parameters and is highlighted in blue when the filter is applied.
 By default, all data is shown:
CI Name=* | CI Type=* | Metric Name(disabled)=*
 To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click **Clear**  to clear the filter.

 Clears the filter parameters.

 Applies the filter parameters.

Include Detail Level Metrics (##) When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (**x**) is the number of detail level metrics available for the currently displayed KMs.

Mouseover

See the following details via mouseover:

Group The Group name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row per Group that a CI is associated with.

Service The Service name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row Service that a CI is associated with.

CI Type The CI Type.


CI Name The CI Name.

Metric Name The name of the metric. This is only included if the **Expand Metrics** checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the CI Type. To see which cache column provides data for this metric, navigate to **Architecture - "RTView KM Defs"**. In the table, look in the **CITYPE** and **SELECTOR** columns to find the row for your metric. The **CACHENAME** column lists the name of the cache containing the metric and the **METRICNAME** column contains the name of the cache column.

Metric Value The value of the metric. This is only included if the **Expand Metrics** checkbox is selected.

Threshold The **Alarm Level** value for the alert associated with the metric. This column is only included if the **Expand Metrics** checkbox is selected. To see which alert is associated with this metric, navigate to **Architecture - "RTView KM Defs"**. In the table, look in the **CITYPE** and **SELECTOR** columns to find the row for your metric. The **AlertName** column lists the name of the alert associated with the metric.
Note: When looking up the alert threshold for a KM, RTView Enterprise Monitor first looks to see if there is an alert override on the alert where the **AlertIndex** exactly matches the CIName (ignoring the **~** and **;** delimiters). If an exact match is found, the override **Alarm Level** is used. If no exact match is found, the **Default Alarm Level** for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.

Threshold % The percent of the **Metric Value** against the **Threshold**. If the **Expand Metrics** checkbox is selected, this is the **Threshold %** of a single KM. If the **Expand Metrics** checkbox is not selected, this is the highest **Threshold %** for all of the KMs on the CI.


Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to **0-10000**. For memory metrics, an exponential scale is applied to the **Threshold %** so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to **0** changing to yellow to orange to red as the value gets closer to **100**. Values at or over **100** are shown as red . To see which **CalcMode** is used for this metric, navigate to **Architecture - "RTView KM Defs"**. In the table, look in the **CITYPE** and **SELECTOR** columns to find the row for your metric. The **CalcMode** column lists the type of scale that is applied to the metric. If blank, no scale is applied.






Quality	<p>Indicates the quality of the data. If the Expand Metrics checkbox is selected, the value is for a single KM on the CI. If the Expand Metrics checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest Quality of those KMs. Possible values are:</p> <p>0 = No KMs are defined for the CI Type (the color is shown as teal).</p> <p>-1 = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray).</p> <p>1 = Data was received for the metric (the color is set based on the Threshold % value).</p>
Time	The time stamp of the data.

Service KM Table

View Key Metrics current data for one or more Services in your CMDB hierarchy in a table.

The **Service KM Table** shows the same information as the [“Service KM Heatmap”](#). Use this display if, for example, you prefer to sort by **Service** or **Threshold %** to identify the Service for which you want to perform proactive health analysis.

The colors of the table rows are determined by the **Threshold %** and **Quality** values. As shown in the color gradient bar , a row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.


-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

Select **Group By Service** to include the **Group** and **Service** columns in the table. Select **Expand Metrics** to include the **Metric Name**, **Metric Value** and **Threshold** columns in the table.

Note: The **CI**s label shows the number of CIs in the table. However, if the CI is associated with multiple Services it is only counted once.

For an overview about the Key Metrics feature, see [“Key Metrics Views”](#).

For Key Metrics definitions by technology, see [“Available KM Metrics and Alerts”](#)

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Drill-down and investigate by double-clicking a row in the table to view details in the corresponding display.

← ↑ KM Table Key Metrics - Current 25-Sep-2015 14:12 Data OK

Owner: APPLICATIONS Area: SECURITY CIs: 55 Env: PRODUCTION

Group: IRIS-COMMON Service: All Services


☒ Group By Service ☐ Expand Metrics No Data No KM 0 50 100 ☐ Show CIs with KMs only (55)

Filter: CI Name=* | CI Type=* | Metric Name(disabled)=*

Group	Service	CI Type	CI Name	Threshold %	Quality	Time
IRIS-COMMON	SCAN-CHECK-1	VMWARE-HOST	vSphereW,esxi-1.west	42.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	VMWARE-VM	vSphereW,VMIRIS1051	68.8	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1001:7222:SCAN-QUEUE	22.2	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1002:7222:SCAN-QUEUE	44.4	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1003:7222:SCAN-QUEUE	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1004:7222:SCAN-QUEUE	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-SFO	23.6	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-LAX	47.2	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-SEA	18.7	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-PDX	18.7	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-SERVER	tcp://MIRIS1051:7222	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1051:7222:CHECK-QUEUE	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	ORACLE	SCAN-DB	28.6	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-2	VMWARE-HOST	vSphereE,esxi-1.east	68.1	1	25-Sep-2015 14:12:31

Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

 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.

Note: The “Up” Arrow (↑) opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking ↑ opens the “Services CI Type Summary” display.

Filter By:

The following filtering options are typically included:



Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

You can also filter KMs by **CI Name**, **CI Type** and **Metric Name**. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item in the display and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click Clear  to clear the filter.

Note: The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.

Fields and Data

This display includes:

Show CIs with KMs only (x)

When selected, hides any CIs that do not have KMs defined. The number following the label (x) is the number of CIs with KMs defined.

Group By Service

When selected, includes the **Service Group** and **Service Name** in the KM data. CIs that are included in multiple Services will appear multiple times, once for each Service they are associated with.

Expand Metrics






When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum **Threshold %** and the minimum **Quality** value for the CI.



The **No Data No KM** indicates the **Quality** value for the data. If no KMs are defined for the CI Type, the **Quality** is set to **0** and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the **Quality** is set to **-1** and the color is shown as gray. If data was received for the metric, the **Quality** is set to **1** and the color is set based on the **Threshold %** value as described above. If the **Expand Metrics** checkbox is selected, this is the Quality of a single KM. If the **Expand Metrics** checkbox is not selected, this is the lowest Quality for all of the KMs on the CI.



The gradient bar is the legend for the table row colors, which are determined by the **Threshold %** and **Quality** values. A row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

Filter:

Shows the current filter parameters and is highlighted in blue when the filter is applied.

By default, all data is shown:

CI Name=* | CI Type=* | Metric Name(disabled)=*

To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click **Clear**  to clear the filter.



Clears the filter parameters.



Applies the filter parameters.

Include Detail Level Metrics (##)

When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (x) is the number of detail level metrics available for the currently displayed KMs.

Group

The Group name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row per Group that a CI is associated with.

Service


The Service name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row Service that a CI is associated with.

CI Type

The CI Type.

CI Name

The CI Name.

Metric Name	The name of the metric. This column is only included if the Expand Metrics checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the CI Type. To see which cache column provides data for this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The CACHENAME column lists the name of the cache containing the metric and the METRICNAME column contains the name of the cache column.
Metric Value	The value of the metric. This column is only included if the Expand Metrics checkbox is selected.
Threshold	The Alarm Level value for the alert associated with the metric. This column is only included if the Expand Metrics checkbox is selected. To see which alert is associated with this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The AlertName column lists the name of the alert associated with the metric. Note: When looking up the alert threshold for a KM, RTView Enterprise Monitor first looks to see if there is an alert override on the alert where the AlertIndex exactly matches the CIName (ignoring the ~ and ; delimiters). If an exact match is found, the override Alarm Level is used. If no exact match is found, the Default Alarm Level for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.
Threshold %	The percent of the Metric Value against the Threshold . If the Expand Metrics checkbox is selected, this is the Threshold % of a single KM. If the Expand Metrics checkbox is not selected, this is the highest Threshold % for all of the KMs on the CI. Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to 0-10000 . For memory metrics, an exponential scale is applied to the Threshold % so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to 0 changing to yellow to orange to red as the value gets closer to 100 . Values at or over 100 are shown as red  . To see which CalcMode is used for this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The CalcMode column lists the type of scale that is applied to the metric. If blank, no scale is applied.
Quality	Indicates the quality of the data. If the Expand Metrics checkbox is selected, the value is for a single KM on the CI. If the Expand Metrics checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest Quality of those KMs. Possible values are: 0 = No KMs are defined for the CI Type (the color is shown as teal). -1 = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray). 1 = Data was received for the metric (the color is set based on the Threshold % value).
Time	The time stamp of the data.


Service KM History






View history heatmap of Key Metric data, over time, for a selected Group and Service.

This is the most important view for analyzing the correlation between a variety of Key Metrics over time that are related to a Service. You would navigate to this view if:

- you have identified a Service in the Alert Impact View that is having degradation right now. You can select the Service and navigate to the **Service KM History** display to determine if there are various factors causing the degradation.
- you have looked at the ["Service KM Heatmap"](#) or the ["Service KM Table"](#) and identified a Service that is about to become degraded. You can navigate to the **Service KM History** display to proactively analyze the Service before issues arise.

Each row in the history heatmap represents a different CI, unless the **Expand Metrics** checkbox is selected, in which case it represents a metric on a CI. The row color shows the **Threshold %** and **Quality** values.

The **Threshold %** value is rounded up to the closest **10** unless the **Quality** is less than **1**, in which case it shows the **Quality**. As shown in the color gradient bar , the color is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

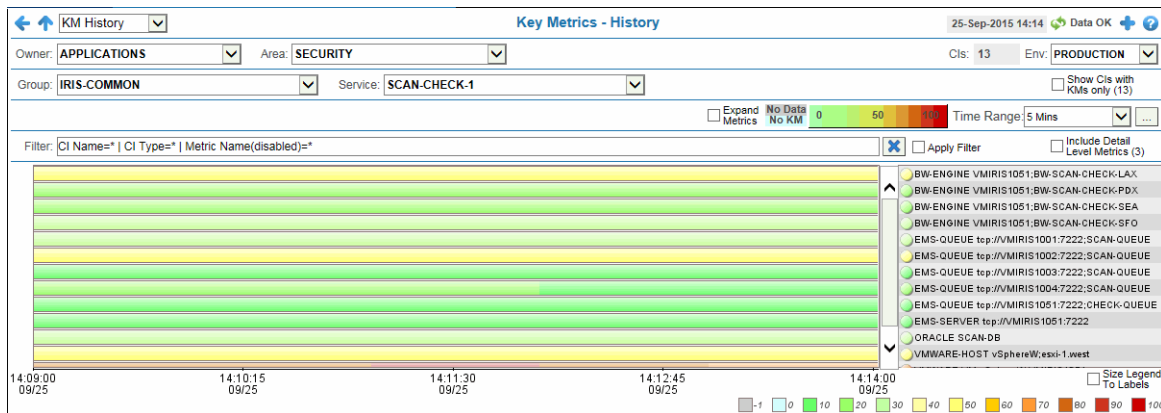
CIs associated with multiple Services are shown once for the first **Group** and **Service** they were associated with, and labeled **+ x more**, where **x** is the number of additional **Groups** and **Services** the CI is associated with.

Select **Expand Metrics** to show each Key Metric in its own row and include the **Metric Name**, **Metric Value** and **Threshold** labels in the mouseover popup window.







For an overview about the Key Metrics feature, see [“Key Metrics Views”](#).


For Key Metrics definitions by technology, see [“Available KM Metrics and Alerts”](#)


Use the available drop-down menus or right-click to filter data shown in the display. Drill-down and investigate by double-clicking a row to view details in the corresponding 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.
-   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.

Note: The “Up” Arrow () opens the most recently viewed display under [“Multi Area Service Views”](#). For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking  opens the [“Services CI Type Summary”](#) display.

Filter By:

The following filtering options are typically included:



Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

You can also filter KMs by CI Name, CI Type and Metric Name. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click Clear  to clear the filter.

Note: The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.

Fields and Data

This display includes:

Show CIs with KMs only (x)

When selected, hides any CIs that do not have KMs defined. The number following the label (x) is the number of CIs with KMs defined.

Expand Metrics






When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum **Threshold %** and the minimum **Quality** value for the CI.




The **No Data No KM** is the legend for the display colors if the **Quality** value for the data is less than **1**. If no KMs are defined for the **CI Type**, the **Quality** is set to **0** and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the **Quality** is set to **-1** and the color is shown as gray. If data was received for the metric, the **Quality** is set to **1** and the color is set based on the **Threshold %** value as described above. If the **Expand Metrics** checkbox is selected, this is the **Quality** of a single KM. If the **Expand Metrics** checkbox is not selected, this is the lowest **Quality** for all of the KMs on the CI.





The gradient bar is the legend for the display colors, which are determined by the **Threshold %** and **Quality** values. A row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.



-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

Time Range

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



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

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

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

Note: To limit the memory used by the KM history displays, the available time ranges are limited by the number of CIs in the selected **Group** and **Service**. This limit can be modified using the **\$rtvKmhHistoryRowLimit** substitution. The **\$rtvKmhHistoryRowLimit** substitution sets the maximum number of rows that can be queried by a history display and this number is used to determine the available time ranges. The default value is **35000**. To change the limit (and the maximum amount of memory used by KM history display), set the following property to a different value: **sl.rtvview.sub=\$rtvKmhHistoryRowLimit:35000**.

Filter:

Shows the current filter parameters and is highlighted in blue when the filter is applied.

By default, all data is shown:

CI Name=* | CI Type=* | Metric Name(disabled)=*

To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click Clear  to clear the filter.



Clears the filter parameters.



Applies the filter parameters.

Include Detail Level Metrics (##)

When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (**x**) is the number of detail level metrics available for the currently displayed KMs.

Mouseover

See the following details via mouseover:

Group

The **Group** name. For CIs that are associated with multiple Groups, the name of the first **Group** the CI was associated with is shown and labeled **+ x more**, where **x** is the number of additional **Groups** the CI is associated with.

Service


The **Service** name. For CIs that are associated with multiple Services, the name of the first **Service** the CI was associated with is shown and labeled **+ x more**, where **x** is the number of additional **Services** the CI is associated with.

CI Type

The CI Type.

CI Name






The CI Name.

Metric Name	The name of the metric. This is only included if the Expand Metrics checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the CI Type . To see which cache column provides data for this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The CACHENAME column lists the name of the cache containing the metric and the METRICNAME column contains the name of the cache column.
Metric Value	The value of the metric. This is only included if the Expand Metrics checkbox is selected.
Threshold	The Alarm Level value for the alert associated with the metric. This column is only included if the Expand Metrics checkbox is selected. To see which alert is associated with this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The AlertName column lists the name of the alert associated with the metric. Note: When looking up the alert threshold for a KM, RTView Enterprise Monitor first looks to see if there is an alert override on the alert where the AlertIndex exactly matches the CIName (ignoring the ~ and ; delimiters). If an exact match is found, the override Alarm Level is used. If no exact match is found, the Default Alarm Level for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.
Threshold %	The percent of the Metric Value against the Threshold . If the Expand Metrics checkbox is selected, this is the Threshold % of a single KM. If the Expand Metrics checkbox is not selected, this is the highest Threshold % for all of the KMs on the CI. Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to 0-10000 . For memory metrics, an exponential scale is applied to the Threshold % so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to 0 changing to yellow to orange to red as the value gets closer to 100 . Values at or over 100 are shown as red  . To see which CalcMode is used for this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The CalcMode column lists the type of scale that is applied to the metric. If blank, no scale is applied.
Quality	Indicates the quality of the data. If the Expand Metrics checkbox is selected, the value is for a single KM on the CI. If the Expand Metrics checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest Quality of those KMs. Possible values are: 0 = No KMs are defined for the CI Type (the color is shown as teal). -1 = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray). 1 = Data was received for the metric (the color is set based on the Threshold % value).
Time	The time stamp of the data.
Size Legend To Labels	When selected, the width of the legend is set to the widest label. When not selected, the width of the legend is set to 20% of the available space and labels that are too wide are clipped.

Service KM History (Alt)

View history heatmap of Key Metric data, over time, for a selected Group and Service. This display shows the same data as the **"Service KM History"** display but contains fewer labels. Each row in the history heatmap represents a different CI, unless the **Expand Metrics** checkbox is selected, in which case it represents a metric on a CI. The row color shows the **Threshold %** and **Quality** values.

As shown in the color gradient bar , the color is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.


CIs associated with multiple Services are shown once for the first **Group** and **Service** they were associated with, and labeled **+ x more**, where **x** is the number of additional **Groups** and **Services** the CI is associated with.

Select **Expand Metrics** to show each Key Metric in its own row and include the **Metric Name**, **Metric Value** and **Threshold** labels in the mouseover popup window.







For an overview about the Key Metrics feature, see [“Key Metrics Views”](#).


For Key Metrics definitions by technology, see [“Available KM Metrics and Alerts”](#).


Use the available drop-down menus or right-click to filter data shown in the display. Drill-down and investigate by double-clicking a row to view details in the corresponding 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.
-   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.

Note: The “Up” Arrow () opens the most recently viewed display under [“Multi Area Service Views”](#). For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking  opens the [“Services CI Type Summary”](#) display.

Filter By:

The following filtering options are typically included:



Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

You can also filter KMs by CI Name, CI Type and Metric Name. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click **Clear**  to clear the filter.

Note: The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.

Fields and Data

This display includes:

Show CIs with KMs only (x)

When selected, hides any CIs that do not have KMs defined. The number following the label (x) is the number of CIs with KMs defined.

Expand Metrics






When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum **Threshold %** and the minimum **Quality** value for the CI.





The **No Data No KM** is the legend for the display colors if the **Quality** value for the data is less than **1**. If no KMs are defined for the **CI Type**, the **Quality** is set to **0** and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the **Quality** is set to **-1** and the color is shown as gray. If data was received for the metric, the **Quality** is set to **1** and the color is set based on the **Threshold %** value as described above. If the **Expand Metrics** checkbox is selected, this is the **Quality** of a single KM. If the **Expand Metrics** checkbox is not selected, this is the lowest **Quality** for all of the KMs on the CI.






The gradient bar is the legend for the display colors, which are determined by the **Threshold %** and **Quality** values. A row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

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



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

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



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

Note: To limit the memory used by the KM history displays, the available time ranges are limited by the number of CIs in the selected **Group** and **Service**. This limit can be modified using the **\$rtvKmhHistoryRowLimit** substitution. The **\$rtvKmhHistoryRowLimit** substitution sets the maximum number of rows that can be queried by a history display and this number is used to determine the available time ranges. The default value is **35000**. To change the limit (and the maximum amount of memory used by KM history display), set the following property to a different value: **sl.rtvview.sub=\$rtvKmhHistoryRowLimit:35000**.

Filter: Shows the current filter parameters and is highlighted in blue when the filter is applied.

By default, all data is shown:

CI Name=* | CI Type=* | Metric Name(disabled)=*

To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click Clear  to clear the filter.



Clears the filter parameters.




Applies the filter parameters.

Include Detail Level Metrics (##) When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (**x**) is the number of detail level metrics available for the currently displayed KMs.

Mouseover

See the following details via mouseover:

Group	The Group name. For CIs that are associated with multiple Groups, the name of the first Group the CI was associated with is shown and labeled + x more , where x is the number of additional Groups the CI is associated with.
Service	The Service name. For CIs that are associated with multiple Services, the name of the first Service the CI was associated with is shown and labeled + x more , where x is the number of additional Services the CI is associated with.
CI Type	The CI Type.
CI Name	The CI Name.

Metric Name	The name of the metric. This is only included if the Expand Metrics checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the CI Type . To see which cache column provides data for this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The CACHENAME column lists the name of the cache containing the metric and the METRICNAME column contains the name of the cache column.
Metric Value	The value of the metric. This is only included if the Expand Metrics checkbox is selected.
Threshold	The Alarm Level value for the alert associated with the metric. This column is only included if the Expand Metrics checkbox is selected. To see which alert is associated with this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The AlertName column lists the name of the alert associated with the metric. Note: When looking up the alert threshold for a KM, RTView Enterprise Monitor first looks to see if there is an alert override on the alert where the AlertIndex exactly matches the CName (ignoring the ~ and ; delimiters). If an exact match is found, the override Alarm Level is used. If no exact match is found, the Default Alarm Level for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.
Threshold %	The percent of the Metric Value against the Threshold . If the Expand Metrics checkbox is selected, this is the Threshold % of a single KM. If the Expand Metrics checkbox is not selected, this is the highest Threshold % for all of the KMs on the CI. Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to 0-10000 . For memory metrics, an exponential scale is applied to the Threshold % so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to 0 changing to yellow to orange to red as the value gets closer to 100 . Values at or over 100 are shown as red  . To see which CalcMode is used for this metric, navigate to Architecture - "RTView KM Defs" . In the table, look in the CITYPE and SELECTOR columns to find the row for your metric. The CalcMode column lists the type of scale that is applied to the metric. If blank, no scale is applied.
Quality	Indicates the quality of the data. If the Expand Metrics checkbox is selected, the value is for a single KM on the CI. If the Expand Metrics checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest Quality of those KMs. Possible values are: 0 = No KMs are defined for the CI Type (the color is shown as teal). -1 = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray). 1 = Data was received for the metric (the color is set based on the Threshold % value).
Time	The time stamp of the data.

Available KM Metrics and Alerts

This section lists available Key Metrics and their associated alerts.

- "Amazon AWS"
- "Host Agent"
- "IBM DB2"
- "IBM MQ"
- "IBM WebSphere"
- "JBoss"
- "Kafka"
- "Oracle Coherence"
- "Oracle Database"
- "Oracle WebLogic"
- "RTViewManager"
- "RTVRULES"
- "Solace"
- "TIBCO ActiveMatrix"
- "TIBCO ActiveSpaces"
- "TIBCO BusinessEvents"
- "TIBCO BusinessWorks (Version 5) Monitor"
- "TIBCO BusinessWorks (Version 6) Monitor"
- "TIBCO EMS"
- "UX"

Amazon AWS

The following KMs are available with the Solution Package for Amazon AWS. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
ACW	AwsEc2InstanceStats	Instance CPU Usage	CPUUtilization / AcwInstanceCpuHigh

Host Agent

The following KMs are available with the Solution Package for RTView Host Agent. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
HOST	HostStats	% CPU Utilization	usedPerCentCpu / HostCpuPercentHigh
HOST	HostStats	% Memory Used	MemUsedPerCent / HostMemoryUsedHigh

IBM DB2

The following KMs are available with the Solution Package for IBM DB2. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
DB2	Db2ResponseTime	Response Time	ResponseTimeMilliSec / Db2ResponseTimeHigh

IBM MQ

The following KMs are available with the Solution Package for IBM WebSphere MQ. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
MQ-BROKER	MqBrokers	Queue Depth	Current queue depth / MqBrokerQueueDepthHigh
MQ-QUEUE	MqQueues	Queue Depth	Current queue depth / MqQueueDepthHigh

IBM WebSphere

The following KMs are available with the Solution Package for IBM WebSphere. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
WAS	WasServerStats	Live Session Count	LiveCount / WasLiveSessionCountHigh
WAS	WasServerStats	WAS CPU %	ProcessCpuUsage / WasJvmCpuHigh
WAS	WasServerStats	Memory Used %	usedMemoryPercent / WasMemoryUsedPercentHigh
WAS-APP	WasServletTotalsByApp	Response Time	responseTime / WasServletResponseTimeHigh
WAS-APP	WasServletTotalsByApp	Requests / sec	DeltatotalRequests / WasServletRequestRateHigh

JBoss

The following KMs are available with the Solution Package for JBoss. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
JBOSS-APP	JbossDeployments	Active Sessions	activeSessions / JbossAppActiveSessionsHigh
JBOSS-SERVER	JbossServerStats	% Process CPU	ProcessCpuLoadPercent / JbossServerProcessCpuLoadHigh
JBOSS-SERVER	JbossDeploymentTotals	Active Sessions	activeSessions / JbossServerActiveSessionsHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)

Kafka

The following KMs are available with the Solution Package for Kafka. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
KAFKA-BROKER	KafkaServerJvm	CpuPercent	CpuPercent / KafkaBrokerCpuPercentHigh
KAFKA-BROKER	KafkaServerJvm	MemoryUsedPercent	MemoryUsedPercent / KafkaBrokerMemoryUsedPercentHigh
KAFKA-CONSUMER	KafkaConsumer	Consumer Lag	records-lag-max / KafkaConsumerLagHigh
KAFKA-CONSUMER	KafkaConsumer	CpuPercent	CpuPercent / KafkaConsumerCpuPercentHigh
KAFKA-CONSUMER	KafkaConsumer	MemoryUsedPercent	MemoryUsedPercent / KafkaConsumerMemoryUsedPercentHigh
KAFKA-PRODUCER	KafkaProducer	CpuPercent	CpuPercent / KafkaProducerCpuPercentHigh
KAFKA-PRODUCER	KafkaProducer	IO Wait Time	io-wait-time-millis-avg / KafkaProducerIoWaitTimeMSHigh
KAFKA-PRODUCER	KafkaProducer	MemoryUsedPercent	MemoryUsedPercent / KafkaProducerMemoryUsedPercentHigh
KAFKA-ZOOKEEPER	KafkaZookeeper	Avg Request Latency	AvgRequestLatency / KafkaZookeeperAvgLatencyHigh

Oracle Coherence

The following KMs are available with the Solution Package for Oracle Coherence. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.) For alert definitions, see *RTView® Oracle® Coherence Monitor User's Guide*.

CI Type	Cache	Selector	Metric / Alert
OC-CACHE	OcCacheTotals	Rate Cache Misses	RateCacheMisses / OcCacheRateCacheMissesHigh This metric is the rate of cache misses against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.

OC-CACHE	OcCacheTotals	Rate Store Reads	<p>RateStoreReads / OcCacheRateStoreReadsHigh</p> <p>The level of this Key Metric is 1. (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)</p> <p>This metric is the rate of store reads (load operations) against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
OC-CACHE	OcCacheTotals	Rate Store Writes	<p>RateStoreWrites / OcCacheRateStoreWritesHigh</p> <p>The level of this Key Metric is 1. (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)</p> <p>This metric is the rate of store writes (store and erase operations) against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
OC-CACHE	OcCacheTotals	Queue Size	<p>QueueSizePos / OcCacheQueueSizeHigh</p> <p>This metric is the cache send queue size for a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
OC-CACHE	OcCacheTotals	Rate Cache Puts	<p>RateTotalPuts / OcCacheRateTotalPutsHigh</p> <p>This metric is the rate of cache puts against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
OC-CACHE	OcCacheTotals	Rate Cache Gets	<p>RateTotalGets / OcCacheRateTotalGetsHigh</p> <p>This metric is the rate of cache gets against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
OC-CACHE	OcCacheTotals	Rate Store Reads	<p>RateStoreReads / OcCacheRateStoreReadsHigh</p> <p>The level of this Key Metric is 1. (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)</p>

OC-CACHE	OcCacheTotals	Rate Store Writes	RateStoreWrites / OcCacheRateStoreWritesHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
OC-CACHE	OcCacheTotals	CacheSize	CacheSize / OCCacheSizeHigh The level of this Key Metric is 0 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.) This metric is the number of objects in a cache for a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.
OC-CLUSTER	OcPacketStats	Packet Loss	SentFailureRate / OcBadCommunicationCluster This metric is the (network/packet) sent failure rate averaged across all of the nodes of a cluster.
OC-CLUSTER NODES	OcNodeTotals	CPU Used %	AvgCpuPercent / OcClusterNodesCPUHigh This metric is the average CPU usage of all the nodes of a given storage class in a cluster. The storage class is represented by the StorageEnabled index column, which can be true or false . Thus metrics for storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = true , and non storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = false . This metric is shown as a trace in the Cluster - Memory/Network Health display. The metric is labeled Avg. CPU% and is displayed (for storage enabled nodes) in the Storage Nodes trend grouping and (for non storage enabled nodes) in the Process Nodes trend grouping.
OC-CLUSTER NODES	OcNodeTotals	Packet Rx Loss	RcvdFailureRate100 / OcClusterNodesRcvdFailureRateHigh This metric is the (network/packet) received failure rate averaged across all of the nodes of a given storage class in a cluster. The storage class is the StorageEnabled index column, which can be true or false . Metrics for storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = true , and non storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = false .

OC-CLUSTER NODES	OcNodeTotals	Memory Used %	MemoryUsedPct100 / OcClusterNodesMemHigh This metric is the memory used percentage averaged across all of the nodes of a given storage class in a cluster. The storage class is the StorageEnabled index column, which can be true or false . Metrics for storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = true , and non storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = false .
OC-CLUSTER NODES	OcNodeTotals	Packet Tx Loss	SentFailureRate100 / OcClusterNodesSentFailureRateHigh This metric is the (network/packet) sent failure rate averaged across all of the nodes of a given storage class in a cluster. The storage class is the StorageEnabled index column, which can be true or false . Metrics for storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = true , and non storage enabled nodes in a cluster are aggregated into a cache row where StorageEnabled = false .

Oracle Database

The following KMs are available with the Solution Package for Oracle Database. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
ORACLE	OraDatabaseAvailability	Response Time	ResponseTimeMilliSec / OraDatabaseResponseTimeHigh

Oracle WebLogic

The following KMs are available with the Solution Package for Oracle WebLogic. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
WLS	WlsJvmStats	JVM CPU %	JvmProcessorLoad / WlsServerCpuHigh
WLS	WlsJvmStats	JVM Memory %	MemoryUsedPercent / WlsServerMemoryUsageHigh
WLS	WlsThreadPoolRuntime	Hogging Threads	HoggingThreadCount / WlsHoggingThreadsHigh
WLS	WlsServerRuntime	Open Sockets	OpenSocketsCurrentCount / WlsOpenSocketsHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
WLS	WlsThreadPoolRuntime	Thread Total Count	ExecuteThreadTotalCount / WlsThreadsTotalHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)

WLS-APP	WlsSessionStats	Open Sessions	OpenSessionsCurrentCount / WlsAppOpenSessionsHigh
WLS-JMS-DEST	WlsJmsDestinationTotals	Messages Pending	MessagesPendingCount / WlsJmsDestinationMessagesPendingHigh
WLS-JMS-SERVER	WlsJmsServerRuntime	Messages Pending	MessagesPendingCount / WlsJmsMessagesPendingHigh

RTViewManager

The following KMs are available with the RTView Manager Solution Package. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
JVM	JvmOperatingSystem	Cpu %	CpuPercent / JvmCpuPercentHigh
JVM	JvmMemory	Memory %	MemoryUsedPercent / JvmMemoryUsedHigh
JVM	JvmThreading	Thread Count	ThreadCount / JvmThreadCountHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
TOMCAT	TomcatWebModuleTotals	Active Sessions	activeSessions / TomcatActiveSessionsHigh
TOMCAT	TomcatWebModuleTotals	Accesses / sec	RateaccessCount / TomcatAccessRateHigh
TOMCAT-APP	TomcatWebModuleStats	Active Sessions	activeSessions / TomcatAppActiveSessionsHigh
TOMCAT-APP	TomcatWebModuleStats	Accesses / sec	RateaccessCount / TomcatAppAccessRateHigh

RTVRULES

The following KMs are available with the RTVRULES Solution Package which comes with RTView EM. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
EM-SERVICE	RtvCmdServiceStats_local	Alert Impact	AlertImpact / RtvEmServiceAlertImpactHigh

Solace

The following KMs are available with the Solution Package for Solace. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
SOLACE-MSGROUTER	SolAppliances	# Msgs Spooled	num-messages-spoiled / SolMsgRouterPendingMsgsHigh
SOLACE-MSGROUTER	SolAppliances	OUT Msgs/sec	total-cl-msgs-sent-per-sec / SolMsgRouterOutboundMsgRateHigh
SOLACE-MSGROUTER	SolAppliances	IN Msgs/sec	total-cl-msgs-rcvd-per-sec / SolMsgRouterInboundMsgRateHigh
SOLACE-VPN	SolVpns	Connections	connections / SolVpnConnectionCountHigh

TIBCO ActiveMatrix

The following KMs are available with the Solution Package for TIBCO ActiveMatrix. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
AMX-SERVICE	AmxServiceTotals	Service Hits/Min	Hits Per Minute / AmxServiceHitRateHigh
AMX-SERVICE	AmxServiceTotals	Service Response Time	Avg. Response Time / AmxServiceResponseTimeHigh
AMX-SERVICE NODE	AmxServices	Node Hits/Min	Hits Per Minute / AmxServiceNodeHitRateHigh
AMX-SERVICE NODE	AmxServices	Node Response Time	Avg. Response Time / AmxServiceNodeResponseTimeHigh

TIBCO ActiveSpaces

The following KMs are available with the Solution Package for TIBCO ActiveSpaces. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
TAS-MEMBER BYSPACE	TasSeeders	Space Util by Seeder	spaceUtilPerSeeder / TasMemberSeederCapacity
TAS-SPACE	TasSpaceStatistics	Gets/sec	RateGets / TasSpaceGetRateHigh
TAS-SPACE	TasSpaceStatistics	Puts/sec	RatePuts / TasSpacePutRateHigh

TIBCO BusinessEvents

The following KMs are available with the Solution Package for TIBCO BusinessEvents. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
TBE-CLUSTER	TbeClusterSummary	Received Events Rate	Received Events Rate / TbeClusterEventsRecvdRateHigh
TBE-CLUSTER	TbeClusterSummary	Rules Fired Rate	totalRateTotalNumberRulesFired / TbeClusterRuleFiringRateHigh
TBE-CLUSTER	TbeClusterSummary	Concept Cache Ops Rate	totalConceptOperationRate / TbeClusterConceptOpRateHigh
TBE-CLUSTER	TbeClusterSummary	Backing Store Ops Rate	totalBkngStoreOpsPerSec / TbeClusterBkngStoreOpRateHigh

TIBCO BusinessWorks (Version 5) Monitor

The following KMs are available with the Solution Package for TIBCO BusinessWorks version 5. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
BW-ENGINE	BwEngines	CPU Used %	CPU % / BwEngineCpuUsedHigh
BW-ENGINE	BwEngines	Memory Used %	PercentUsed / BwEngineMemUsedHigh
BW-PROCESS	BwProcesses	AverageElapsed	Process Avg Elapsed Time / BwProcessAvgElapsedTimeHigh
BW-PROCESS	BwProcesses	RateCreated / sec	Processes Created/sec / BwProcessCreatedRateHigh
BW-PROCESS	BwProcesses	TotalCpuPercent	Process Total CPU Percent / BwProcessTotalCpuPercentHigh
BW-PROCESS	BwProcesses	Process Exec Time / sec	RateTotalExecution / BwProcessExecutionTimeHigh
BW-SERVER	BwServers	CPU Used %	CPU Usage % / BwServerCpuUsedHigh

TIBCO BusinessWorks (Version 6) Monitor

The following KMs are available with the Solution Package for TIBCO BusinessWorks version 6. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
BW6-APPNODE	Bw6AppNodes	CPU Used %	Used CPU Percentage / Bw6AppNodeCpuUsedHigh
BW6-APPNODE	Bw6AppNodes	Mem Used %	Used Memory Percentage / Bw6AppNodeMemUsedHigh
BW6-APP	Bw6ProcessTotalsByApp	App Created / sec	RateCreated / Bw6AppProcessCreatedRateHigh
BW6-APP	Bw6ProcessTotalsByApp	App Exec Time / sec	RateTotal Execution / Bw6AppProcessExecutionTimeHigh
BW6-PROCESS	Bw6Processes	Process Created / sec	RateCreated / Bw6ProcessCreatedRateHigh
BW6-PROCESS	Bw6Processes	Process Exec Time / sec	RateTotal Execution / Bw6ProcessExecutionTimeHigh

TIBCO EMS

The following KMs are available with the Solution Package for TIBCO EMS. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
EMS-QUEUE	EmsQueues	Pending Msgs	pendingMessageCount / EmsQueuesPendingMsgsHigh
EMS-QUEUE	EmsQueues	In Msgs / sec	inboundMessageRate / EmsQueuesInMsgRateHigh
EMS-QUEUE	EmsQueues	Out Msgs / sec	outboundMessageRate / EmsQueuesOutMsgRateHigh
EMS-QUEUE	EmsQueues	Consumers	consumerCount / EmsQueuesConsumerCountHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
EMS-SERVER	EmsServerInfo	Pending Msgs	pendingMessageCount / EmsServerPendingMsgsHigh
EMS-SERVER	EmsServerInfo	In Msgs / sec	inboundMessageRate / EmsServerInMsgRateHigh
EMS-SERVER	EmsServerInfo	Out Msgs / sec	outboundMessageRate / EmsServerOutMsgRateHigh

EMS-SERVER	EmsServerInfo	Msg Mem %	messageMemoryPct / EmsServerMemUsedHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
EMS-SERVER	EmsServerInfo	Connections	connectionCount / EmsServerConnectionCountHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
EMS-SERVER	EmsServerInfo	Async DB Size	asyncDBSize / EmsServerAsyncDBSizeHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
EMS-SERVER	EmsServerInfo	Sync DB Size	syncDBSize / EmsServerSyncDBSizeHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
EMS-TOPIC	EmsTopics	Pending Msgs	pendingMessageCount / EmsTopicsPendingMsgsHigh
EMS-TOPIC	EmsTopics	In Msgs / sec	inboundMessageRate / EmsTopicsInMsgRateHigh
EMS-TOPIC	EmsTopics	Out Msgs / sec	outboundMessageRate / EmsTopicsOutMsgRateHigh
EMS-TOPIC	EmsTopics	Consumers	consumerCount / EmsTopicsConsumerCountHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)
EMS-TOPIC	EmsTopics	Subscribers	subscriberCount / EmsTopicsSubscriberCountHigh The level of this Key Metric is 1 . (Level 0 KMs are always displayed. Level 1 KMs are displayed if the Include Detail Level Metrics checkbox is checked.)

UX

The following KMs are available with the Solution Package for UX. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
UX-URL	UXURLData	Response Time	MostRecentTime / UXURLResponseSlow

VMWare vSphere

The following KMs are available with the Solution Package for VMWare vSphere. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
VMWARE-HOST	VmwHostSystems	CPU Usage	cpu.usage.average / VmwHostCpuUtilizationHigh
VMWARE-HOST	VmwHostSystems	Memory Usage	mem.usage.average / VmwHostMemoryUsageHigh
VMWARE-VM	VmwVirtualMachines	CPU Usage	cpu.usage.average / VmwVmCpuUtilizationHigh
VMWARE-VM	VmwVirtualMachines	Memory Usage	mem.usage.average / VmwVmMemoryUsageHigh

Component Views


These displays present the lowest level view of CMDB contents--the component level. In these displays, alert states for components are shown by Service and Area in tabular and heatmap formats, while highlighting the most critical alert state for each component. Data can be filtered by Areas, Services, Groups, Regions and Environment. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.

Use these displays to determine whether a component is malfunctioning. Displays in this View are:

- **"CI / Service Tree View"**: Table of CMDB contents for all component-level details by Service for all Owners, Areas, Groups, Regions and Environments (without the option to filter).
- **"CI / Service Table"**: Table of CMDB contents for all component-level details by Service for all Owners, Areas, Groups, Regions and Environments (without the option to filter).
- **"CI / Type Heatmap"**: Heatmap of CMDB contents organized by CType, with the option to filter by Owner, Area, Group, Environment and alert Metric, and show CI Names.
- **"CI / Type Table"**: Table of CMDB contents for all component-level details for all Areas, Services, Groups, Regions and Environments, with the option to filter by Owner and one or all Areas, Groups and Environments.

CI / Service Tree View

View the contents of the CMDB hierarchically ordered in a navigation tree. Each row in the table is a different CI (for example, **localhost;RTVMGR_DATASERVER**).

Make a selection from the **Owner** drop-down menu, then use the navigation tree to filter data in the **Service CI Table**. The navigation tree, which provides a visual of the CMDB hierarchy, provides further filtering to the **Area**, **Group**, and **Environment** drop-down menus. Click Sort  to order column data.

Tree

Service Components24-Sep-2015 11:24Data OK

Owner: InfrastructureArea: All AreasGroup: All Service GroupsEnv: QA

Service CI Table - Current StatusNumber of Components: 5,487

Infrastructure

Middleware

IBM-MQ

TOMCAT

SOLACE

TIBCO-BW

TIBCO-AMX

TIBCO-BE

TIBCO-EMS

Processes

JVM

Servers

Hosts

Databases

CIType	CIName	Severity	Alerts	Criticality	Environment	Regi
MQ-BROKER	vmrh5-1		1	E	DEMOSITE	AMER
MQ-BROKER	vmrh5-2		1	E	DEMOSITE	AMER
MQ-BROKER	vmrh5-3		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_01		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_02		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_03		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_04		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_05		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_06		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_07		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_08		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_09		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_10		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_01		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_02		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_03		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_04		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_05		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_06		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_07		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_08		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_09		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_10		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_01		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_02		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_03		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_04		0	E	DEMOSITE	AMER

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.

Row Color Code:

Tables with colored rows indicate the following:

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

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.




Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.


Fields and Data

This display includes:

Number of Components	The total number of CIs currently in the table.
CIType	The type of CI.
CIName	The name or address of the CI.
Severity	<p>The maximum level of alerts for the CI. Values range from 0 to 2, where 2 is the greatest Alert Severity:</p> <ul style="list-style-type: none">  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
Criticality	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E, where A is the highest Criticality. This value is used to determine the value for Alert Impact.
Environment	The Environment for the CI.
Region	The name of the Region for the CI.
City	The name of the City for the CI.
Country	The name of the Country for the CI.
SiteName	The name of the Site for the CI.
OSType	The operating system currently running on the CI.
City	The name of the City for the CI.
Country	The name of the Country for the CI.

CI / Service Table

View the contents of the CMDB, without filtering, in a tabular format. Each row in the table is a different CI (for example, **localhost;RTVMGR_DATASERVER**).

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


← CMDB ▾ All Components By Service CMDB 24-Sep-2015 11:22 Data OK + ?

Service CI Table - Current Values Number of Rows: 5,784

Owner	Area	ServiceGroup	ServiceName	CIType	CIName	Se
Jerelyn Parker	Backends	IBM	MQ	MQ-BROKER	vmrh5-1	
Jerelyn Parker	Backends	IBM	MQ	MQ-BROKER	vmrh5-2	
Jerelyn Parker	Backends	IBM	MQ	MQ-BROKER	vmrh5-3	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_01	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_02	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_03	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_04	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_05	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_06	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_07	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_08	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_09	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_10	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_01	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_02	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_03	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_04	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_05	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_06	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_07	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_08	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_09	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_10	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_01	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_02	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_03	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_04	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_05	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_06	

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.

Row Color Code:

Tables with colored rows indicate the following:

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

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.




Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Fields and Data

This display includes:

Number of Rows	The current total number of rows in the table.
Service CI Table	
Owner	The Owner the CI is associated with.
Area	The Area the CI is associated with.
ServiceGroup	The Group the CI is associated with.
ServiceName	The Service the CI is associated with.
CIType	The type of CI.
CIName	The name or address of the CI.
Severity	<p>The maximum level of alerts for the CI. Values range from 0 to 2, where 2 is the greatest Alert Severity:</p> <ul style="list-style-type: none">  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
Criticality	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views - CI Service Table display, which range from A to E , where A is the highest Criticality. This value is used to determine the value for Alert Impact.
Environment	The Environment for the CI.
City	The name of the City for the CI.
Country	The name of the Country for the CI.
Region	The name of the Region for the CI.
SiteName	The name of the Site for the CI.

CI / Type Heatmap

View heatmap of alert states for CIs in all or one Area, Group or Environment. The heatmap organizes CIs by CI Type, and uses color to show the most critical alert state for each. Each rectangle in the heatmap represents a CI (for example, **localhost;RTVMGR_DATASERVER**).

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Double-click (or right-click and select **Drill Down**) a rectangle in the heatmap to view details relevant to the CI Type. By default, this display shows all Areas, Groups, and Environments and alert Impact.



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.
- 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.

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Metric:


Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:


Alert Impact


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

Alert Severity


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

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


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

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

Alert Count

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


Criticality

























The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from **1** to **5**, as indicated in the color gradient  bar, where **5** is the highest Criticality.


Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the **Component Views** - "[CI / Service Table](#)" display, which range from **A** to **E**, where **A** is the highest Criticality (level **5** maps to a Criticality of **A** and level **1** maps to a Criticality of **E** with equally spaced intermediate values).

CI / Type Table

View tabular list of all CIs by CType, as well as their alert metrics (Impact, Severity and Count, for one or all Areas, Groups or Environments). Each row in the table is a different CI (for example, **localhost;RTVMGR_DATASERVER**). The row color represents the most critical alert state for the CI.




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

All Components By Type				
24-Sep-2015 11:16 Data OK				
Owner: Infrastructure	Area: All Areas	CIs: 64	Env: QA	
Group: All Service Groups				
CType	CName	Severity	AlertCount	AlertImpact
EMS-QUEUE	tcp://192.168.200.132:7222;queue.sample		3	2
EMS-QUEUE	tcp://192.168.200.132:7222;sample		3	2
EMS-SERVER	tcp://192.168.200.132:7222		2	2
EMS-TOPIC	tcp://192.168.200.132:7222;sample		2	2
EMS-TOPIC	tcp://192.168.200.132:7222;topic.sample		2	2
EMS-TOPIC	tcp://192.168.200.132:7222;topic.sample.exported		2	2
EMS-TOPIC	tcp://192.168.200.132:7222;topic.sample.imported		2	2
HOST	QATB:SLHOST-WIN3		2	2
HOST	QATB:SLHOST-WIN4		1	1
JVM	localhost:ALERTHISTORIAN		1	1
JVM	localhost:ALERT_SERVER		3	2
JVM	localhost:AMXMON-HISTORIAN		1	2
JVM	localhost:AMXMON-SLHOST-WIN3		2	1
JVM	localhost:AMXMON-SLHOST-WIN4		2	1
JVM	localhost:BWMON-HISTORIAN		1	2
JVM	localhost:BWMONITOR-WIN-8		1	2
JVM	localhost:CONFIG_SERVER		2	1
JVM	localhost:DISPLAYSERVER		2	1
JVM	localhost:DISPLAYSERVER_DARKSTYLES		2	1
JVM	localhost:EMSMON-HISTORIAN		1	2
JVM	localhost:EMSMON-SLHOST-WIN3		2	2
JVM	localhost:EMSMON-SLHOST-WIN4		2	1
JVM	localhost:EMSMONITOR-WIN-8		2	2
JVM	localhost:MISCMON-HISTORIAN		1	2



Row Color Code:

Tables with colored rows indicate the following:

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

Filter By:

The following filtering options are typically included:

Owner: Choose an Owner to see metrics for Areas associated with that Owner.

Area: Choose an Area to see metrics for Groups associated with that Area and Owner.

Group: Choose a Group to see metrics for Services associated with that Group, Area and Owner.

Service: Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

Env: Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

Fields and Data

This display includes:

CI Count The total number of CIs listed in the table. This value is determined by the selections made from display drop-down menus. The totals number for each Environment are also shown.




CI Table

This table lists all CIs for the selected Group. Each row in the table is a CI. Each CI can have multiple alerts. Click a CI to view alerts for the CI in the lower table.

CIType The type of CI.

CIName The name or address of the CI.

Severity The maximum level of alerts for the CI. Values range from 0 to 2, where 2 is the greatest Alert Severity:

-  One or more alerts exceeded their ALARM LEVEL threshold.
-  One or more alerts exceeded their WARNING LEVEL threshold.
-  No alert thresholds have been exceeded.

Alert Count The total number of critical and warning alerts for the CI.

Alert Impact The product of the maximum Alert Severity multiplied by the maximum Criticality of alerts. Values range from **0** - **10**, where **10** is the highest Alert Impact.

Metric Explorer

The Metric Explorer (MX) is a tool for creating and viewing custom dashboards, referred to as *MX Views*. An MX View contains a trend graph with up to five traces which you can configure to show numeric metrics from any EM solution package. While EM provides out-of-the-box Views of metric data, there might not be a single display that shows all the metrics that are critical to a single application. MX allows end-users to create Views containing the metrics that are important to them. The MX Views your end-users create are accessed from the MX **View** drop-down menu (rather than the navigation tree as RTView Enterprise Monitor Views are). Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.

Displays in this View are:


- “Metric Explorer”:

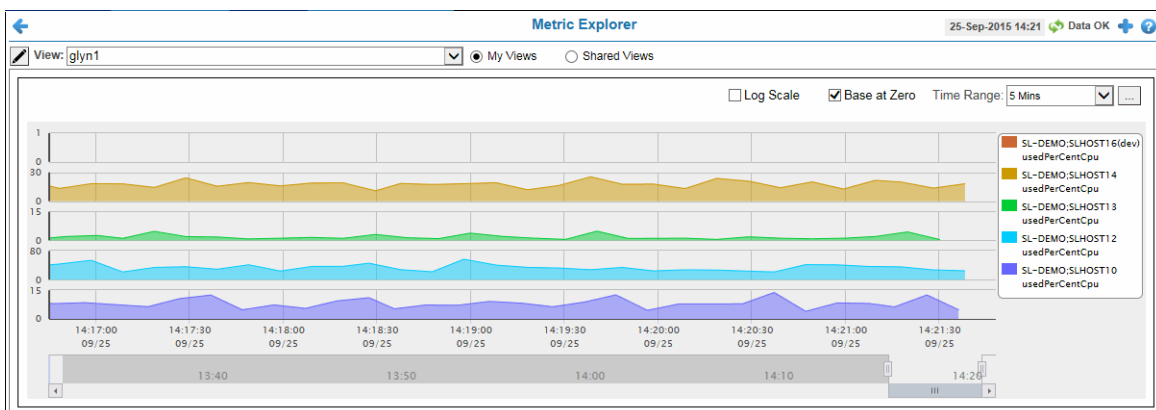
Note: The Metric Explorer was added in RTView Enterprise Monitor version 1.5.0. For instructions about adding the Metric Explorer to applications created with versions older than 1.5.0, see the RTView Enterprise Monitor Upgrade Notes.


Metric Explorer

View your previously created MX Views. Select an MX View from the **View** drop-down menu. The contents of the **View** drop-down menu depend on whether you choose **My Views** or **Shared Views**. Choose **My Views** to see public and private MX Views owned by you. Choose **Shared Views** to see public MX Views owned by you and other users. A public MX View is an MX View where the creator chose the **Share View with Others** option. The creator of the MX View is the owner.

Each MX View has options to apply **Log Scale**, **Base at Zero** and **Time Range** to your graphs.

To create or edit an MX View click Edit  to open the edit pane. For details, see “[Creating MX Views](#)”.





Fields and Data

Options include:



Open the edit pane.

View

Select an MX View from the **View** drop-down menu.

My Views

Choose **My Views** to see public and private MX Views owned by you in the **View** drop-down menu.

Shared View

Choose **Shared Views** to see public MX Views owned by you and other users. A public MX View is an MX View where the creator chose the **Share View with Others** option. The creator of the MX View is the owner.

Log Scale

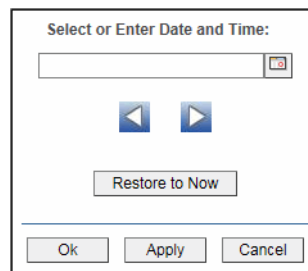
Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale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

Base at Zero

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

Time Range

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




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

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






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



Creating MX Views


Click Edit  to open the edit pane. If an MX View is already selected, click New  to start a new MX View.



Select the Service containing the metric you want to see from the **Metrics Options/Service Tree**. The **Service Tree** is filtered by the environment in the **Env** menu. If necessary select the environment containing your Service. When you select a Service, all available metrics for that Service are listed in the **Metric Tree**. Because the metric will be displayed in a trend graph, only numeric metrics with history are listed in the **Metric Tree**. Select the metric you want to see from the **Metrics Options/Metric Tree** and click **Add Metric**. The metric is added to the **Selected Metrics** list and the MX View preview (in the right panel). Add up to five metrics to your MX View.



To change the order in which the metrics are displayed in the graph use the Up  and Down  arrows. To remove a metric, select it in the **Selected Metrics** and click Trash . To add a label to your metric, select it in the **Selected Metrics** list and enter your label text in the **Label** field. Click Apply  to apply the label, or Cancel  to cancel the label.

Click **Save** and enter a descriptive MX View name. Click **Share View with Others** to make your MX View public, otherwise, the MX View is only available to you. Click Confirm  to write the MX View to the database. Click Cancel  to return to edit mode. Click **Done** to return to the **Metric Explorer** page. The MX View you created is added to the **View** drop-down menu.

To create a new MX View with the Service already selected, select a Service from a **Service Summary Views** display and click MX  (or the table context menu). This opens the MX edit pane with the Service already selected in the MX edit pane **Service Tree**. If you selected a CI Type or CI, these are also already selected in the MX edit pane **Metric Tree**. This spares you from having to search for the Service, CI Type or CI in the **Service** and **Metrics Trees**. The displays from which you can use this feature are:

- **Service Summary Views** - "Service By CI Type"
- **Service Summary Views** - "Service Summary"





Editing MX Views

In the Metric Explorer, select the MX View you want to edit and click Edit . The edit pane opens with the selected MX View in edit mode. To delete the MX View click Trash . To save your MX View under a new name, click **Save As**. Add, remove, reorder or label metrics as described in the "Creating MX Views" section (above). Select a metric in the **Selected Metrics** list and click on **Search** to update the selection in the **Service Tree** and **Metric Tree** to the values used when that metric was added to the MX View. This is useful when you want to see which Service contains a metric so you can add more metrics from the same Service.

When you are finished editing your metric, you can click **Cancel** to cancel your changes or **Save** to save your changes. To edit another MX View, select it from the **View** drop-down menu. Click **Done** to return to the Metric Explorer page.







Note: When you edit an MX View you do not own a copy of the MX View is automatically created and you are prompted to enter a name for the MX View when you save it.

View Options


	Create a new MX View.
	Delete the selected MX View.
Save	Save the selected MX View. If this is an existing MX View, the save is done immediately. If this is a new MX View, the Name field becomes available and you must enter a name and click Confirm Save to save your MX View.
Save As	Save the selected MX View under a new name. The Name field becomes available and you must enter a name and click Confirm Save to save your MX View.
Done	Close the edit pane. This option is available when you do not have unsaved changes.
Cancel	Cancel your edits.
Name	Enter a name for your MX View. This field is available when saving a new MX View or after you click Save As .
	Confirm that you want to save your MX View after you enter a name. This option is available when saving a new MX View or after you click Save As .
	Cancel the save. This is available when saving a new MX View or after you click Save As .
Share View	Select to make your MX View public. Public MX Views are available to all users in the View drop-down menu when the Shared Views option is selected. Deselect to make this MX View only available to you.

Metric Options

Env	Select an Environment to filter the items in the Service Tree.
Service Tree	The CMDB service model (Owner, Area, Group, Service). Select a Service to populate the Metric Tree with metrics for that Service. The Services in the Service Tree are filtered by the following login substitutions: \$rtvOwnerMask , \$rtvAreaMask , \$rtvGroupMask and \$rtvServiceMask . For details, refer to the <i>RTView Enterprise Monitor Configuration Guide</i> .
Metric Tree	The available metrics for the selected service. The tree hierarchy is CI Type, CI name, Metric (cache: metric). The tree only contains numeric metrics with history.
Add Metric	Add the selected metric to the MX View. When a metric is added to the MX View, it appears in the graph.
Selected Metrics	The list of metrics for this MX View.

	Move the metric up in the list of selected metrics.
	Move the metric down in the list of selected metrics.
	Set the selection in the Service and Metric trees to the values used when you added the selected metric to the MX View. Note: If your CMDB has changed such that the Service you used to add this metric no longer exists, the search button will fail
	Delete the selected metric from the MX View.
Label	Enter a label to use for the selected metric. This label is not applied until you click on the confirm label button. This label is used in the graph legend.
	Confirm the label you entered for the selected metric.
	Discard the label you entered for the selected metric (revert back to the previously applied value).

Limitations

- The Search  button fails without an error if the Service that was selected when you initially added the metric is no longer in your CMDB. To fix this, delete the metric and add it again from a Service that is currently in your CMDB. **Note:** The missing Service only makes the Search button fail. It does not cause any problems with viewing the metric.
- When you try to add a metric to an MX View that already contains that metric, it will not be added again. In the Viewer, an error message will come up saying that the metric is already in the MX View. In the Thin Client, no error is shown.
- MX Views are limited to five metrics. After a view contains five metrics, the **Add Metric** button is disabled.
- There is no indicator that shows if the MX database or Central Configuration Server are off-line in the MX configuration display. Any changes you save when either the MX database or Central Configuration Server are off line will be lost.
- When you save an MX View, RTView writes to both the View Table and the Metrics Table to the database even if only one or the other changed.
- When you save an MX View, the MX Configuration UI temporarily reverts back to the previous version of the MX View for one update, then updates with the latest changes.
- By default, MX attaches to the history_combo table for the metric history. If the cache is not configured with a history_combo table, the Metric Explorer will instead make a one-time attachment to the history table. In this case, toggling the **Log Scale** check-box will cause all points plotted after the initial history query to be lost. On the next update of current data a straight line will be drawn from the last history point to the new current data point.

Alert Views

These displays present detailed information about all alerts that have occurred in your RTView Enterprise Monitor system (all Owners and all Areas). The type of alerts that appear in these displays depends on the solution packages installed on your RTView Enterprise Monitor system. Displays in this View are:

- **"RTView Alerts Table"**: Shows current alert data. Use this time-ordered tabular view to track, manage and assign alerts.
- **"Alert History Table"**: Shows historical alert data. Use this time-ordered tabular view to track alert status changes.

RTView Alerts Table

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

The color coded navigation tree shows the contents of the CMDB hierarchically ordered. Choose a node to filter alerts shown in the table. The **Alerts Table** only shows alerts associated with the node you select. A green indicator means the node has no associated alerts. A red indicator means the node has one or more associated alerts.

Service name labels are appended with the Environment and number of alerts. For example, the following illustrates that the **TBE** Service currently has no (**0**) associated alerts in the **PRODUCTION** Environment.

▼ ✓ TIBCO-AS
 ✓ TAS-MEMBER (PRODUCTION)

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

Click a column heading to sort the table on that column data.

Optionally, you can use the **\$rtvUserShowDualTables** substitution to add a table that lists alerts owned by the logged in user.

The screenshot shows the 'Alerts Table' window. At the top, there's a 'Current' dropdown and 'Alert Group: All'. Below that are search and filter fields. The table itself has columns: First Occ, Last Occ, Count, Sup, Owner, Alert Name, Primary Service, and CI. The rows are color-coded: red for 'High Warning Limit exceeded', yellow for 'High Alert Limit exceeded', and green for 'High Alert Limit exceeded'. The status bar at the bottom shows 'Data OK' and '23-Mar-2017 12:04'.

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.

The row color indicates the following:




Row Color Code:

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
- Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** options are disabled.

Fields and Data

This display includes:

Field Filter	<p>Select a table column from the drop-down menu to perform a search in: Alert Name, Alert Text, Alert Class, Service, CI, Closed Reason, Closed, CompId, Count, First Occ, ID, Last Occ, Owner, Primary Service, Sup, TicketGroup, TicketID.</p> <p>Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table).</p>
Clear	Clears the Field Filter and Search Text entries.
Search Text	Enter the (case-sensitive) string to search for in the selected Field Filter .
CMDB Filter	<p>Shows the selected Owner, Area, Group, Service and Environment filters. By default, all components of the CMDB (*) are included in the search.</p> <p>These CMDB Filter fields are populated when you click Open Alerts Table , which is accessible from the Multi Area Service Views displays, to open the Alerts Table in a new window. The filters selected in the All Management Areas and Multi Area Service Views displays are applied to the Alerts Table (that opens in the new window). NOTE: When you use the navigation tree (in the left panel) to open the Alerts Table display, the Environment filter is applied to the display if it has a value other than * (asterisk).</p>
Clear CMDB Filter	Clears all of the values in the CMDB Filter (Owner, Area, Group, Service and Environment filters). NOTE: This action is not applied to any other display.
RegEx	Toggles the Search Text field to accept Regular Expressions for filtering.
All	Click to show all alerts in the table: Open and Closed alerts.
Open	Click to only show Open alerts in the table.
Closed	Click to only show Closed alerts in the table.
Owner Filter	Select the alert Owner to show alerts for in the table.
	<p>All Shows alerts for all Owners in the table: Not Owned and Owned By Me alerts.</p> <p>Not Owned Shows only alerts without Owners in the table.</p> <p>Owned By Me Shows only alerts for the current user in the table.</p>
Alert Settings Conn OK	<p>The Alert Server connection state:</p> <p> Disconnected.</p> <p> Connected.</p>
Total	X/Y where X is the total number of alerts in the table with all selected filters applied. Y is the number of alerts in the table with only the CMDB and Cleared filters applied.
Critical	<p>Check to show alerts in the table that are currently in a critical state. NOTE: You must check Critical to see alerts that are in a critical state.</p> <p>X/Y where X is the total number of critical alerts in the table with all selected filters applied. Y is the number of alerts in the table with only the CMDB Filter and Cleared filters applied.</p>

Warning	<p>Check to show alerts in the table that are currently in a warning state. NOTE: You must check Warning to see alerts that are in a warning state.</p> <p>X/Y where X is the total number of warning alerts in the table with all selected filters applied. Y is the number of alerts in the table with only the CMDB and Cleared filters applied.</p>
Suppressed	<p>Check to show alerts in the table that are suppressed. The Suppressed count is not impacted by the Critical and Warning filters. It is impacted only by the CMDB Filter and the Owner Filter. NOTE: You must check Suppressed to see Suppressed alerts in the table.</p>
Own	<p>Click to assign an Owner for the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row. For details, see Configure User and Role Management.</p>
Suppress	<p>Click to suppress the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row. For details, see Configure User and Role Management.</p>
UnSuppress	<p>Click to unsuppress the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row or when you select a row. For details, see Configure User and Role Management.</p>
Close	<p>Click to close the alert. This option is only visible to users with Administrator privileges. This option is disabled when you select a gray row or you select a row where the Primary Service is not in the \$rtvManageableCompID list for the logged in user. For details, see Configure User and Role Management.</p>

Alerts Table

This table lists all active alerts for the current filters. The table is empty unless you check **Critical**, **Warning**, or both. Filter the list using the search fields and drop-down menus (in the upper portion of the display). To view details about an alert, select an alert and click **Details** (in the bottom right portion of the display) to open the **Alert Detail** dialog. To view details about the CI source of the alert, select an alert and click **Go To CI** (in the bottom right portion of the display) to open its Summary display.

	First Occ	The date and time the alert first occurred.
	Last Occ	The date and time the alert last occurred.
	Count	The number of times the alert was generated.
	Sup	When checked, the alert has been suppressed by a user.
	Owner	The named owner assigned by the administrator.
	Alert Name	The name of the alert.
	Primary Service	The name of the Service with which the alert is associated.
	CI	The CI alert source.
	Alert Text	Description of the alert.
	AlertClass	An optional alert field which can be used when integrating with other alerting systems.
	CompID	An optional alert field which can be used when integrating with other alerting systems.
	TicketID	An optional alert field which can be used when integrating with other alerting systems.
	TicketGroup	An optional alert field which can be used when integrating with other alerting systems.
Columns	Id	When checked, shows the ID column in the table.
	Closed	When checked, shows the Closed column in the table.
	Closed Reason	When checked, shows the Closed Reason column in the table.
	Alert Index	When checked, shows the Alert Index column in the table.
Go To CI	Select an alert from the Alerts Table , then click Go To CI to view details for the selected CI in the Summary display.	
Annotate	Select one or more alerts from the Alerts Table , then click Annotate to open the Set Owner and Comments dialog and enter comments or change alert owner. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row or when you select a row where the Primary Service is not in the \$rtvManageableCompID list for the logged in user. For details, see Configure User and Role Management .	
	ID	Lists the alert IDs, separated by semicolons, for the alerts selected from the Alert Table .
	Source	Lists the name of the back-end Data Server reporting the alert, separated by semicolons.
	Enter Owner	Enter the name of the owner for one or more alerts, click Set Owner of One Alert to assign the Owner, then click Close . By default, this field displays the current user name.
	Enter Comment	Enter a comment for one or more alerts, click Add Comment on One Alert to apply the Comment, then click Close . By default, this field displays previously entered comments. The text appears in the Comments field for the alert.


Set Owner	Applies the name of the alert owner in the Enter Owner field for one or more alerts.
Add Comment	Applies the comment in the Enter Comment field for one or more alerts.
Clear Comments	Removes all comments for one or more alerts.
Close	Closes the dialog.
Options	Select a single alert from the Alerts Table , then click Options to open the Alert Options dialog. This dialog is provided for customizing your own alert options. This option is disabled when you select a gray row or more than one row.
Details	Select a single alert from the Alerts Table , then click Details to open the Alert Detail window and view alert details. This option is disabled when you select a gray row or more than one row.

Alert History Table

Use this display to track the history of any alert that has occurred in your RTView Enterprise Monitor system. There is one row in the table for each update to each alert. The table is limited to **20,000** rows. If there are more than **20,000** rows in the selected time range, the newest **20,000** rows are shown.

The color coded navigation tree shows the contents of the CMDB hierarchically ordered. Choose a node to filter alerts shown in the table. The **Alert History Table** only shows alerts associated with the node you select. A green indicator means the node has no associated alerts. A red indicator means the node has one or more associated alerts.

Service name labels are appended with the Environment. For example, the following illustrates that the **TAS-MEMBER** Service currently has no alerts in the **PRODUCTION** Environment.



To filter the table, select a table column from the **Field Filter** drop-down menu. In the **Search Text** field, enter the (case-sensitive) string to search for in the selected **Field Filter**, then click **<Enter>**. Click **Clear** to clear the **Field Filter** and **Search Text** fields.

The **Count** label shows two values: the filtered row count / the total row count.

Click a column heading to sort the table by the column data.

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.



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.

Note: The **Count** field in the title bar of this display shows two values: the filtered row count and the unfiltered row count.

The row color indicates the most critical alert state for the row, as follows:

Row Color Code:

Tables with colored rows indicate the following:

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

Fields and Data

This display includes:


Field Filter

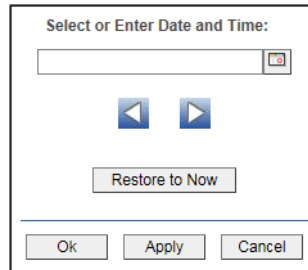
Select a table column from the drop-down menu to perform a search in: **Alert Name, Alert Text, Cleared Reason, Ctr, ID, Owner, Sev, Source, Sup, ID or Time.**

Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table).


Clear



Clears entries in the **Alert Name Filter** field and all table data.

- Search Text** Enter the (case-sensitive) string to search for in the selected **Field Filter**.
- RegEx** Toggles the **Search Text** field to accept Regular Expressions for filtering.
- Sort by ID + Time** When checked, table rows are sorted by the **Time** and **ID** columns.
- 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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

Alerts Table

This table lists all alerts for all Owners and all Areas that have occurred in your RTView Enterprise Monitor system. Filter the list by alert names using the **Alert Name Filter** drop-down menu.

Time	The date and time the alert first occurred.
ID	The unique string identifier for the alert.
Clear	When checked, the alert has been cleared by a user.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Alert Index	Lists the Alert Indexes, separated by tildes (~), for the alert.
Alert Text	Descriptive text about the alert.
Cleared Reason	DATA UPDATE: The metric returned to normal thresholds. MANUAL: A user cleared or closed the alert manually.
Sev	The severity level of the alert.
Source	The name of the back-end Data Server reporting the alert.

Administration

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

- **"Alert Administration"**: Displays active alerts and provides interface to modify and manage alerts.
- **"Alert Admin Audit"**: Track modifications of alerts throughout your system, such as alert threshold modifications.
- **"Alert Action Audit Trail"**: Track alert management throughout your system, including the name of the user who performed the action, the time the action was performed and what the action was.

Alert Administration

Set global or override alert thresholds. Alert settings are global by default. Only users logged in with the admin or super roles can save changes to alert thresholds. For details, see

Configure User and Role Management.

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

Global Thresholds

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

The manner in which global alerts are applied depends on the Solution Package. For example, the EMS Monitor Solution Package has queue alerts, topic alerts and server alerts. When a queue alert is applied globally, it is applied to all queues on all servers. Likewise, a server alert applies to all servers, and a topic alert applies to all topics on all servers.

Override Thresholds

Setting override alerts allows you to set thresholds for a single resource (for example, a single server). Override alerts are useful if the majority of your alerts require the same threshold setting, but there are other alerts that require a different threshold setting. For example, you might not usually be concerned with execution time at a process level, but perhaps certain processes are critical. In this case, you can apply alert thresholds to each process individually.

To apply an individual alert you Index the Monitored Instance or resource (such as a message queue, in the case of the EMS Monitor package). The Index Types available are determined by the Solution Package installed. For example, with the EMS Monitor package you can set an alert for a specific topic on a specific server--the PerServerTopic Index option--rather than for all topics on all servers.

For information about setting override alerts, see [“Tabular Alert Administration”](#).

Note: To filter the alerts shown in the **Administration - Alert Administration** display by solution package, use the `$rtvAlertPackageMask` substitution. For details, refer to the *RTView Enterprise Monitor Configuration Guide*.

Alert Administration 23-Sep-2015 16:15 Data OK

Alert Filter: Clear

Alert	Warning Level	Alarm Level	Duration	Alert Enabled	Override Count
AcwlInstanceCpuHigh	50	75	30	<input checked="" type="checkbox"/>	
AcwlInstanceDiskReadBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AcwlInstanceDiskReadOpsHigh	100	200	30	<input checked="" type="checkbox"/>	
AcwlInstanceDiskWriteBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AcwlInstanceDiskWriteOpsHigh	100	200	30	<input checked="" type="checkbox"/>	
AcwlInstanceNetworkReadBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AcwlInstanceNetworkWriteBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AmxServiceHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeFaultRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeMovingAvgHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeMovingAvgResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	
Bw6AppNodeCpuUsedHigh	50	80	30	<input type="checkbox"/>	
Bw6AppNodeMemUsedHigh	50	80	30	<input type="checkbox"/>	
Bw6AppProcessCreatedRateHigh	50	80	30	<input checked="" type="checkbox"/>	
Bw6AppProcessElapsedTimeHigh	100	200	30	<input type="checkbox"/>	

Settings for Selected Alert

Name: Warning Level: Duration (Secs.):

Description: Alarm Level: Enabled: ☒

Save Settings

Tabular Alert Options

The Warning Level, Alert Level and Alarm Enabled values on this screen can be overridden for each alert

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.
- 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.

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.


Fields and Data


This display includes:

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

**Alert
Setting
s**

The Alert Server connection state:

 Disconnected.

 Connected.

Active Alert Table


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

Alert	The name of the alert.
Warning Level	The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.
Alarm Level	The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.
Duration (Secs)	The amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution.
Alert Enabled	When checked, the alert is enabled globally.
Override Count	The number of times thresholds for this alert have been defined individually in the Tabular Alert Administration display. A value of: -0 indicates that no overrides are applied to the alert. -1 indicates that the alert does not support overrides.

Settings for Selected Alert

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

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

Name	The name of the alert selected in the Active Alert Table .
Description	Description of the selected alert. Click Calendar  for more detail.
Warning Level	Set the Global warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value. NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the warning to occur sooner, increase the Warning Level value. To set the warning to occur later, reduce the Warning Level value.
Alarm Level	Set the Global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value.
Duration	Set the amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution. This setting is global.
Enabled	Check to enable alert globally.

Save Settings

Click to apply alert settings.

Override SettingsClick to open the **Tabular Alert Administration** display to set override alerts on the selected alert.**Tabular Alert Administration**

Set override alerts (override global alert settings). This display opens when you select an alert in the **Alert Administration** display and then select **Override Settings**.

For step-by-step instructions setting thresholds for individual alerts, see **Setting Override Alerts**.

Fields and Data

This display includes:

Alert Settings Conn OK

The connection state.



No servers are found.



One or more servers are delivering data.

Override Settings For Alert:(name)

This table lists and describes alerts that have override settings for the selected alert. Select a row to edit alert thresholds. The selected item appears in the Index field. Edit settings in the Alert Settings fields, then click Save Settings.

Index Type	<p>Select the type of alert index to show in the Values table. Options in this drop-down menu are populated by the type of alert selected, which are determined by the Package installed. For example, with the EMS Monitor package the following Index Types are available:</p> <ul style="list-style-type: none"> • PerServer: Alert settings are applied to a specific server. • PerQueue: Alert settings are applied to the queue on each server that has the queue defined. • PerServerQueue: Alert settings are applied to a single queue on a specific server. • PerTopic: Alert settings are applied to the topic on each server that has the topic defined. • PerServerTopic: Alert settings are applied to a single topic on a specific server.
Index	The value of the index column.
Override Settings	When checked, the override settings are applied.
Alert Enabled	When checked, the alert is enabled.
Index Type	Select the index type. The index type specifies how to apply alert settings. For example, to a queue (topic or JVM, and so forth) across all servers, or to a queue on a single server. NOTE: Options in this drop-down menu are populated by the type of alert selected from the Alert Administration display. Index Types available depend on the Package installed.
Index	The selected index column to be edited. This field is populated by the selection made in the Unassigned Indexes table.
Unassigned Indexes	This table lists all possible indexes corresponding to the Index Type chosen in the drop-down list. Select a row to apply individual alert thresholds. The selected item appears in the Index field. Edit settings in the Alert Settings fields, then click Add .
Add	Click to add changes made in Alert Settings , then click OK to confirm.
Remove	Click to remove an alert selected in the Index Alert Settings table, then click OK to confirm.
Save Settings	Click to save changes made to alert settings.
Alert Settings	Select a topic, server or queue from the Unassigned Indexes table and edit the following settings.
Warning Level	<p>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.</p> <p>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.</p> <p>Click Save Settings to save settings.</p>

Alarm Level	Set the alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value. Click Save Settings to save settings.
Alert Enabled	Check to enable the alert, then click Save Settings .
Override Settings	Check to enable override global setting, then click Save Settings .

Back to Alerts Returns to the **Administration - Alert Administration** display.

Setting Override Alerts

Perform the following steps to set an override alert. Index Types available depend on the Solution Package installed. In this example, we use the EMS Monitor Package to illustrate.


Note: To turn on an alert, both **Alert Enabled** and **Levels Enabled** must be selected.

To turn on/off, change threshold settings, enable/disable or remove an alert on a single resource:

1. In the **Alert Administration** display, select a tabular alert in the **Active Alert Table** and click **Override Settings**. The **Tabular Alert Administration** display opens.

Note: Alerts that do not support overrides have a value of **-1** for the **Override Count** column and the **Override Settings** option is not present when you select such an alert.

2. In the **Tabular Alert Administration** display, select the Index type from the **Index Type** drop-down menu (options are populated by the type of alert you previously selected). For example, with the EMS Monitor package, select PerServerQueue, PerServerTopic or PerServer. NOTE: If you select PerServerQueue or PerServerTopic, the alert settings are applied to the queue or topic on a single server.
3. In the **Unassigned Indexes** table, select the item you want to apply an override alert setting to, click **Add** and **OK** in the confirmation dialog. After a few moments the override setting appears in the **AlertLevels** table.
4. Select the item in the **AlertLevels** table.
5. In the Alert Settings panel (lower right), if needed, modify the Warning Level and Alarm Level settings.
6. In the **Alert Settings** panel, set the following as appropriate.
 - To turn on the alert for this index with the given thresholds:
Alert Enabled Select this option.



Fields and Data

This display includes:

Audit Conn OK	<p>The Alert Server connection state.</p> <ul style="list-style-type: none"> ● Disconnected. ● Connected.
TIME_STAMP	The date and time of the modification.
USER	The user name of the administrator who made the modification.
ACTION	The type of modification made to the alert, such as UPDATED .
ALERTNAME	The name of the alert modified.
INDEXTYPE	<p>The type of alert Index.</p> <p>Index Type refers to the manner in which alert settings are applied and vary among CI Types. For example, the JVM CI Type has a PerJvm Index Type, the EMS CI Type has PerServer, PerTopic and PerQueue Index Types which apply alerts to servers, topics and queues, respectively.</p>
ALERTINDEX	The index of the alert which identifies its source.
WARNINGLEVEL	<p>The warning threshold value for the alert at the time this modification was made, as indicated in the TIME_STAMP column.</p> <p>The warning level is a threshold that, when exceeded, a warning is executed.</p>
ALARMLEVEL	<p>The alarm threshold value for the alert at the time this modification was made, as indicated in the TIME_STAMP column.</p> <p>The alarm level is a threshold that, when exceeded, an alarm is executed.</p>
DURATION	<p>The duration value for the alert at the time this modification was made, as indicated in the TIME_STAMP column.</p> <p>The alert duration is the amount of time (in seconds) that a value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution.</p>
ENABLED	When checked, indicates the alert was enabled at the time this modification was made, as indicated in the TIME_STAMP column.
USEINDEX	When checked, indicates the alert override was enabled at the time this modification was made, as indicated in the TIME_STAMP column. For details about alert overrides, see Alert Administration .

Alert Action Audit Trail

The **Alert Action Audit Trail** display shows all user actions concerning alert management, including the name of the user who performed the action, the time the action was performed and what the action was. This display can help managers of the RTView Enterprise Monitor solution determine how and when user interactions have impacted the alert system and help manage users so that best practices for alert handling are maintained.

TIME_STAMP	USER	ACTION_TYPE	ACTION	TARGET	VALUE	ALERT_NAME
10/01/15 16:56:29	admin	Event Management	Set Owner	2764	admin	EmsServerRouteState
10/01/15 16:56:29	admin	Event Management	Set Owner	2562	admin	EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Set Owner	2385	admin	EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Set Owner	2339	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2304	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2256	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2096	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2039	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	2004	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1956	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1796	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1761	admin	EmsServerAsyncDBSize
10/01/15 16:56:29	admin	Event Management	Set Owner	1732	admin	EmsQueuesProducerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1375	admin	EmsQueuesProducerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1358	admin	EmsQueuesConsumerC
10/01/15 16:56:29	admin	Event Management	Set Owner	1001	admin	EmsQueuesConsumerC
10/01/15 16:56:29	admin	Event Management	Clear Alert	2764		EmsServerRouteState
10/01/15 16:56:29	admin	Event Management	Clear Alert	2562		EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Clear Alert	2385		EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Clear Alert	2339		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2304		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2256		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2096		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2039		EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Clear Alert	2004		EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Clear Alert	1956		EmsTopicsConsumerCo

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.

Fields and Data

This display includes:

Action Audit Conn OK

The Alert Action database connection availability.

The connection to the Alert Action database is not available.

The connection to the Alert Action database is available.

Time_Stamp

The time the action was performed.

User

The log in name of the user who performed the action.

Action_Type

The type of action which was performed.

Action

The action which was performed.

Target

The alert ID on which the action was performed.

Value

Any value associated with the alert action.

Alert_Name

The name of the alert on which the action was performed.

Alert_Index

The index of the alert which identifies its source.

CMDB Administration

This display allows you to modify your Service Data Model.

- **"CMDB Admin"**: View or modify your Service Data Model.

CMDB Admin

Use this display to setup, view or modify your Service Data Model (CMDB), including: adding, renaming, deleting or merging your CMDB hierarchical elements (Owners, Areas, Groups or Services), associating CIs with Services and assigning or modifying CI attributes (such as Criticality). Only users logged in with the admin or super roles can apply changes in this display.

This display requires administrator permissions.

The **CI List for Selected Service** (upper) table lists the CIs that are associated with the Service selected (from the **Service** drop-down menu).

The **Available Components** (lower) table is not part of the CMDB. The **Available Components** table lists all available CIs for the CI Type (selected from the **Selected CI Type** drop-down menu) that are in your RTView Enterprise Monitor system--whether or not they are in the CMDB. Filter this list using the **CIName Filter** field.

You add CIs to the CMDB by associating them with an Owner, Area, Group, and Service. To do so, select the CI Type from the **Selected CI Type** drop-down menu, choose one or more CIs from the **Available Components** table, then click **Add CI**.

It is not necessary to restart the Configuration Server after making changes to the Service Data Model using the **CMDB Admin** display.

Creating a new Service, Group, Area or Owner:

Select the CI Type from the **Selected CI Type** drop-down menu, choose one or more CIs from the **Available Components** table, then click **Add CI To...** Assign a new or existing Owner, Area, Group or Service, review your entries and click **OK**. Your changes are visible in drop-down menus and displays.

Associating CIs with a Service:

This option is useful when you want to define which CIs are to be monitored for Services. CIs can be associated with more than one Service, Group, Area or Owner. Select the Owner, Area, Group and Service to which you want to associate one or more CIs using the drop-down menus. The **CI List Table** (the upper table) populates with all CIs already associated with the Owner, Area, Group and Service you select. Select the CI Type of the CI(s) you want to associate. The **Available Components** table (the lower table) populates with all CIs that are categorized as that CI Type. Select one or more CIs in the **Available Components** table, set the Criticality and other optional assignments using the drop-down menus (on the right). Click **Add CI** to associate the CI(s) with the Service. A row is added for each associated CI to the **CI List Table**. Your changes are visible in the drop-down menus and displays.

Renaming a Service, Group, Area or Owner:

This option is useful when, for example, a *new* Owner is replacing a retiring Owner, a name is misspelled or a more relevant name is required. Select the relevant Owner, Area, Group or Service using the drop-down menus, then click the corresponding **Manage (Owner, Area, Group or Service)** option for what you are renaming. The **Manage (Owner, Area, Group or Service)** dialog opens. In the **Manage (Owner, Area, Group or Service)** dialog, type the new name in the **New Name** field, click **Rename** and **OK**. Your changes are visible in the drop-down menus and displays.

Deleting a Service, Group, Area or Owner:

This option is useful when, for example, an Owner, Area, Group or Service and all the CIs associated with it are not relevant in your RTView Enterprise Monitor system. When you delete a Service, Group, Area or Owner everything underneath it (lower CMDB levels and associated CIs) is also removed from the CMDB database and displays. Select the relevant Owner, Area, Group or Service using the drop-down menus, then click the corresponding **Manage (Owner, Area, Group or Service)** option for what you are deleting. The **Manage (Owner, Area, Group or Service)** dialog opens. In the **Manage (Owner, Area, Group or Service)** dialog click **Delete** and **OK**. Your changes are visible in the drop-down menus and displays.

Important: There is no option to undo a deletion from the CMDB. To restore a deletion you must recreate the Owner, Area, Group or Service and the CIs must be re-associated.

Moving a Service, Group or Area:

This option is useful when, for example, an Area belongs under a different Owner, a Group belongs under a different Area or a Service belongs under a different Group. When you move a Service, Group or Area (Owners cannot be moved) everything underneath it (lower CMDB levels and associated CIs) moves with it. Select the Area, Group or Service you want to move using the drop-down menus, then click the relevant **Manage (Area, Group or Service)** option for what you are moving. The **Manage (Area, Group or Service)** dialog opens. In the **Manage (Area, Group or Service)** dialog, select the new Owner, Area, Group or Service to move to from the **New (Area, Group or Service)** drop-down menus, click **Move** and **OK**. Your changes are visible in the drop-down menus and displays.

Merging Services, Groups, Areas or Owners:

This option is useful when, for example, an *existing* Owner is taking over for a retiring Owner. When you merge a Service, Group, Area or Owner its name changes to that of the target Service, Group, Area or Owner, and everything underneath it (lower CMDB levels and associated CIs) goes with it. Select the Area, Group or Service you want to merge using the drop-down menus, then click the relevant **Manage (Area, Group or Service)** option for what you are merging. The **Manage (Area, Group or Service)** dialog opens. In the **Manage (Area, Group or Service)** dialog, select an existing Owner, Area, Group or Service to merge to in the **New Name** field, click **Merge** and **OK**. Your changes are visible in the drop-down menus and displays.

Deleting a CI:

Select a CI from the **CI List Table**, click **Delete** and **OK**. The CI is removed from the CMDB database and displays. Your changes are visible in the drop-down menus and displays.

Applying Criticality value to multiple CIs:

In the **CI List Table** select a CI that has the Criticality value you want to apply to all CIs in the **CI List Table**, click **Update Criticality like selected CI** and **OK**. The **Criticality** column for all CIs is updated. Your changes are visible in the drop-down menus and displays.

Changing CI attributes

In the **CI List Table** select the CI you want to modify attributes for, use the **Environment**, **Region**, **SiteName**, **Criticality**, **City**, **Country** and **OSType** drop-down menus to apply attributes, then click **Update** and **OK**. The **CI List Table** is updated. Your changes are visible in the drop-down menus and displays.

By default, the Owner named **Infrastructure** is created. **Infrastructure** organizes all available CIs collected through all Data Servers configured under RTView EM by technology. This default organization can be disabled if needed.

CMDB - Administration 24-Sep-2015 11:45 Data OK

Owner: Jerelyn Parker Manage Owner Source: RTV_CMDB

Area: Systems Manage Area

Group: Databases Manage Group

Service: IBM DB2 Manage Service

Update Criticality like selected CI

Environ: DEMOSITE

Region: Region: Criticality: A SiteName: Headquarters City: Country: OSType:

Update Delete

Selected CI Type: MQ-QUEUE CName Filter: Regex Add CI Add CI To...

CI List for Selected Service - select a CI to see detail and to edit:

CIType	CIName	Criticality	Region	Env
VMWARE-HOST	vSphere2:slesxi-1.sldemos-hq.local	B		QA
VMWARE-VM	vSphere2:2008S-WIN14	B		QA

Available Components (CIs):

conn	Name	CIName	Data Servi
vmrh5-1	TEST_Q_01	vmrh5-1;TEST_Q_01	MQMON-64-OL7-3
vmrh5-1	TEST_Q_02	vmrh5-1;TEST_Q_02	MQMON-64-OL7-3
vmrh5-1	TEST_Q_03	vmrh5-1;TEST_Q_03	MQMON-64-OL7-3
vmrh5-1	TEST_Q_04	vmrh5-1;TEST_Q_04	MQMON-64-OL7-3
vmrh5-1	TEST_Q_05	vmrh5-1;TEST_Q_05	MQMON-64-OL7-3

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.

Fields and Data

This display includes:

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

	<p>Manage Area Opens a dialog that enables you to Delete, Rename or Merge the Area.</p> <p>Delete removes the Area from the CMDB database as well as all CMDB data and CIs associated with the Area.</p> <p>Rename Changes all records for the Area to a new name. Rename is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Merge Changes all records for the Area to a different, already existing name in the CMDB. Merge is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Move Changes all records for the Area to a different, already existing name in the CMDB that you choose from the New Area drop-down menu.</p>
Group	Select a Group to filter by. The Group selected populates the Service drop-down menu.
	<p>Manage Group Opens a dialog that enables you to Delete, Rename or Merge the Group.</p> <p>Delete removes the Group from the CMDB database as well as all CMDB data and CIs associated with the Group.</p> <p>Rename Changes all records for the Group to a new name. Rename is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Merge Changes all records for the Group to a different, already existing name in the CMDB. Merge is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Move Changes all records for the Group to a different, already existing name in the CMDB that you choose from the New Group drop-down menu.</p>
Service	Select a Service to edit, then click Update .
	<p>Manage Service Opens a dialog that enables you to Delete, Rename or Merge the Service.</p> <p>Delete removes the Service from the CMDB database as well as all CMDB data and CIs associated with the Service.</p> <p>Rename Changes all records for the Service to a new name. Rename is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Merge Changes all records for the Service to a different, already existing name in the CMDB. Merge is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Move Changes all records for the Service to a different, already existing name in the CMDB that you choose from the New Service drop-down menu.</p>

CI List Table

This table lists all CIs associated with the selected Service. Each table row is a different CI. Select a CI to see its attributes in the drop-down menus at the right of the table. Use the **OSType**, **Region**, **SiteName**, **Criticality**, **City** and **Country** drop-down assign attributes, then click **Update**. To associate CIs with the Service, select one or more CIs from the **Available Components** table, then click **Add CI** (to associate the CI(s) with the selected Service.) or **Add CI To...** (to create a new Service and associate the CI(s) with it).

CIType	The type of CI. For example, server or application.
CIName	A unique Identifier for the CI.
Criticality	<p>The importance level of the CI in your organization. Values range from A to E, where A is the highest Criticality and E is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate the Alert Impact (maximum Alert Severity multiplied by the maximum Criticality equals Alert Impact).</p> <p>Criticality values are listed in the Component Views - CI Service Table display. Criticality values are also shown in heatmaps and tables.</p>
Region	The name of the Region for the CI.
Environment	The name of the Environment for the CI.
SiteName	The name of the Site for the CI.
OSType	The operating system on the CI.
City	The name of the City for the CI.
Country	The name of the Country for the CI.
Update Criticality like selected CI	Updates the Criticality attribute assigned to all CIs in the CI List table to match the selected CI level.
Environ	Select or type the Environment for the CI selected in the CI List Table , or the CI selected in the Available Components and added into the CI List Table .
Region	Select or type the region for the CI selected in the CI List Table , or the CI selected in the Available Components and added into the CI List Table .
SiteName	Select or type the site name for the CI selected in the CI List Table , or the CI selected in the Available Components and added into the CI List Table .
Criticality	<p>Specify the importance level of a Service or a CI for your organization. Select a Service or a CI and set the Criticality value from A to E, where A is the highest Criticality and E is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate Alert Impact (maximum Alert Severity multiplied by the maximum Criticality equals Alert Impact).</p> <p>Criticality values are listed in the Component Views - CI Service Table display. Criticality values are also shown in heatmaps and tables.</p>
Country	Select or type the country for the CI selected in the CI List Table , or the CI selected in the Available Components and added into the CI List Table .
OSType	Select or type the operating system for the CI selected in the CI List Table , or the CI selected in the Available Components and added into the CI List Table .
Update	Updates the CI selected in the CI List Table with attributes selected from the drop-down menus (on the right).
Delete	Removes the selected CI from the CMDB database.

Available Components Table

This table lists all available CIs in your RTView Enterprise Monitor system whether they are in the CMDB or not. Each row in the table is a different CI (for example, a server or a process). Select one or more CIs to associate with the currently selected Service, then click **Add CI** (to associate the CI(s) with the selected Service.) or **Add CI To...** (to create a new Service and associate the CI(s) with it).

Selected CI Type	Select the type of CI to include in the Available Components table. All CIs of this type are listed. A CI can be associated with multiple Services.
-------------------------	--

CIName Filter	Enter a string to filter the list of available components.
Regex	Check to enable Regex filtering.
Add CI	<p>Associates the CI selected in the Available Components table with the selected Service, and applies the attributes selected from the drop-down menus (on the right) to the CI.</p> <p>To associate a CI with the currently selected Service, select a CI from the Available Components table, use the drop-down menus on the right (Environ, Region, SiteName, etc.) to modify attributes for the CI, click Add CI and then click Update. The CI appears in the CI List Table.</p>
Add CI To...	<p>Creates a new Service and associates the selected CI with it.</p> <p>To create a new Service and associate a CI with it, select a CI from the Available Components table, use the drop-down menus on the right (Environ, Region, SiteName, etc.) to modify attributes for the CI, click Add CI To..., enter the name of the new Service, then click Update. The new Service is added to the list of Services and the CI appears in the CI List Table.</p>

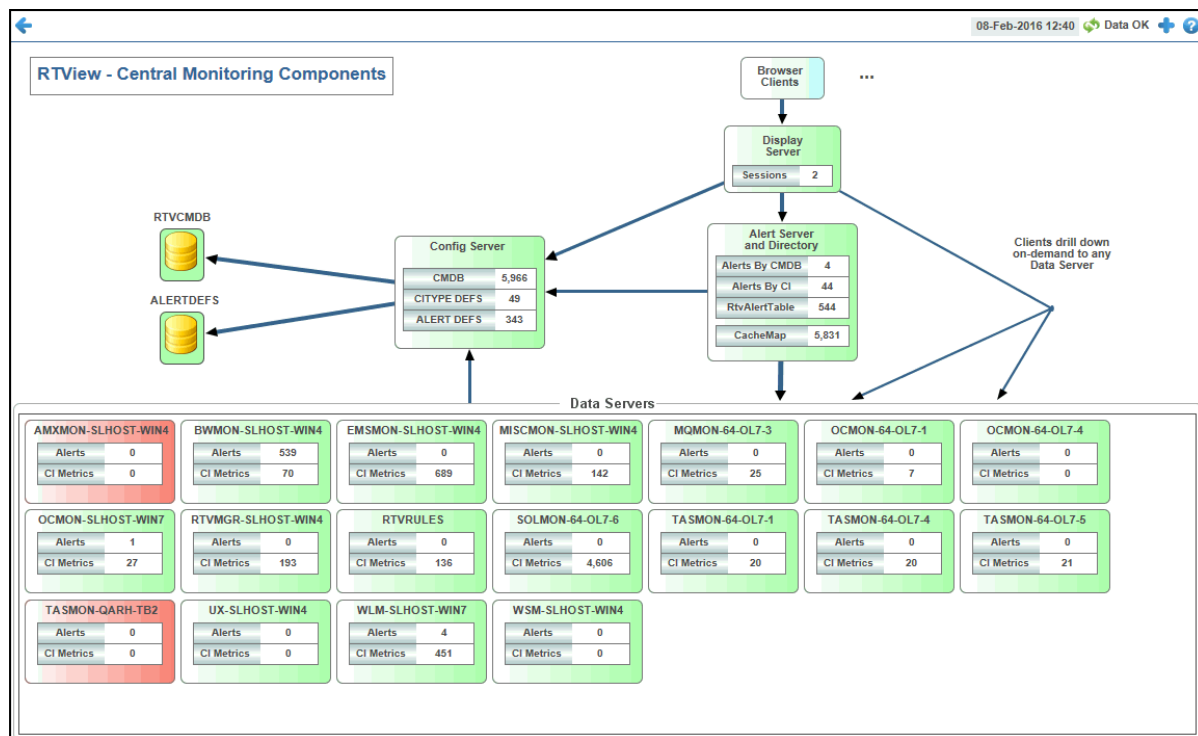
Architecture

These displays provide a view of RTView Enterprise Monitor component connectivity, mapping between component types, and component level connection and performance information. The Architecture displays are provided with RTView Enterprise Monitor. Displays in this View are:

- ["System Overview"](#) Topology map of the main RTView Enterprise Monitor components. Objects are color-coded to show component status.
- ["RTView Data Servers"](#): Configuration and connection details for RTView Data Servers.
- ["Data Server Summary"](#): Connection and query statistics for RTView Data Servers.
- ["RTView History Table Statistics"](#): Performance of historical data being stored from caches with history.
- ["RTView Cache Tables"](#): Configuration and alert details for RTView Cache Tables.
- ["RTView CI Stats Tables"](#): Alert details for RTView Cache Tables by CI.
- ["RTView CI Type Defs"](#): CI Type definitions, cache map and alert map by CI Type.
- ["RTView KM Defs"](#): Key Metrics definitions for all CI Types.
- ["About"](#): This display shows details about the RTView Enterprise Monitor version and data sources available to your system.

System Overview

View the topology of the central RTView Enterprise Monitor monitoring components and their current connection state. Each object represents a component which are color-coded to indicate component status. Red indicates the component stopped running. Green indicates the component is running.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

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.

Fields and Data

This display includes:


Config Server

The Configuration Server provides configurations to all central RTView Enterprise Monitor components.

- CMDB** The number of CIs in the CMDB.
- CITYPE DEFS** The current number of CITYPE definitions.
- ALERTDEFS** The current number of alert settings and override definitions.

Alert Server and Directory

The Alert and Directory Server centralizes access to all alerts sent by remote Data Servers, and maintains a directory table of CI locations. The CI location is the name of the source Data Server.



Fields and Data

This display includes:

Local Connections to Data Server

This table lists all Data Servers and detailed connection information. Select a Data Server to view further details (in the lower tables).

Name	The Data Server name.
Connected	When checked, the connection is currently connected.
Status	The Data Server connection status.
Connection String	The host name and port number for TCP connections, or the URL for servlet connections.
Rcv Cnt	The number of data updates received from that Data Server.
ReceiveTime	The time that data was last received.
Config	The RTView version running on the Data Server.

Data Server Manager

This table shows connection information for the Data Server selected from the **Local Connections to Data Server** table.

NumberOf Clients	The number of clients currently connected to the Data Server.
ServingData	When checked, the Data Server is currently serving data.
Connection Request Count	The number of client requests to connect to the Data Server.
Connection Request FailedCount	The number of client requests to connect to the Data Server that were unable to connect.

Data Server Clients

This table shows information for clients connected to the Data Server selected from the **Local Connections to Data Server** table.

ClientID	A unique string identifier for the client.
Address	The client IP address.
Duration	The client session length of time.
Host	The address of the client host.
Last Data Sent	The amount of data, in bytes, the Data Server last sent to the client.
Total Data Sent	The total amount of data, in bytes, the Data Server has sent to the client.

Data Server Summary

View Data Server connection status, cache table sizes and database query metrics. Use the available drop-down menus or right-click to filter data shown in the display.

RTView DataServer - Summary23-Sep-2015 14:20Data OK

Data Server: <Default>

Connection Status					
Connected	Status	Connection String	Receive Count	Receive Time	Config

Alert Table ViewAlert AdminHistory Tables

RTView Cache Tables	
CacheTable	Rows
RtvAlertTableLocal	19,909
RtvMxCacheDefsWithCo	1,429
RtvTabTreeCache	488
RtvMxCacheDefsRaw	234
RtvMxCacheDefs	169
RtvCmdbServiceTable	56
RtvMxCacheInfoByServ	56
RtvDataServerConnecti	20
RtvCmdbGroupTable_I	17
RtvCmdbAreaTable_loc	8
RtvCmdbOwnerTable_I	2
JmxStatsTotals	1
RtvAlertMapByCI	0
RtvAlertSourceStats	0
RtvAlertStatsByCI	0
RtvAlertStatsByCIAndA	0

Database Queries (running on selected Data Server)								
Database	Conn	Count	Active	ExecTime	Rows	RunTime	Status	
ALERTDEFS	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
PROPDB	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
RTVCMDB	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
RTVCONFIG	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
RTVHISTORY	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	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 OKData 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 12:04Current 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.

Fields and Data

This display includes:

Data Server Select a Data Server from the drop-down menu to view details for in the display.

Connection Status

This table shows connection details for the selected Data Server.

Connected	When checked, the Data Server is currently connected.
Status	The Data Server connection status.
Connection String	The host name and port number for TCP connections, or the URL for servlet connections.
Rcv Cnt	The number of data updates received from that Data Server.
ReceiveTime	The time that data was last received.
Config	The RTView version running on the Data Server.

Alert Table View Select to view or manage current alerts for the selected Data Server in the **RTView Alerts Table** display.

Alert Admin Select to view or manage alert thresholds for the selected Data Server in the **Alert Administration** display.

History Tables Select to view database table statistics for each cache for the selected Data Server in the **"RTView History Table Statistics"** display.

RTView Cache Tables

This table lists Cache Tables and their size, in number of rows, for the selected Data Server. Select a Cache Table to view details in the **RTView Cache Tables** display.

Use this data for debugging. This display is typically used for troubleshooting with SL Technical Support.

CacheTable	The name of the Cache Table.
Rows	The current number of rows in the Cache Table.

Database Queries

This table lists the databases and query details for the selected Data Server. Each table row describes a different query.

Database	The name of the database.
Conn	When checked, the database is currently connected.
Count	The number of query requests from current Data Server.
Active	When checked, the query is currently running.
ExecTime	The amount of time, in milliseconds, to execute the query.
Rows	The number of rows the query created.
RunTime	The time the query was executed.
Status	The latest result status of the query.
Query	The query that was executed.

RTView History Table Statistics

This display opens when you click **History Tables** from the **Architecture - "Data Server Summary"** display. View information about the performance of historical data being stored from caches with history. Use this display to verify your tables are growing as expected by:

- seeing how many rows are in the database table (**Row Count**).
- seeing how many rows are added at each update period (**Delta**).
- verifying that the range of the data stored in the table is consistent with defined compaction rules and that behavior is as expected. To do this, compare the time of **First Entry** and **Last Entry** and verify the dates match the defined compaction interval (for example, **2 weeks** by default). For this verification, you must first confirm the historian has been operating for at least the defined compaction time interval, otherwise the range of data will be shorter.

DataServer: EMSMOH-SLDEMOS		RTView History Table Statistics					10-Dec-2014 07:14	Data OK	+	?
Cache Name / DB Table Name	Row Count	Delta	Distinct	First Entry	Last Entry	Current				
EmsAdmStats EMS_ADMSTATS	8,337	1	0	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00					
EmsDurables EMS_DURABLES	395,397	53	53	18-Sep-2014 01:15:00	10-Dec-2014 07:15:00					
EmsQueueTotalsByServer EMS_QUEUETOTALS	181,236	181,236	24	18-Sep-2014 01:15:00	10-Dec-2014 05:25:00					
EmsQueues EMS_QUEUES	3,958,595	296	665	15-Sep-2014 02:00:00	10-Dec-2014 07:16:00					
EmsRouteCountsByServer EMS_ROUTECOUNTS	124,790	15	16	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00					
EmsRoutes EMS_ROUTES	208,014	25	26	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00					
EmsServerInfo EMS_SERVERINFO	199,807	24	25	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00					
EmsTopicTotalsByServer EMS_TOPICTOTALS	183,187	22	23	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00					
EmsTopics EMS_TOPICS	7,132,737	621	37,369	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00					
JvmMemory JVM_MEMORY	20,737	0	13	27-Nov-2014 00:15:00	10-Dec-2014 07:13:00					
JvmOperatingSystem JVM_OPERATINGSYSTEM	20,214	17	13	27-Nov-2014 00:15:00	10-Dec-2014 07:21:00					

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.

Fields and Data

This display includes:

- Cache Name / DB Table Name** The name of the cache and the name of the database table. Mouse-over to see the **Index** columns for the cache.
- Row Count** The number of rows in the table.
- Delta** The number of rows added since the last update.
- Distinct** The number of distinct indexes in the table.

Fields and Data

This display includes:

- Data Server** Select a Data Server from the drop-down menu to view details for in the display.
- Max Rows** Enter the maximum number of rows to include in the lower table, then click Enter.
- History Tables** Select to include all defined history tables in the **RTView Cache Tables** list.

RTView Cache Tables

This table lists cache tables for the selected Data Server. Select a cache table to view details in the lower table.

CacheTable	The name of the cache table.
TableType	The type of cache table.
current	This table is a current table which shows the current values for each index.
current_condensed	This table is a current table with primary compaction configured.
history	This table is a history table.
history_condensed	This table is a history table with primary compaction configured.
history_combo	This table is a history table with primary compaction configured, and which is also configured to store rows of recent raw data followed by rows of older condensed data.
Rows	The number of rows currently in the table.
Columns	The number of columns currently in the table.
Memory	The amount of space, in bytes, used by the table.

(Lower Table)

This table shows the contents of the selected cache table. Available columns vary by cache. For example, a JVM cache table might provide **BootClassPath** and **InputArgument** columns, and a Tomcat cache might provide **RateAccess** and **cacheMaxSize** columns.

Rows	The number of rows currently in the table.
-------------	--

RTView CI Stats Tables

View details for components that currently have an active warning or alarm alert.

RTView CI Stats Tables

25-Sep-2015 11:34 Data OK

Alert Stats By CI					
time_stamp	CITYPE	CINAME	MaxSeverity	AlertCount	
09/25/15 11:13:33	BW-ENGINE	slhpux11(simon);domainsimon.BWApp-7.Procs	1	1	
09/25/15 11:13:33	BW-ENGINE	slapm(slapm);domainslapm.BWApp-5.Procs	2	1	
09/25/15 11:13:33	BW-ENGINE	slvmrh2(slapm);domainslapm.BWEngine.Process Archive	2	1	
09/25/15 11:13:33	BW-ENGINE	slsl4-64(simon);domainsimon.BWApp-4.Procs	2	1	
09/25/15 11:13:33	BW-ENGINE	slapm(slapm);domainslapm.BW Engine.Process Archive	2	1	
09/25/15 11:13:33	BW-ENGINE	slapm(slapm);domainslapm.BWApp.Space.Procs-1	2	1	
09/25/15 11:13:33	BW-ENGINE	slsl4-64(simon);domainsimon.BWApp-5.Procs	1	1	
09/25/15 11:13:33	BW-ENGINE	slsl4-64(simon);domainsimon.BWApp-10.Procs	1	1	

Cache Map By CITYPE

time_stamp	CITYPE	CACHENAME	Source
09/25/15 09:33:53	ACW	AwsEc2InstanceStats	MISCMON-DATA-1
09/25/15 01:32:30	BW-ENGINE	BwEngines	Z-SIMDATA-1
09/25/15 06:02:09	BW-ENGINE	BwEngines	BWMON-SLDEMOS
09/25/15 01:32:30	BW-PROCESS	BwProcesses	Z-SIMDATA-1
09/25/15 06:02:09	BW-PROCESS	BwProcesses	BWMON-SLDEMOS
09/25/15 01:32:30	BW-SERVER	BwServers	Z-SIMDATA-1
09/25/15 06:02:09	BW-SERVER	BwServers	BWMON-SLDEMOS

Cache Map By CI

CI Type Filter: All CI Types

Count: 2475

time_stamp	CITYPE	CINAME	Data ServerName	Expired
------------	--------	--------	-----------------	---------

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.

Fields and Data

This display includes:

Alert Stats By CI

This table provides summary alert details for all CIs that currently have active warning or alarm alerts.

time_stamp	The date and time this table row of data was last updated. Format: MM/DD/YY HH:MM:SS <month>/ <day>/<year> <hours>:<minutes>:<seconds>
CITYPE	The component type.
CINAME	The name of the component.
MaxSeverity	The most critical alert state of all current alerts for this component.
AlertCount	The number of current warning and alarm alerts for this component.

Cache Map By CITYPE

This table provides mapping of all component types to caches.

time_stamp	The date and time this table row of data was last updated. Format: MM/DD/YY HH:MM:SS <month>/ <day>/<year> <hours>:<minutes>:<seconds>
CITYPE	The component type.
CACHENAME	The name of the cache associated with the component type.
Source	The name of the Data Server alert sending data for that component type.

Cache Map By CI

This table provides the location of all CIs.

CI Type Filter:	Select the CI Type to filter by in this table, or select All CI Types .
Count	The number of CIs currently in this table.
time_stamp	The date and time this table row of data was last updated. Format: MM/DD/YY HH:MM:SS <month>/ <day>/<year> <hours>:<minutes>:<seconds>
CITYPE	The component type.
CIName	The name of the component.
DataServerName	The name of the Data Server which sent this CI.
Expired	When checked, data has not been received within the time specified in the solution package that is hosting the data. If the solution package is configured to delete expired data, this row will be deleted if no data is received within the time specified for deletion. See the documentation for the solution package that is hosting the data for information on how to configure expiration and deletion times.

RTView CI Type Defs

This display provides component type definitions and shows the mapping of component types to caches as well as component types to alerts.

RTView CI Type Definitions						
13-Oct-2015 11:08 Data OK						
CI Type Definitions						
CITYPE	INDEXMAP	INDEXNAMES	RTVDISPLAY	CIVARMAP	DEFAULTQUALITY	OW
ACW	1	Dimension	acw_instance_summary	\$awsEc2InstanceId		1 Infrastr
AMX-HOST	1	AMX Host	amx_host_summary	\$amxHost		1 Infrastr
AMX-NODE	1,2	AMX Host;Node	amx_node_summary	\$amxHost;\$amxNode		1 Infrastr
AMX-SERVICE	1,2	Application;Service	amx_service_summary	\$amxApplication;\$amxService		1 Infrastr
AMX-SERVICENODE	1,2,3,4	AMX Host;Node;Application;Ser...	amx_servicenode_summary	\$amxHost;\$amxNode;\$amxApp...		1 Infrastr
BW6-APP	1,2,3	Domain;AppSpace;Application...	bw6_app_summary	\$bw6domain;\$bw6appspace;\$b...		1 Infrastr
BW6-APPNODE	1,2,3	Domain;AppSpace;AppNode	bw6_appnode_summary	\$bw6domain;\$bw6appspace;\$b...		1 Infrastr
BW6-PROCESS	1,2,3,4,5	Domain;AppSpace;AppNode;Ap...	bw6_process_summary	\$bw6domain;\$bw6appspace;\$b...		1 Infrastr
BW-ENGINE	1,2	AgentName;MicroAgentName	bw_engine_summary	\$bwserver;\$bwengine		1 Infrastr

Cache Map By CITYpe		Alert Map By CITYpe	
CITYPE	CACHENAME	CITYPE	ALERTNAME
ACW	AwsEc2InstanceStats	ACW	AcwInstanceCpuHigh
AMX-NODE	AmxNodes	ACW	AcwInstanceDiskReadBytesHigh
AMX-SERVICE	AmxServiceTotals	ACW	AcwInstanceDiskReadOpsHigh
AMX-SERVICEN...	AmxServices	ACW	AcwInstanceDiskWriteBytesHigh
BW6-APP	Bw6Apps	ACW	AcwInstanceDiskWriteOpsHigh
BW6-APP	Bw6ProcessTotalsByApp	ACW	AcwInstanceNetworkReadBytesHigh
BW6-APPNODE	Bw6AppNodes	ACW	AcwInstanceNetworkWriteBytesHigh
BW6-PROCESS	Bw6Processes	AMX-SERVICE	AmxServiceNodeHitRateHigh
BW-ENGINE	BwEngines	AMX-SERVICE	AmxServiceNodeResponseTimeHigh
BW-PROCESS	BwProcesses	AMX-SERVICE	AmxServiceNodeFaultRateHigh
BW-SERVER	BwServers	AMX-SERVICE	AmxServiceHitRateHigh
DB2	Db2DbSummary	AMX-SERVICE	AmxServiceResponseTimeHigh
DB2	Db2ResponseTime	AMX-SERVICE	AmxServiceFaultRateHigh
EM-SERVICE	RtvCmdDbServiceTable_local	AMX-SERVICE	AmxServiceNodeNotRunning
EM-SERVICE	RtvCmdDbServiceStats_local	BW6-APP	Bw6AppProcessCreatedRateHigh

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.

Fields and Data

This display includes:

CI Type Definitions

This table provides definitions for all CI Types.

CITYPE	The component type.
INDEXMAP	Number of indexes and the order in which they are used to form the CI Name.
INDEXNAMES	Semicolon-separated list of the index columns.
RTVDISPLAY	The name of the RTView display to drill-down to from the Alerts Table to see summary data for this CI Type. This is the target of the Go To CI button in the Alerts Table and in the Service Summary display.
CIVARMAP	The names of substitutions that must be set to drill-down to the display.
DEFAULTQUALITY	A flag indicating whether the lack of data is considered an error condition or not.
OWNER	The Owner the CITYPE is associated with, when the CMDB is populated automatically from CIs of this type.
AREA	The Area the CITYPE is associated with.
SERVICEGROUP	The SERVICEGROUP the CITYPE is associated with, when the CMDB is populated automatically from CIs of this type.

Cache Map By CITYPE

This table provides mapping of component types to caches for all component types.

CITYPE	The type of CI.
CACHENAME	The name of the cache associated with the component type.

Alert Map By CITYPE

This table provides mapping of component types to alerts.

CITYPE	The type of CI.
ALERTNAME	The name of the alert.

RTView KM Defs

This display shows the Key Metrics definitions for all CI Types. For details, see [“Available KM Metrics and Alerts”](#).

CITYPE	CACHENAME	SELECTOR	METRICNAME	AlertName
ACW	AwsEc2InstanceStats	Instance CPU Usage	CPUUtilization	AwsInstanceCpuHigh
AMX-SERVICE	AmxServiceTotals	Service Hits/Min	Hits Per Minute	AmxServiceHitRateHigh
AMX-SERVICE	AmxServiceTotals	Service Response Time	Avg. Response Time	AmxServiceResponseTimeHigh
AMX-SERVICENODE	AmxServices	Node Hits/Min	Hits Per Minute	AmxServiceNodeHitRateHigh
AMX-SERVICENODE	AmxServices	Node Response Time	Avg. Response Time	AmxServiceNodeResponseTimeHigh
BW6-APP	Bw6ProcessTotalsByApp	App Created / sec	RateCreated	Bw6AppProcessCreatedRateHigh
BW6-APP	Bw6ProcessTotalsByApp	App Exec Time / sec	RateTotal Execution	Bw6AppProcessExecutionTimeHigh
BW6-APPNODE	Bw6AppNodes	CPU Used %	Used CPU Percentage	Bw6AppNodeCpuUsedHigh
BW6-APPNODE	Bw6AppNodes	Mem Used %	Used Memory Percentage	Bw6AppNodeMemUsedHigh
BW6-PROCESS	Bw6Processes	Process Created / sec	RateCreated	Bw6ProcessCreatedRateHigh
BW6-PROCESS	Bw6Processes	Process Exec Time / sec	RateTotal Execution	Bw6ProcessExecutionTimeHigh
BW-ENGINE	BwEngines	CPU Used %	CPU %	BwEngineCpuUsedHigh
BW-ENGINE	BwEngines	Memory Used %	PercentUsed	BwEngineMemUsedHigh
BW-PROCESS	BwProcesses	Process Exec Time / sec	RateTotalExecution	BwProcessExecutionTimeHigh
BW-SERVER	BwServers	CPU Used %	CPU Usage %	BwServerCpuUsedHigh
DB2	Db2ResponseTime	Response Time	ResponseTimeMilliSec	Db2ResponseTimeHigh
EM-SERVICE	RtvCmdServiceStats_local	Alert Impact	AlertImpact	RtvEmServiceAlertImpactHigh
EMS-QUEUE	EmsQueues	Pending Msgs	pendingMessageCount	EmsQueuesPendingMessageCountHigh
EMS-QUEUE	EmsQueues	In Msgs / sec	inboundMessageRate	EmsQueuesInMsgRateHigh
EMS-QUEUE	EmsQueues	Out Msgs / sec	outboundMessageRate	EmsQueuesOutMsgRateHigh
EMS-QUEUE	EmsQueues	Consumers	consumerCount	EmsQueuesConsumerCountHigh

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 , 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.

Fields and Data

This display includes:

RTView Key Metrics Definitions

This table provides Key Metrics definitions for all CI Types.

CITYpe	The component type.
CACHENAME	The name of the cache that contains the Key Metric.
SELECTOR	The name used for this Key Metric in the Metric Name column of Key Metric displays.
METRICNAME	The name of the cache column that contains this Key Metric.
ALERTNAME	The name of the alert associated with this Key Metric. When blank, the Key Metric is not configured for inclusion in Key Metric displays.

CalcMode

The calculation used for the **Threshold %** value. The base value is calculated as the percent of the Key Metric value between **0** and the **ALARMLEVEL** of the associated alert. If the **CalcMode** is blank, this value is used. If the **CalcMode** is:

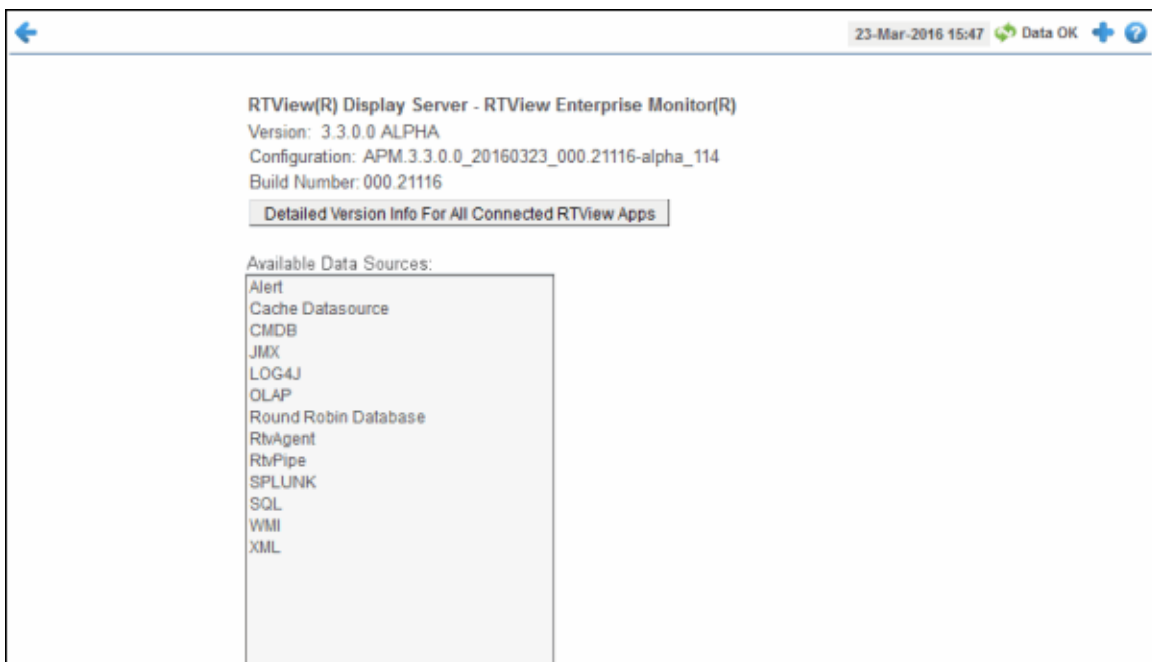
- **exp**, the value is adjusted so that lower values are diminished. Typically, this is used for memory metrics.
- **inverse**, the value is calculated in reverse of the standard thresholds. This is used when the associated alert is a low threshold alert.
- **invpct**, the value is calculated in reverse of the standard threshold and is assumed to be percent and therefore a value between **0** and **100**. This is used when the associated alert is a low threshold alert against a percent.
- **log**, a logarithmic algorithm is applied.

Level

The Key Metric level. Level **0** KMs are always displayed. Level **1** KMs are displayed if **Show More Metrics** is selected.

About

Get RTView Enterprise Monitor version and configuration information including a list of all available data sources.



Property Views

These displays show how your properties are configured and the values for all connected RTView processes. The displays are located under the **ADMIN** tab. Displays in this View are:

- [“Properties Configuration”](#): Table of properties configuration settings, per connection.
- [“System Properties”](#): Table of system properties for RTView processes, per connection.
- [“Applied Properties”](#): Table of all properties that were applied to RTView processes, per connection.
- [“All Properties”](#): Table of all properties that were read from the properties files and database regardless of whether or not the RTView process uses them.
- [“Properties Descriptions”](#): Table of all properties that are supported by RTView processes, per connection.

Properties Configuration

This display shows properties configuration information. The **Last Property Read Time** shows the last time that properties were read for the RTView process specified by the selected **Connection**.

Source: Data Server ▼	Connection: ALERT_SERVER ▼
Last Property Read Time: Mar 23, 2016 10:53:12 AM	


Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

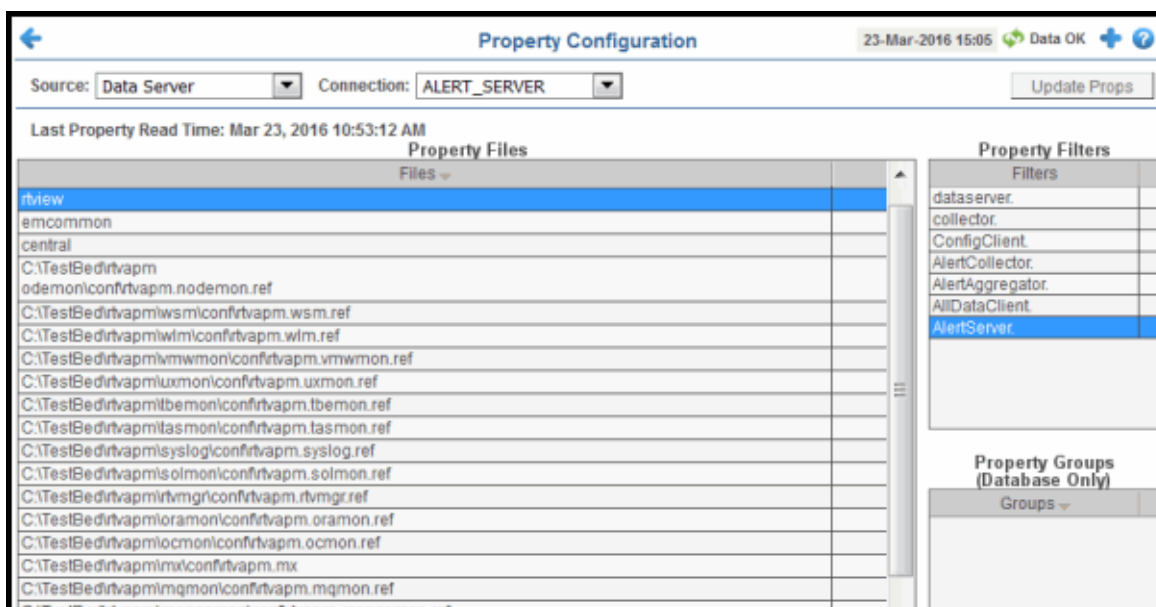
- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTView** is only visible when the **Source** is **RTView JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.







The **Property Files** table shows all of the properties files that were read by the RTView process specified by the selected **Connection** in the order they were read. The **Property Filters** table shows all filters that are applied to the properties. **Property Groups** shows all property groups that are applied to the properties. **Property Groups** are only used when reading properties from a database.

Click **Update Props** to have the RTView process specified by the selected **Connection** re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the [“Properties Descriptions”](#) display to find out if a specific property supports updates.

Right-click/**Export** to create a PDF image of the display. Click Sort  to order column data.



Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   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.

Note: The “Up” Arrow () opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking  opens the “Services CI Type Summary” display.

Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

Fields and Data

This display includes:

Update Props

Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the “Properties Descriptions” display to see if a specific property supports updates.

Last Property Read Time

The last time that properties were read for the RTView process specified by the selected **Connection**.

Property Files (table)	List of all properties files that were read by the RTView process specified by the selected Connection in the order they were read.
Property Filters (table)	List of all filters that are applied to the properties.
Property Groups	List of all property groups that are applied to the properties. Property Groups are only used when reading properties from a database.

System Properties


This display shows the System properties for the RTView process specified by the selected Connection.


Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.







Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.


Click **Update Props** to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the ["Properties Descriptions"](#) display to find out if a specific property supports updates.


Right-click/**Export** to create a PDF image of the display. Click Sort  to order column data.



Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   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.

Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

Fields and Data

This display includes:

- Update Props** Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the "[Properties Descriptions](#)" display to see if a specific property supports updates.
- System Properties (table)** List of all system properties for the RTView process specified by the selected **Connection**.
- Property** The name of the property.
- Value** The property setting.

Applied Properties

This display shows all properties that were applied to the RTView process specified by the selected **Connection**.

There are several reasons a property specified in a properties file might not be applied to an RTView process:

- the filter doesn't match.
- it was overridden in another property file.
- it was specified in a file that is not used by the RTView process.
- it was a property that is not supported in that RTView process (ex, a builder specific property would not be applied to a data server process).

You can filter the **Applied Properties** table using the **Filter Column** and **Field Value** fields. The **Clear Filter** button clears the filter. Double-click on a row in the table to drill down to the ["All Properties"](#) display filtered by the **Property Name** for that row.


Note: The double-click feature is not supported on iPad. iPad users can access the ["All Properties"](#) display from the navigation tree.

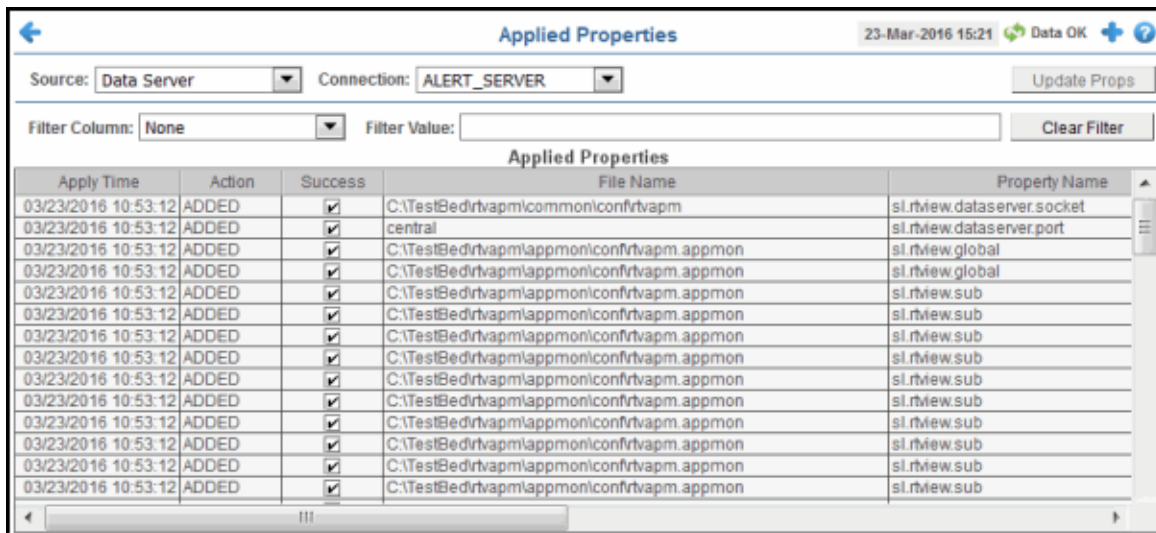
Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.







Click **Update Props** to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the ["Properties Descriptions"](#) display to find out if a specific property supports updates.


Right-click/**Export** to create a PDF image of the display. Click Sort  to order column data.



Apply Time	Action	Success	File Name	Property Name
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\common\conf\rtvapl...	sl.rview.dataserver.socket
03/23/2016 10:53:12	ADDED	✓	central	sl.rview.dataserver.port
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.global
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.global
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub
03/23/2016 10:53:12	ADDED	✓	C:\TestBed\rtvapl\appmon\conf\rtvapl.appmon	sl.rview.sub

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.
-   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.

Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

Fields and Data

This display includes:

- Update Props** Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the ["Properties Descriptions"](#) display to see if a specific property supports updates.
- Filter Column:** Select a column to filter the **Applied Properties** table.
- Filter Value:** Enter a string to filter the **Applied Properties** table.
- Clear Filter** Clears the filter.

Applied Properties (table)

- Apply Time** The last time this property was applied.

Action	Describes what occurred at Apply Time . <ul style="list-style-type: none"> • ADDED: Property was added. • REMOVED: Property was removed. • CHANGED: Property was modified.
Success	When the box is checked the Action was successful.
File Name	The source of this property. For properties read from a database this value is database .
Property Name	The name of the property after the property filter has been applied.
Property Value	The value of the property.
Handler	The RTView Handler that uses this property.

All Properties

This display shows all properties that were read from the properties files and database regardless of whether or not the RTView process uses them. There are several reasons a property specified in a properties file might not be applied to an RTView process:

- the filter doesn't match.
- it was overridden in another property file.
- it was specified in a file that is not used by the RTView process.
- it was a property that is not supported in that RTView process (ex, a builder specific property would not be applied to a data server process).

You can filter the **All Properties** table using the **Filter Column** and **Field Value** fields. The **Clear Filter** button clears the filter. Double-click on a row in the table to drill down to the ["All Properties"](#) display filtered by the **Property Name** for that row.

Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.

Click **Update Props** to have the RTView process specified by the selected **Connection** re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the “[Properties Descriptions](#)” display to find out if a specific property supports updates.

Order	File Name	Property Name	Property Value
0	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.cp	C:\TestBed\rtvapi\rtview\lib\rtvssa.jar
1	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.cp	C:\TestBed\rtvapi\common\lib\rtvapi_common.jar
2	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.cp	C:\TestBed\rtvapi\common\lib\gms\rtvutils.jar
3	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.cp	C:\TestBed\rtvapi\rtview\lib\rtvdebug.jar
4	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jvm	-Xmx256m
5	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jvm	-Xms128m
6	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.cmd_line	-nolibco
7	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.stylesheet	rtv_darkstyles,rtv_flat,rtv_html5
8	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.sql.dbretry	40000
9	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.global	rtv_global_vars.rtv
10	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.global	rtv_global_trendrange.rtv
11	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.xml.xmlsource	rtv_constants.xml 0 rtv_constants.xml 0 1
12	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jmx.jmxconn	local -- 'URL' local -- false
13	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.dsenable	jmx
14	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jmx.jmx_metrics_period	10000
15	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jmx.jmx_minreconnecttime	30
16	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jmx.jmx_mbeans_change_dyn	false
17	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jmx.jmxdsShowConnectionOnl	true
18	C:\TestBed\rtvapi\common\conf\rtvapi	sl.rtv.jvm	-Dcom.sl.rtv.customRtvApplManagerClassName=com.

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.
- 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.

Order	The order in which this property was read. For properties that support a single value that are specified multiple times, the one with the highest Order value will be applied.
File Name	The source of this property. It will be database for properties read from a database.
Property Name	The name of the property after the property filter has been applied.
Property Value	The value of the property.
Original Property Name	The name of the property before the property filter was applied. This will match the literal property string in your properties file.

Filter By:

Source: Select the **Source** of the connection to the RTView process for which you want to see property information.

Connection: Select the **Connection** to the RTView process for which you want to see property information.

Fields and Data

This display includes:

All Properties (table)	Update Props	Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the "Properties Descriptions" display to see if a specific property supports updates.
	Filter Column:	Select a column to filter the Applied Properties table.
	Filter Value:	Enter a string to filter the Applied Properties table.
	Clear Filter	Clears the filter.
	Order	The order in which this property was read. For properties that support a single value that are specified multiple times, the one with the highest Order value will be applied.
	File Name	The source of this property. It will be database for properties read from a database.
	Property Name	The name of the property after the property filter has been applied.
	Property Value	The value of the property.
	Original Property Name	The name of the property before the property filter was applied. This will match the literal property string in your properties file.

Properties Descriptions

This display shows one row for each property that is supported for the RTView process specified by the selected Connection.

Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.

Property	Multi	Updates	Handler	Deprecated	Depreciated
sl.rview.alert.actionauditdataserver	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.actionauditdb	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.actionaudittable	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.alertclearedcommand	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.alertcleartime	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.alertcommand	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.alertinitdelay	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.cleansettingstable	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.commentcommand	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.commentlimit	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.config	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.createDbTables	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.customAlertActionHandlerClass	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.custom_alertdef_prop	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.custom_event_attr	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.enableactionaudit	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.enablebuffer	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.enabled	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.enablessa	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.exitOnPersistInitFailed	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.history	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	
sl.rview.alert.lutupdatesnewdata	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>	

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.

Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

Fields and Data

This display includes:

- Update Props** Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the ["Properties Descriptions"](#) display to see if a specific property supports updates.

All Properties (table)

- Property** The name of the property

Multi	Box is checked if this property supports multiple values.
Updates	Box is checked if this property supports updates.
Handler	The name of the RTView Handler that uses this property.
Deprecated	Box is checked if this property is deprecated.
Deprecation Info	If the property is deprecated, this lists the currently supported property to use instead.

Diagram Views


The RTView Enterprise Monitor comes with the Diagram Generator, a feature that auto-creates diagram displays which mirror your system components and hierarchy.

When you monitor applications with complex architectures, it is often very beneficial to visualize the health of individual components within the context of the application hierarchy. This allows you to understand the complete architecture which is supporting the application as well as understand how individual components may affect the behavior of other components. RTView Enterprise Monitor allows you to manually construct such views, but at times this might become too time consuming to maintain if there are many applications to model or the architecture is constantly changing. The Diagram Generator is a feature which allows for the automatic generation of these application diagrams using application meta-data, without having to manually construct them.

The Diagram Generator, located under **CUSTOM Tab/Diagram Views**, is comprised of several displays which you use to create your diagram displays.

This section includes the following Diagram Generator instructions and displays:

- "Steps to Create a Diagram Display"
- "Create an Object Template Display"
- "Node Administration Display"
- "Link Admin Display"
- "Diagram Properties Admin Display"
- "Add Diagrams to your Project"
- "View Diagram Displays"
- "Optional Diagram Display Customizations"



The Diagram Generator feature uses a database table of your nodes and database table of your links to create the diagrams. For example, the diagram below was generated from the tables shown next to it. The order of the nodes in the table controls the order of the nodes in the diagram.

The screenshot shows the RTView Enterprise Monitor interface. The top navigation bar includes tabs for SERVICE TREE, SERVICE VIEWS, COMPONENTS, ALERTS, ADMIN, and CUSTOM. The CUSTOM tab is selected, and the 'Sample Diagram' is displayed. The diagram shows a flow from Node 1 to Node 2, then to a group of three nodes (Node 3A, Node 3B, Node 3C), and finally to Node 4 and Node 5. Below the diagram, three tables are shown: Nodes, Links, and Diagram Props.

Nodes Table:

ID	Type	Properties	Layout
G1	Group DefaultGroup	label=Group 1	
N1	DefaultNode	label=Node 1~dd=rtv_diagram_test2.rtv~mouseOverText=dd to diagram	
N2	DefaultNode	label=Node2	group=G1
N3	GreenNode	label=Node 3	tier=SL_1~group=G1
N3A	GreenNode	label=Node 3A	tier=SL_1~group=G1
N3B	GreenNode	label=Node 3B	tier=SL_1~group=G1
N4	DefaultNode	label=Node 4	
N5A	DefaultNode	label=Node 5A	tier=SL_2

Links Table:

ID	Type	Properties	Start Node	End Node
L1	DefaultLink		N1	N2
L2	DefaultLink		N2	N3
L3	DefaultLink		N2	N3A
L4	DefaultLink		N2	N3B
L5	DefaultLink		N4	N6
L6	GreenLink		N6	N4
L7	GreenLink		N6	N7
L8	GreenLink		N7	N8

Diagram Props Table:

Name	Value
hGap	30
objectTemplateDisplay	test_clone_source.rtv
orientation	VERTICAL
hAlignment	CENTER
vAlignment	CENTER
margins	40,10,10,10
wrapDiagram	false
wrapSpacing	20

Diagram Generator Demo

You can view a demo of the Diagram Generator, located under **CUSTOM/Sample Diagram**. This demo shows a simple diagram as well as the UI used to construct diagrams. The icons in the sample diagram come from a sample object template display. When creating diagrams for your project, you will create an object template display with icons that are appropriate for the process you wish to diagram.

Steps to Create a Diagram Display

To create a custom diagram display using the Diagram Generator:

1. If you are using EM 3.5 or earlier, see **Upgrading the emsample** project is already configured to include the Diagram Generator and no setup is required.
2. [“Create an Object Template Display”](#) using the Display Builder. This step is optional. The object template display serves as your palette of objects (icons, links, shapes, and so forth) for building your diagrams. A sample object template display is built-in that can be used for simple diagrams and to demo the Diagram Generator. When creating diagrams for your project, you should create an object template display with icons that are appropriate for the process you wish to diagram. If you are going to create a custom object template display, you must create it and add it to the Diagram Properties before defining nodes and links for your diagrams.
3. In the RTView Enterprise Monitor, open the [“Node Administration Display”](#) display, located under **CUSTOM Tab/Diagram Views/Diagram Admin** and add nodes to your diagram.
4. Open the [“Link Admin Display”](#) display and add links to your diagram.
5. Open the [“Diagram Properties Admin Display”](#) display and format the layout of your diagram display.
6. [“Add Diagrams to your Project”](#) to publish the diagram display.
7. [“View Diagram Displays”](#) to confirm settings.
8. [“Optional Diagram Display Customizations”](#): These customizations are not required.

Create an Object Template Display

This section describes how to create an object template display using the Display Builder. The object template display serves as your palette of objects (icons, links, shapes, and so forth) that you use to build your diagrams.

Assumptions:

- You have familiarized yourself with the [“Diagram Generator Demo”](#).
 - You are familiar with using the Display Builder.
1. Create a display in the Display Builder that contains all the icons you want to use for the nodes, links and icons in your diagram displays.
 2. For each object, specify a user friendly name in the **objName** field (it must be alpha-numeric but can contain under-bar). The **objName** field will be referenced in the **Type** field when you add nodes and links to your diagram. You will be able to override any of the properties by specifying them in the **Properties** field.
 3. Save this file to the **RTViewEnterprise/emsample/servers/central** directory.
 4. Use the [“Diagram Properties Admin Display”](#) to set the **objectTemplateDisplay** property to the name of this file.

For assistance, contact Technical Support.

Proceed to [“Node Administration Display”](#).

Node Administration Display

Use this Diagram Generator display to create new diagrams and add or edit nodes in existing diagrams. Diagram definitions are stored in the DIAGRAMS database.

To define a new diagram, enter a new diagram name in the **Diagram** field, then fill in the fields in the **Enter Values** section to define the first node in the diagram.

To add a new node to a diagram, enter the name of the diagram to which you want to add the node, then fill in the fields in the **Enter Values** section, described below, and click **Add Node**.

To edit an existing node, select it in the table to populate the **Enter Values** fields at the bottom of the display, make your changes and click **Update Node**.

The screenshot shows the 'Node Table Administration' window in RTView Enterprise Monitor. The top navigation bar includes 'SERVICE TREE', 'SERVICE VIEWS', 'COMPONENTS', 'ALERTS', 'ADMIN', and 'CUSTOM'. The 'CUSTOM' tab is active, showing 'Custom Views' on the left and 'Node Table Administration' on the right. The 'Node Table Administration' section has a 'Diagram Filter' set to 'All' and a 'Click on a row to edit...' link. The table lists nodes with the following data:

Order	Diagram	ID	Type	Layout	Properties
0	sample_diagram	N1	DefaultNode		label=Node 1
1	sample_diagram	N2	DefaultNode		label=Node 2
2	sample_diagram	N3	GreenNode	tier=3	label=Node 3A
3	sample_diagram	N4	GreenNode	tier=3	label=Node 3B
4	sample_diagram	N5	GreenNode	tier=3	label=Node 3C
5	sample_diagram	N6	DefaultNode	group=G1	label=Node 4
6	sample_diagram	N7	DefaultNode	group=G1	label=Node 5
7	sample_diagram	G1	Group DefaultGroup		label=Group 1

Below the table is the 'Enter Values' section with the following fields and buttons:

- Diagram:** sample_diagram
- ID:** N1
- Layout:** (empty)
- Type:** DefaultNode
- Is Group:** (checkbox)
- Properties:** label=Node 1
- Buttons:** Update Node, Insert Node, Delete Node, Cancel, Move Node Up, Move Node Down, Preview

For each node in your diagram display, fill in the following fields:

Diagram

Required. The name of the diagram. **Note:** The diagram name GLOBAL is reserved for global properties in Diagram Props. The value in this column will be used to identify this diagram when you add this diagram to the navigation tree.

ID

Required. Must be unique within the diagram across nodes and links. Use this value in the Link Node1 and Node2 fields to refer to this node. The value must be alpha-numeric but may contain underbars (_).

Layout

Optional. One or more layout options for the icon. All icons with the same tier=value will be positioned in a single tier (row if the diagram property orientation=VERTICAL, column if orientation=HORIZONTAL). Each tier is centered along the diagram's centerline unless only one node in a tier has a link to another tier. In that case the node with the link is centered on the diagram's centerline.

All icons with the same group=value will have a group object drawn behind them. The group value must be the ID of a node whose type is defined as a Group.

To specify both a tier and a group, separate them with a ~. For example:

tier=T1~group=G1

Type	Required. The name of the object in the objectTemplateDisplay file to use as the icon for this node. When you select a type from the list, you will see a preview of it to the right of the Type field.
Is Group	Optional. Check to specify that this node is a Group. Groups are only drawn if at least one node references them in their Layout field. They are drawn behind the nodes that reference them and their extent is set to the combined extent of all nodes that reference them. In wrapped diagrams, if the nodes in the group break across multiple tiers, the group object will be broken across the tiers as well.
Properties	<p>Optional. One or more properties to set on the node icon delimited by ~. Syntax is propName=propVal~propName2=otherPropVal.</p> <p>Note that property values must be specified as they are saved in .rtv files, which is not necessarily the same as they are shown in the Object Properties dialog in the Display Builder.</p> <p>In addition to properties on the RTView object, you can also specify hGap or vGap to override the diagram property hGap or vGap for this object. The hGap is applied to the left of an object and the vGap is applied above an object.</p> <p>For example, you must specify the font index instead of the font name for font properties, and the color index instead of the color for color properties.</p>

Use the following buttons to save changes to the database and to preview the diagram display (after the changes have been saved to the database):

Update Node	Save changes to the selected node to the database. This is only enabled if the selected node is already in the database.
Insert Node	Insert a new node to the database. This is only enabled if the selected node is not in the database.
Delete Node	Delete the selected node from the database. This is only enabled if the selected node is already in the database.
Cancel	Clear the Enter Values fields.
Move Node Up	Move the selected node up in the diagram. Nodes are laid out in the diagram according to their order.
Move Node Down	Move the selected node down in the diagram. Nodes are laid out in the diagram according to their order.
Preview	Open a window showing the selected diagram as it is saved in the database. Changes to the diagram will not update an open preview window. To update the diagram in the preview window, close and reopen the window.

Link Admin Display

Use this display to add or edit links in your diagrams. To add a link, enter the name of the diagram containing the nodes you want to link, then fill in the fields below for each link you want to add. To edit an existing link, select it in the table.

Link Table Administration

Diagram Filter: All

Diagram	ID	Type	Node 1	Node 2	Properties
sample_diag	L1	DefaultLink	N1	N2	
sample_diag	L2	DefaultLink	N2	N4	
sample_diag	L3	DefaultLink	N4	N6	
sample_diag	L4	GreenLink	N3	N4	arrow1VisFlag=1
sample_diag	L5	GreenLink	N4	N5	arrow1VisFlag=1

Enter Values

Diagram: sample_diagram ID: L2 Type: DefaultLink

Node 1: N2 Node 2: N4

Properties:

Update Link Insert Link Delete Link Cancel Preview

Diagram

Required. The name of the diagram. **Note:** The diagram name GLOBAL is reserved for global properties in Diagram Props.

ID

Required. Must be unique within the diagram across the node and link tables. Value must be alpha-numeric but may contain under-bars (_).

Type

Required. The name of the link object in the objectTemplateDisplay to use for this link.

Node 1

Required. The ID of the start node for the link.

Node 2

Required. The ID of the end node for the link.

Properties

Optional. One or more properties to set on the icon delimited by ~. Syntax is propName=propVal~propName2=otherPropVal

Use the following buttons to save link changes to the database and to preview the diagram display (after the changes have been applied to the database):

Update Link

Save changes to the selected link to the database. This is only enabled if the selected link is already in the database.

Insert Link

Insert a new link to the database. This is only enabled if the selected link is not in the database.

Delete Link

Delete the selected link from the database. This is only enabled if the selected link is already in the database.

CancelClear the **Enter Values** fields.**Preview**

Opens a window showing the selected diagram as it is saved in the database. Changes to the diagram will not update the preview window. To update the preview, close and reopen it.

Diagram Properties Admin Display

Use this Diagram Generator display to configure “[Diagram Properties](#)” for your diagrams. Diagram properties are settings that are applied to the diagrams as a whole, such as orientation, alignment and spacing. Properties that use GLOBAL for the diagram name are applied to all diagrams. You can override a diagram property for a single diagram by using the name of that diagram in the **Diagram** field.

The screenshot shows the 'Diagram Properties Administration' window in RTView Enterprise Monitor. The left sidebar contains a 'Custom Views' menu with options like 'Diagram Views', 'Sample Diagram', 'Diagram Admin', 'Node Admin', 'Link Admin', and 'Diagram Prop Admin'. The main area displays a table of properties for 'GLOBAL' diagrams. Below the table is an 'Enter Values' section with input fields for 'Diagram', 'Property', and 'Value', and buttons for 'Update Property', 'Insert Property', 'Delete Property', and 'Cancel'.

Diagram	Property Name	Property Value
GLOBAL	hAlignment	CENTER
GLOBAL	orientation	VERTICAL
GLOBAL	vAlignment	CENTER

Enter Values

Diagram: GLOBAL

Property: vAlignment

Value: CENTER

Buttons: Update Property, Insert Property, Delete Property, Cancel

Diagram

Required. The name of the diagram or GLOBAL if it should be applied to all diagrams.

Property

Required. The name of the property.

Value

Required. The value of the property.

Use the following buttons to save diagram property changes to the database:

Update Property

Save changes to the selected property to the database. This is only enabled if the selected property is already in the database.

Insert Property

Insert a new property to the database. This is only enabled if the selected property is not in the database.

Delete Property

Delete the selected property from the database. This is only enabled if the selected link is already in the database.

Cancel

Clear the Enter Values fields.

Diagram Properties

Use the following properties in the **Property/Value** fields in the [“Diagram Properties Admin Display”](#) display to configure diagram display.

Property Name	Description
hGap	Horizontal space between nodes in pixels. This can be overridden per-node in the node properties. Default is 28.
vGap	Vertical space between nodes in pixels. This can be overridden per-node in the node properties. Default is 22.
objectTemplateDisplay	The name of the file containing the nodes and links to use for the diagrams. The name of each node and link in this file corresponds to the name in the Type field of the node and link tables.
orientation	The orientation of the diagram. Default is VERTICAL which lays the nodes out in the order specified from the top of the display to the bottom with nodes in the same tier laid out left to right. HORIZONTAL lays out the nodes in the order specified from left to right with nodes in the same tier laid out top to bottom.
hAlignment	Controls the horizontal alignment of the diagram in the available space. Default is CENTER which centers the diagram in the available space. Options are: LEFT – Position the diagram at the left of the available space. CENTER – Position the diagram in the center of the available space. CENTERLINE – Position the centerline of the diagram in the center of the available space. This option is only valid if orientation=VERTICAL and wrapDiagram=false. If orientation=HORIZONTAL or wrapDiagram=true, CENTER will be used instead. RIGHT – Position the diagram at the right of the available space.
vAlignment	Controls the vertical alignment of the diagram in the available space. Default is CENTER which centers the diagram in the available space. Options are: TOP – Position the diagram at the top of the available space. CENTER – Position the diagram in the center of the available space. CENTERLINE – Position the centerline of the diagram in the center of the available space. This option is only valid if orientation=HORIZONTAL and wrapDiagram=false. If orientation=VERTICAL or wrapDiagram=true, CENTER will be used instead. Bottom – Position the diagram at the bottom of the available space.
wrapDiagram	If true, the diagram wraps into columns if orientation=VERTICAL or rows if orientation=HORIZONTAL. Default is false.
wrapSpacing	The number of pixels between columns/rows if wrapDiagram = true. Default is 20.
margins	Sets the minimum amount of space between the edge of the display and the diagram in pixels. You can either specify one value that will be used for all 4 margins or a comma separated list of 4 values in this order: top, left, bottom, right. Default is 40,10,10,10 which sets the top margin to 40 pixels and the left, bottom and right margins to 10 pixels.
deleteSavedDiagramNodes	This option applies to diagrams that were manually edited as described in “Edit Diagrams in the Display Builder” . If true, delete any diagram nodes that were saved to the display in the Display Builder. Default is false.

Proceed to [“Add Diagrams to your Project”](#) to publish your displays.

Add Diagrams to your Project

This section describes how to add diagrams to your EM project. After you define one or more diagrams as described in [“Node Administration Display”](#), you can add a display for each diagram to the EM navigation tree. The navigation tree entry for each diagram should look like this:

```
<node label="Diagram 1" display="rtv_diagram_cache" subs="$diagramName:diagram1
$diagramTitle:'Diagram 1'"/>
```

The value for **label** is the label to use in the navigation tree. The value for **display** is **rtv_diagram_cache** unless you have a custom diagram background as described in [“Customize the Diagram Background Display”](#), in which case you should use the name of that display instead.

The **subs** values are as follows:

- **\$diagramName** – Set this to the name of your diagram. This corresponds to the value in the DIAGRAM column in the database.
- **\$diagramTitle** – Set this to the value to use for the title label in the diagram display.

Note: You must use single quotes around any substitution values that contain spaces.

Proceed to [“View Diagram Displays”](#).

View Diagram Displays

After you add one or more diagram displays to your EM project as described in the [“Add Diagrams to your Project”](#), open them in the navigation tree. Note that:

- Diagram definitions are only read when the display is opened. If you edit the diagram definition for an open diagram display, you must re-open the diagram display to see the changes.
- When you resize the window, the diagram display auto-resizes to fill the available space, and also positions the diagram in the available space according to the **margin** and **alignment** Diagram Properties. If you resize the window smaller than 800x576 or the area required to display the diagram (whichever is larger) scrollbars auto-appear.
- Diagrams with the **wrapDiagram** Diagram Property set to true reposition nodes to use the available space when the window is resized.

Optional Diagram Display Customizations

This section includes:

- [“Edit Diagrams in the Display Builder”](#):
- [“Customize the Diagram Background Display”](#):
- [“Customize the Diagram Database”](#):

Edit Diagrams in the Display Builder

To manually edit your generated diagram, use the Display Builder to open the diagram in the **RTViewEnterprise/emsample/servers/central** directory.

Note: If you created a custom diagram background display as described in [“Customize the Diagram Background Display”](#), use the name of that file instead of **rtv_diagram_cache** in the instructions below. Run the Display Builder in the **RTViewEnterprise/emsample/servers/central** directory as follows (where **diagramName** is the name of the diagram you want to modify):

```
runb_appmon -sub:$diagramName:diagramName rtv_diagram_cache
```

Edit the diagram and save the display as **rtv_diagram_cache_diagramName**, replacing the **diagramName** with the name of your diagram. Update the corresponding navigation tree entry to use the new display name.

Important: Do NOT to save these changes to **rtv_diagram_cache.rtv** or these nodes will show up in all of your diagrams.

When you view this diagram, any saved diagram nodes and links that are no longer in the diagram definition will be removed and any new nodes in the diagram definition will be added to the bottom left corner. You need to position those new nodes by hand in the Display Builder. The Properties from the database will be applied to diagram nodes that were saved in the Display Builder. The diagram will still be positioned in the Display Viewer according to the **alignment** and **margin** Diagram Properties when the window is resized. However, for diagrams where **wrapDiagram** is set to true, the diagram will not be re-wrapped to fit the available space.

Customize the Diagram Background Display

To create a custom version of the diagram background display, open **rtv_diagram_cache.rtv** in the Display Builder from the **central** directory of your EM project as follows:

```
runb_appmon rtv_diagram_cache
```

Modify the display and save it under a new name in the **central** directory. The name must start with **rtv_diagram**. When adding diagram displays to the navigation tree as described in [“Add Diagrams to your Project”](#), use the name of this file instead of **rtv_diagram_cache**.

When modifying the display, use the following guidelines:

- Do not change the **Resize Mode**. It must be set to **Crop**.
- When you resize this display in the viewer or thin client, objects will be positioned according to their anchor properties.
- Do not remove the **dg_include_cache.rtv** entry from the list of included files. This file reads the diagram data and creates the data structures required to generate the diagrams.

Customize the Diagram Database

Diagram definitions are stored in the DIAGRAM database. By default, an HSQLDB database is used. Schemas for all supported databases are provided in **RTVAPM_HOME\dg\dbconfig**. To change to another supported database, use the schema for your database to create the diagram tables and add this property to the central properties file for your project (**central.properties** in **emsample**) replacing the user name, password, URL and driver with the appropriate information for your database:

```
ConfigCollector.sl.rtvview.sql.sqldb=DIAGRAMS sa - jdbc:hsqldb:hsq://localhost:9099/  
rtvdiagram org.hsqldb.jdbcDriver - false true
```


CHAPTER 3 RTView DataServer for IBM

The RTView DataServer for IBM provides a way to create connections and modify default configuration settings for the various solution packages and sends collected data to RTView Central, which contains the displays associated with the RTView DataServer for IBM that help you to monitor the health and performance across your IBM components.


RTView Central contains the following Views and their associated displays that will be populated with data collected via the RTView DataServer for IBM:

- "IBM DB2"
- "IBM MQ"
- "IBM WebSphere"

The RTView *DataCollector* for IBM is also available for use with the RTView DataServer for IBM. RTView DataCollector for IBM is used for collecting and sending data to one or more data servers. The RTView DataCollector for IBM is also useful if you need to distribute data collection.

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

The Solution Package for IBM® DB2 includes a high level heatmap and tabular displays as well as drilldown views to access real-time and historical performance metrics for each DB2 Database in your monitored services and applications.



The following IBM DB2 Views can be found under **Components** tab > **Databases**> **IBM DB2 Database**:

- "DB2 Instances"
- "DB2 Members"
- "All DB2 Databases"
- "Single Database"

Displays in this View are:


- "All Instances Heatmap"
- "All Instances Table"
- "Instance Summary"




All Instances Heatmap

View current alert status and performance metrics of all or just one of your IBM DB2 instances. Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **Response Time**, **I/O Wait Time %**, **Network Wait Time %**, **Agent Wait Time %**, **Avg Deadlocks per Activity**, **Avg Lock Escalations per Activity**, **Avg Lock Timeouts per Activity**, **Avg Lock Waits per Activity**, **Rows Read per Rows Returned**, **Activites/sec**, **App Requests/sec**, **App Commits/sec**, **App Rollbacks/sec** or **Buffer Pool Hit Ratio %**.

Answer questions such as, Are any instances reaching a state of critical health? Do I need to allocate more tablespace to an instance? Are response times slow on any instances? Are application deadlocks causing bottlenecks on any instances? Is processing load evenly distributed across instances?

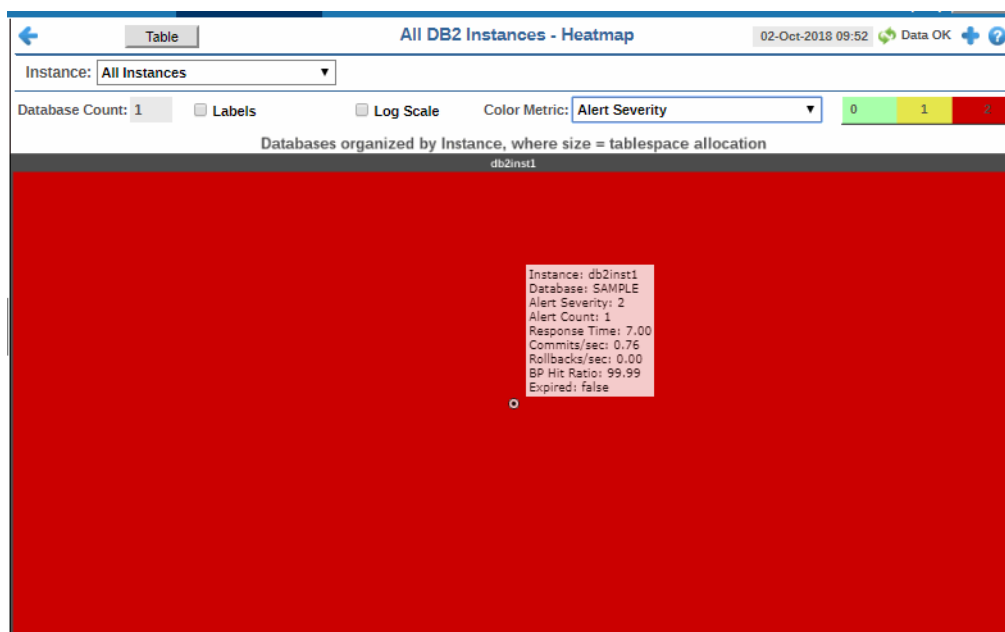
Each rectangle in the heatmap represents a different instance, where the rectangle color indicates the most critical alert state for items associated with that instance, and the rectangle size represents the tablespace allocation size for the instance.


Each metric has its own color gradient bar legend that maps values to colors. By default, the Alert Severity metric is shown, which is the current alert severity for items associated with the rectangle. Values range from **0** - **2**, as indicated in the color gradient  bar:

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










Other performance metrics include wait times, application requests and rollbacks.

Use the **Labels** check-box ☒ to include or exclude labels in the heatmap, use the **Log Scale** check-box ☒ to apply log scale and mouse over a rectangle to see additional metrics. Click a rectangle to drill down on an instance in the ["Instance Summary"](#) display.








Fields and Data:


Instance:	Select an instance.
Labels	Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
Log Scale	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Metric	Choose a metric to view in the display. For details about the data, refer to vendor documentation.
Alert Severity	<p>The current alert severity for items associated with the rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  (2) Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  (1) Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  (0) Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning unacknowledged alerts for items associated with the rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
Response Time	<p>The average response time, in milliseconds, for items associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of the Db2ResponseTimeHigh alert. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.

I/O Wait Time %

The percentage of the wait time being used by I/O processes. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2ResponseTimeHigh** alert threshold. The middle value in the gradient bar indicates the middle value of the range.


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


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


 Green indicates that no metrics have exceeded their alert thresholds.


The percentage wait time taken by I/O operations.

Network Wait Time %


The percentage of the wait time being used by network processes. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2ResponseTimeHigh** alert threshold. The middle value in the gradient bar indicates the middle value of the range.


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


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


 Green indicates that no metrics have exceeded their alert thresholds.

Agent Wait Time %


The percentage of the wait time being used by agent processes. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2ResponseTimeHigh** alert threshold. The middle value in the gradient bar indicates the middle value of the range.


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


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


 Green indicates that no metrics have exceeded their alert thresholds.

Avg Deadlocks per Activity


The average number of application deadlocks per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2DeadlockRateHigh** alert. The middle value in the gradient bar indicates the middle value of the range.


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


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


 Green indicates that no metrics have exceeded their alert thresholds.

Avg Lock Escalations per Activity

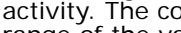
The average number of application deadlocks that were escalated per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2DeadlockRateHigh** alert. The middle value in the gradient bar indicates the middle value of the range.


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


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


 Green indicates that no metrics have exceeded their alert thresholds.

Avg Lock Timeouts per Activity

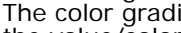
The average number of application lock timeouts per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2LockWaitTimeHigh** alert. The middle value in the gradient bar indicates the middle value of the range.


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


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


 Green indicates that no metrics have exceeded their alert thresholds.

Avg Lock Waits per Activity


The average number of application lock waits per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2LockWaitTimeHigh** alert. The middle value in the gradient bar indicates the middle value of the range.

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


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

 Green indicates that no metrics have exceeded their alert thresholds. The average number of application deadlock timeouts per activity.

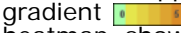
Rows Read per Rows Returned

The total number of rows read per number of rows returned. The color 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 value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


Activites/sec

The rate of activities (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.


App Requests/sec

The rate of application requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.


App Commits/sec


The rate of application commits (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.


App Rollbacks/sec


The rate of application rollbacks (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.

Buffer Pool Hit Ratio %

The percentage Tablespace used by the buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2UsedTablespaceHig** alert threshold. The middle value in the gradient bar indicates the middle value of the range.

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


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


 Green indicates that no metrics have exceeded their alert thresholds.

All Instances Table

This display provides a list of all DB2 instances, configuration details about each instance and their alert status. Each row in the table is a different DB2 instance.

Click a column header to sort column data in numerical or alphabetical order. Investigate by clicking a row to view details for an instance in the ["Instance Summary"](#) display.

All DB2 Instances - Table										
Instance Count: 1		All DB2 Instances								
Instance Name	Alert Status	Partitionable?	# Partitions	# DB's Monitored	Pointer Bit Size	Service Level	Release #	Build Level	Fix Pack #	Temp Fix ID
db2inst1		0	1	1	64	DB2 v10.5.0.1	0602010E	special_31160	1	IP23526_31160



Instance Count	The number of instances in the table.
Instance Name	The name of the instance.
Alert Status	<p>● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</p> <p>● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</p> <p>● Green indicates that no metrics have exceeded their alert thresholds.</p>
Partitionable?	<p>Indicates whether the instance can be partitioned.</p> <p>0 is No</p> <p>1 is Yes</p>
# Partitions	The number of partitions for the instance.
#DB's Monitored	The number of databases monitored.
Pointer Bit Size	The pointer bit size (e.g. 64).
Service Level	The installed DB2 software release version.
Release #	The installed DB2 software release number.
Build Level	The installed DB2 software build number.
Fix Pack #	The number of the installed DB2 Fix Pack.
Temp Fix ID	The IBM fix ID.

Instance Summary

Select an instance **Instance** to see the following information for that instance:

CPUs	The number of CPUs and the number of CPUs in use.
CPU Speed	The processor speed.
Total RAM	The total amount of RAM, in megabytes.
Total Virtual Memory	The total amount of virtual memory, in megabytes.
CPU Load Short	Amount of processor load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
CPU Load Medium	Amount of processor load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
CPU Load Long	Percentage of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
% CPU Usage	The percentage of CPU used.

Instance Databases Table

Each row is a different database. Column values describe the database.

Database Count: The number of databases in the table.

Database	The name of the database.
Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this row will be deleted if no data is received within the time specified for deletion.
Connected	When checked, the database is connected.
Severity	The alert status:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Response Time	The response time, in milliseconds.
I/O Wait Time %	The percentage wait time taken by I/O operations.
Network Wait Time %	The percentage wait time taken by the network.
Agent Wait Time %	The percentage wait time taken by agents.
Avg Deadlocks per Activity	The average number of application deadlocks per activity.
Avg Lock Escalations per Activity	The average number of application deadlock escalations per activity.
Avg Lock Timeouts per Activity	The average number of application deadlock timeouts per activity.
Avg Lock Waits per Activity	The average number of application deadlock waits per activity.
Rows Read per Rows Returned	The number of rows read per number of rows returned.
Activites/sec	The number of activities per second.

App Requests/sec	The number of application requests per second.
Commits/sec	The number of application commits per second.
Rollbacks/sec	The number of application rollbacks per second.
Buffer Pool Hit Ratio %	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
Activity Wait Time %	The percentage wait time taken by activities.
Avg Request CPU Time	The average amount of CPU time used by requests, in seconds.
Compile Proc Time %	The percentage of time used for compiling processes.
Routine Time Request %	The percentage of time used for routine request processes.
Section Time %	The percentage of time used for section processes.
Section Sort Time %	The percentage of time used for sorting section processes.
BP Hit Ratio %	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
Transaction Time %	The percentage of time used for transaction processes.
Utils Proc Time %	The percentage of time used for utilities processes.
Timestamp	The data and time this data was last updated.

DB2 Members

Displays in this View are:

- ["All Members Table"](#): Get configuration details and utilization metrics for host members of an instance.
- ["All Members Grid"](#): View trend graph for CPU load and CPU utilization for each member of an instance.
- ["Member Summary"](#): Investigate TableSpaces and buffer pool utilization metrics for a host member.

All Members Table

Select an instance **Instance** to see a list of all host members of an instance.

Each row is a different host that shows configuration details and utilization metrics for each. Details include OS, RAM, CPU load (short, medium, long) and virtual memory.

Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for a host member in the “Member Summary” display.

DB2 All Members Table

15-Oct-2018 14:22 Data OK

Instance: db2inst1

Member Count: 1

Host Name	Expir...	Instance	Mem...	OS	CPUs	CPU Speed	Total R...	Total Virt...	CPU L...	CPU Lo...	CPU L...	% Us
SLHOST93	<input type="checkbox"/>	db2inst1	0	Linux 2.6.32(x86_64)	1 of 1 in use	2,300	3,835	7,867	0.12	0.07	0.02	

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.

Instance Members Table

Each row is a different host member. Column values describe the host.

Member Count:	The number of members in the table.
Host Name	The name of the host.
Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this row will be deleted if no data is received within the time specified for deletion.
Instance	The name of the instance that the host is a member of.
Member	The name of the member.
OS	The installed operating system.
CPUs	The number of CPUs and the number of CPUs in use.
CPU Speed	The processor speed.
Total RAM	The total amount of RAM, in megabytes.

Total Virtual Memory	The total amount of virtual memory, in megabytes.
CPU Load Short	Amount of processor load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
CPU Load Medium	Amount of processor load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
CPU Load Long	Percentage of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
% CPU Usage	The percentage of CPU used.
Timestamp	The data and time this data was last updated.

All Members Grid

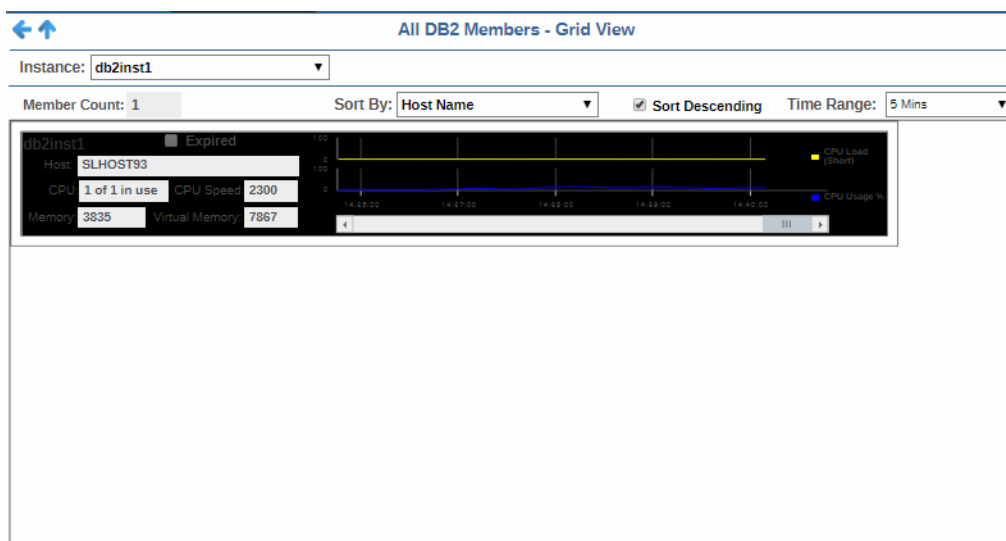
This display provides a grid view of a subset of data in the [“All Members Table”](#) display.

Each grid object is a different host member of the selected instance. Each has a trend graph that traces the **CPU Load Short** and **% CPU Usage** for the host member. Choose a **Time Range** for the trend graph to trace, or select **All Data** to include all available data in the trace. Scroll forward and backward in the trend graph and click **Reset** to return to the original state.

Show **Expired** members first using **Sort By: Expired**. Or show members with most severe **CPU Load Short** values first, and most severe **CPU Usage** values first.

Toggle **Sort Descending** to order the grid objects.

Investigate performance metrics details for a host member in the [“Member Summary”](#) display by clicking on the (left or right side of the) grid object.



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.
- 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.

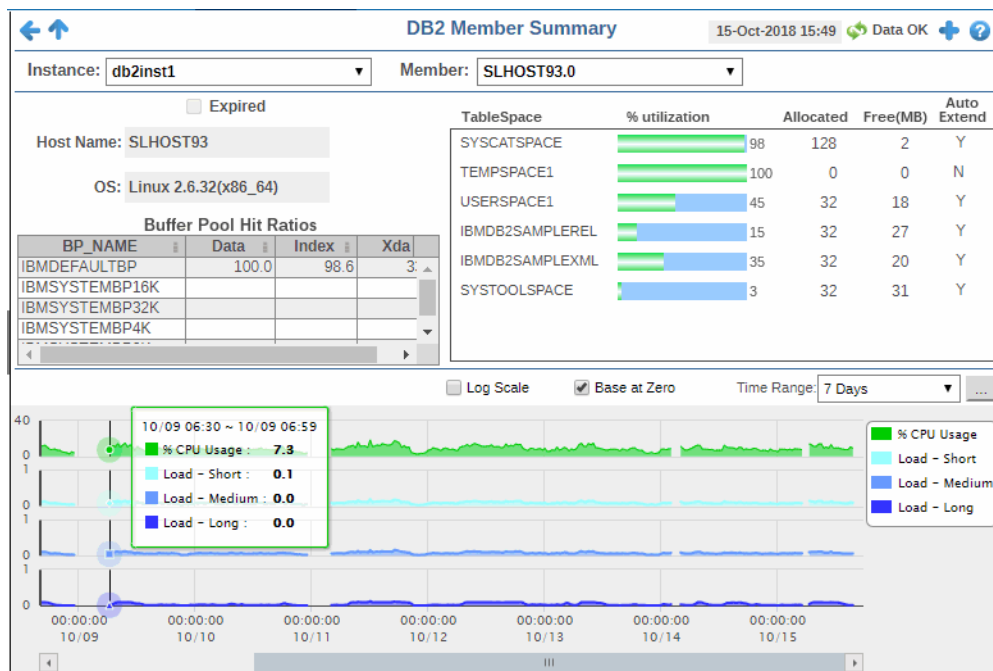
Instance:	Select an instance.														
Member Count:	The number of members in the display.														
Sort By:	Orders the grid objects as follows: <ul style="list-style-type: none"> • Instance: Instance names. • Host Name: Host member names. • CPU Load - Short: Current host member CPU Load - Short values, from most to least critical. • CPU Usage: Current host member CPU Usage values, from most to least critical. • Expired: Host members that are in an expired state. 														
Sort Descending	Toggle on to order the grid objects.														
Time Range	Choose a time range for the trend graph to trace, from 2 minutes to 7 days, or choose All Data to include all available data. The selected time range applies to all grid objects in the display.														
Grid Objects	Each grid object is a different host member and values describe each host member. <table> <tr> <td>Expired</td><td>When checked, performance data has not been received within the time specified by your administrator for the Expire Time. If your administrator has also set the Delete Time, this grid object will be deleted if no data is received within the time specified for deletion.</td></tr> <tr> <td>Host</td><td>The name of the host.</td></tr> <tr> <td>CPU</td><td>The number of CPUs and the number of CPUs in use.</td></tr> <tr> <td>CPU Speed</td><td>The processor speed.</td></tr> <tr> <td>Memory</td><td>The total amount of memory, in megabytes.</td></tr> <tr> <td>Virtual Memory</td><td>The total amount of virtual memory, in megabytes.</td></tr> <tr> <td>Trend Graph</td><td> <ul style="list-style-type: none"> • CPU Load Short Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes). • CPU Usage % Traces the percentage of CPU used. </td></tr> </table>	Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this grid object will be deleted if no data is received within the time specified for deletion.	Host	The name of the host.	CPU	The number of CPUs and the number of CPUs in use.	CPU Speed	The processor speed.	Memory	The total amount of memory, in megabytes.	Virtual Memory	The total amount of virtual memory, in megabytes.	Trend Graph	<ul style="list-style-type: none"> • CPU Load Short Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes). • CPU Usage % Traces the percentage of CPU used.
Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this grid object will be deleted if no data is received within the time specified for deletion.														
Host	The name of the host.														
CPU	The number of CPUs and the number of CPUs in use.														
CPU Speed	The processor speed.														
Memory	The total amount of memory, in megabytes.														
Virtual Memory	The total amount of virtual memory, in megabytes.														
Trend Graph	<ul style="list-style-type: none"> • CPU Load Short Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes). • CPU Usage % Traces the percentage of CPU used. 														

Member Summary






Use this display to study trends for a single host member, such as CPU usage and loads (Short, Medium and Long).

Select an **Instance** and a **Member** from the drop-down menus. Check status of the selected host member **TableSpace** allocations and current **Buffer Pool Hit Ratios**.

Choose a **Time Range** or click  to specify your own. Mouse-over to see additional details..



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 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.

Instance: Choose an instance.

Member: Choose a member.

Expired When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

Host Name The name of the selected host member.

OS The operating system on the selected host member.

Buffer Pool Hit Ratios

Each row in the table in a different buffer pool. Values describe values for the selected host member.

BP_NAME The name of the buffer pool.

Data

Index

Xda

TableSpace

Values describe TableSpace values for the selected host member.

TableSpace	The name of the TableSpace.
% Utilization	The percentage of TableSpace used.
Allocated	The amount of TableSpace allocated.
Free(MB)	The amount of free TableSpace, in megabytes.
Auto Extend	Indicates whether auto extend is enabled. Y/N


Trend Graph


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




- **CPU Load Short** Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
- **CPU Load Medium** Traces the amount of CPU load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
- **CPU Load Long** Traces the amount of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
- **CPU Usage %** Traces the percentage of CPU used.

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

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

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. By default, the time range end point is the current time. To enter a specific time range, click the associated ellipsis button .



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

All DB2 Databases

Displays in this View are:

- **"All Databases Table"**: View a tabular list of performance metrics for all databases in an instance.
- **"All Databases Grid"**: View trend graphs of databases Commits/sec, Rollbacks/sec and Response Times.

All Databases Table

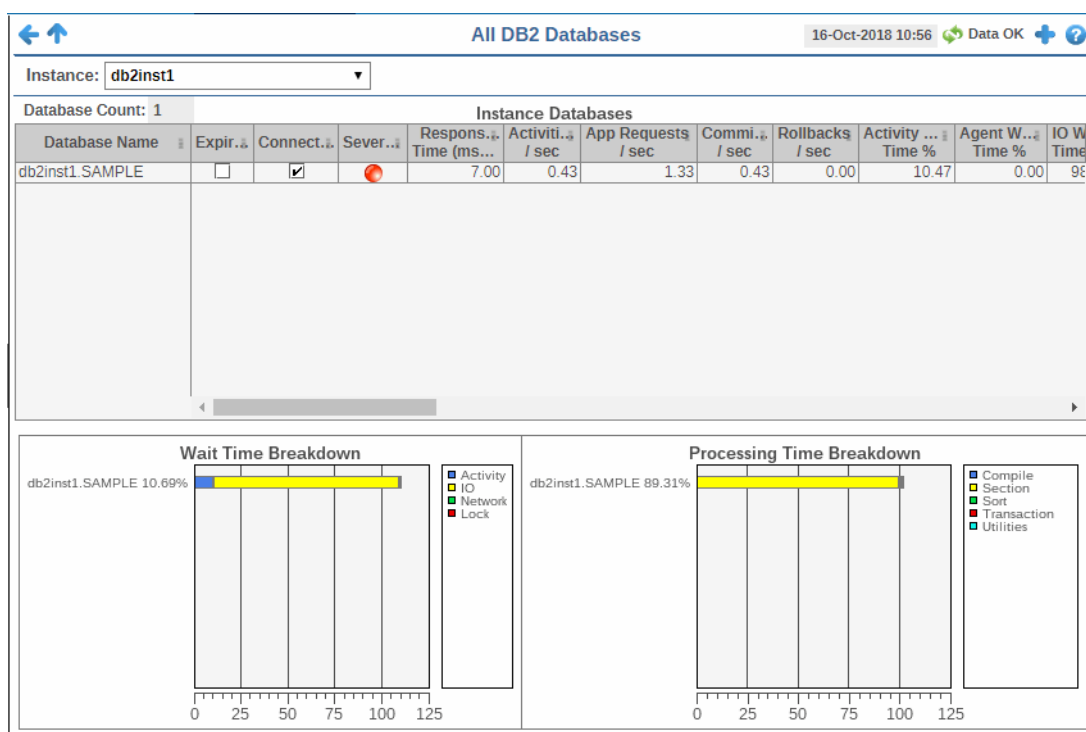
Select an **Instance** and view a list of all its databases and their performance metrics.

Use this display to identify which databases are having or causing issues for the instance.

Sort the list of databases by **Alert Severity** or by any other table column value.


In the bar graphs, view **Wait Time Breakdown** for the instance (by Activity, IO, Network and Lock), as well as **Processing Time Breakdown** (by Compile, Section, Sort, Transaction and Utilities) for the instance.

Click a row to drill-down to details in the **"Database Summary"** display.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.



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.

Instance Databases Table

Each row is a different database on the selected instance. Column values describe the database.

Database Count: The number of databases in the table.

Database	The name of the database.
Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this row will be deleted if no data is received within the time specified for deletion.
Connected	When checked, the database is connected.
Severity	The alert status:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Response Time	The response time, in milliseconds.
I/O Wait Time %	The percentage wait time taken by I/O operations.
Network Wait Time %	The percentage wait time taken by the network.
Agent Wait Time %	The percentage wait time taken by agents.
Avg Deadlocks per Activity	The average number of application deadlocks per activity.
Avg Lock Escalations per Activity	The average number of application deadlock escalations per activity.
Avg Lock Timeouts per Activity	The average number of application deadlock timeouts per activity.
Avg Lock Waits per Activity	The average number of application deadlock waits per activity.
Rows Read per Rows Returned	The number of rows read per number of rows returned.
Activites/sec	The number of activities per second.
App Requests/sec	The number of application requests per second.
Commits/sec	The number of application commits per second.
Rollbacks/sec	The number of application rollbacks per second.
Buffer Pool Hit Ratio %	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
Activity Wait Time %	The percentage wait time taken by activities.
Avg Request CPU Time	The average amount of CPU time used by requests, in seconds.
Compile Proc Time %	The percentage of time used for compiling processes.
Routine Time Request %	The percentage of time used for routine request processes.

Section Time %	The percentage of time used for section processes.
Section Sort Time %	The percentage of time used for sorting section processes.
BP Hit Ratio %	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
Transaction Time %	The percentage of time used for transaction processes.
Utils Proc Time %	The percentage of time used for utilities processes.
Activity Time (ms)	The amount of time, in milliseconds, used for activity processes.
Agent Idle Time (ms)	The amount of time, in milliseconds, that agents were in an idle state.
CPU Time (ms)	The amount of time, in milliseconds, used for CPU processes.
Processing Time (ms)	The amount of time, in milliseconds, used for processing requests.
Request Time (ms)	The amount of time, in milliseconds, used for processing requests.
DB Processing %	The percentage of time used for database processes.
DB Wait %	The percentage of time used for database waits.
# Connections	The current number of database connections.
Timestamp	The data and time this data was last updated.
Wait Time Breakdown	Shows the percentage of total wait time used for the instance categorized by: <ul style="list-style-type: none"> • Activity • IO • Network • Lock
Processing Time Breakdown	Shows the percentage of total processing time for the instance used by the following types of actions: <ul style="list-style-type: none"> • Compile • Section • Sort • Transaction • Utilities

All Databases Grid

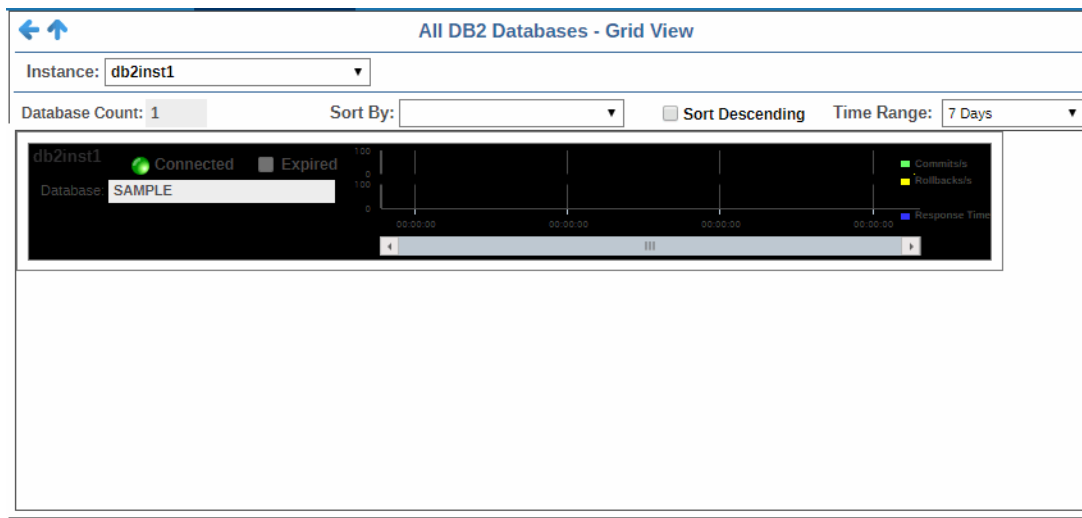
This display provides a grid view of a subset of data in the [“All Databases Table”](#) display.

Each grid object is a different database on the selected instance. Each has a trend graph that traces the **Commits per second**, **Rollbacks per second** and **Response Time** for the database. Choose a **Time Range** for the trend graph to trace, or select **All Data** to include all available data in the trace. Scroll forward and backward in the trend graph and click **Reset** to return to the original state.

Show **Expired** databases first using **Sort By: Expired**. Or sort by instance name, database name or connected state.

Toggle **Sort Descending** to order the grid objects.

Investigate performance metrics details for a database in the “Database Summary” display by clicking on the (left or right side of the) grid object.



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.
- 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.

- Instance:** Select an instance.
- Database Count:** The number of databases in the display.
- Sort By:** Orders the grid objects as follows:
- **Instance:** Instance names.
 - **Database:** Database names.
 - **Connected:** Databases that are connected.
 - **Expired:** Databases that are in an expired state.
- Sort Descending** Toggle on to order the grid objects.
- Time Range** Choose a time range for the trend graph to trace, from 2 minutes to 7 days, or choose **All Data** to include all available data. The selected time range applies to all grid objects in the display.
- Grid Objects** Each grid object is a different database and values describe each database.
- Connected** The database connection status:
- Red indicates that the database is disconnected.
 - Green indicates that the database is connected.
- Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this grid object will be deleted if no data is received within the time specified for deletion.

Database

The name of the database.

Trend Graph

- **Commits/s** Traces the number of commits per second.
- **Rollbacks/s** Traces the number of rollbacks per second.
- **Response Time** Traces the amount of time to respond.

Single Database

Displays in this View are:

- "Database Summary":
- "Partition Table":
- "Partition Heatmap"
- "Partition Grid"
- "Reports"
- "Trends"
- "DB2 Log"

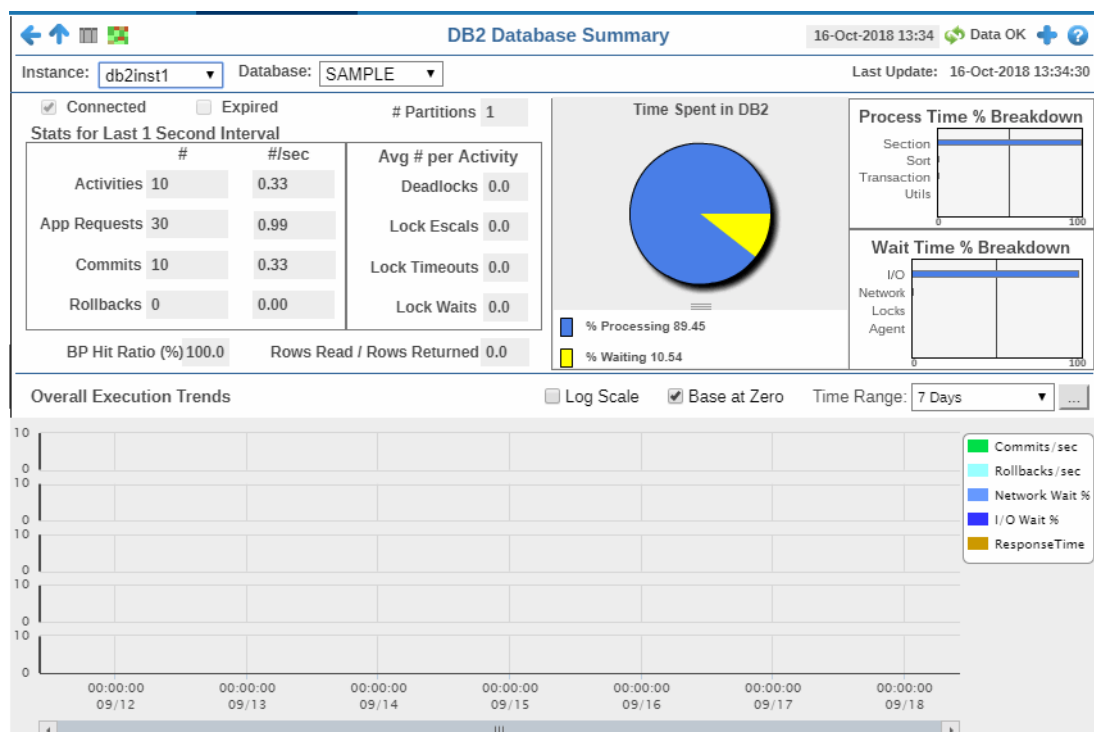
Database Summary


Use this display to investigate the performance and health of a database.

Select an **Instance** and a **Database**. Check main performance statistics such as database load, **Commits per second**, **Processing Time Breakdown** and **Wait Time Breakdown**

View trend graphs tracing **Commits** and **Rollbacks** per second, **Wait** and **Response** times, among others.

Choose a **Time Range** or click to specify your own.





Instance: Choose an instance.

Database: Choose a database.

Connected When checked, the database is connected.

Expired When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

Partitions The number of partitions on the database.

Stats for Last 1 Second Interval

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

Avg # per Activity

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

Time Spent in DB2

% Processing

% Waiting

Process Time % Breakdown

Shows the percentage of total processing time for the instance used by the following types of actions:

- Section
- Sort
- Transaction
- Utilities

Wait Time % Breakdown

Shows the percentage of total wait time used for the instance categorized by:

- Activity
- IO
- Network
- Locks

BP Hit Ratio %

The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.

Rows Read / Rows Returned

The number of rows read per number of rows returned.

Overall Execution Trends

For the selected database, trend graphs trace as follows:

- **Commits/sec** Traces the number of commits per second.
- **Rollbacks/sec** Traces the number of rollbacks per second.
- **Network Wait %** Traces the percentage of wait time used by network operations.
- **I/O Wait %** Traces the percentage of wait time used by I/O operations.
- **Response Time** Traces the total amount of wait time.


Log Scale


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




Base at Zero

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

Time Range

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



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

Partition Table

View a list of all partitions on a database, performance metrics for each partition, as well as setup details such as the product associated with the partition and service level.

Use this display to investigate partitioning issues on a database. Select an **Instance** and a **Database**. Sort the list by column values such as **DB2 Status**, **Rollbacks** per second and many others.

DB2 Database Partition Table

18-Oct-2018 10:06

Data OK


Instance: db2inst1

Database: SAMPLE

Last Update: 18-Oct-2018 10:06:30

Database Partitions

Partition #	Expir...	DB2 Status	DB Status	DB Start Time	Local Cons	Remote Cons	App C... Cons	SQL Fa... / sec	Rows Read / sec	Rows Selected / sec	Rows Chan. / sec
0	<input type="checkbox"/>	ACTIVE	ACTIVE	2018/05/30 10:16:29 PDT	0	11	4	0.00	3.46	4.08	0.0



Database Partitions Table

Each row is a different partition on the selected database. Column values describe the partition.

Instance: Select an instance.

Database: Select a database.

Partition #	The partition number.
Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this row will be deleted if no data is received within the time specified for deletion.
DB2 Status	The current DB2 status. For example, ACTIVE.
DB Status	The current database status. For example, ACTIVE.
DB Start Time	The date and time the database was last started.
Local Cons	The number of local connections.
Remote Cons	The number of remote connections.
App Cur Cons	The number of currently connected applications.
SQL Faults/sec	The number of SQL faults per second.
Rows Read /sec	The number of rows read per second.
Rows Selected /sec	The number of rows selected per second.
Rows Changed /sec	The number of rows changed per second.
SQL Selects /sec	The number of SQL selects per second.
Commits /sec	The number of commits per second.
Rollbacks /sec	The number of rollbacks per second.
Update/Del/Ins/Stmts / sec	The number of updates, deletions, insertions and statements per second.
Avg Sort Time/Transaction	The average amount of time for sorting transactions.
Product Name	The name of the product.
Service Level	The service level for the product.
Pool Data Hit Ratio	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.

Pool TmpData Hit Ratio	Refer to vendor documentation for details.
Pool TmpIndex Hit Ratio	Refer to vendor documentation for details.
Pkg Cache Inserts/K-Trans	Refer to vendor documentation for details.
Lock Wait Time/K-Trans	Refer to vendor documentation for details.
Dirty Steal Triggers/K-Trans	Refer to vendor documentation for details.
Deadlocks & Lock Timeouts/K-Tran	Refer to vendor documentation for details.
Avg Log Write Time / Trans	Refer to vendor documentation for details.
% Agent Usage	The percentage used by agents.
Agents Registered (Top)	The number of registered agents.
Rows Read Returned Ratio	The number of rows read per number of rows returned.
Select %	Refer to vendor documentation for details.
Phys Buffer Pool Read Ratio	Refer to vendor documentation for details.
Phys Buffer Pool Write Ratio	Refer to vendor documentation for details.

Partition Heatmap




View current alert status and performance metrics of all partitions on a DB2 database.

Answer questions such as, Are any partitions on this database reaching a state of critical health? Do I need to allocate more tablespace to any partitions? Is processing load and number of connections evenly distributed across partitions?

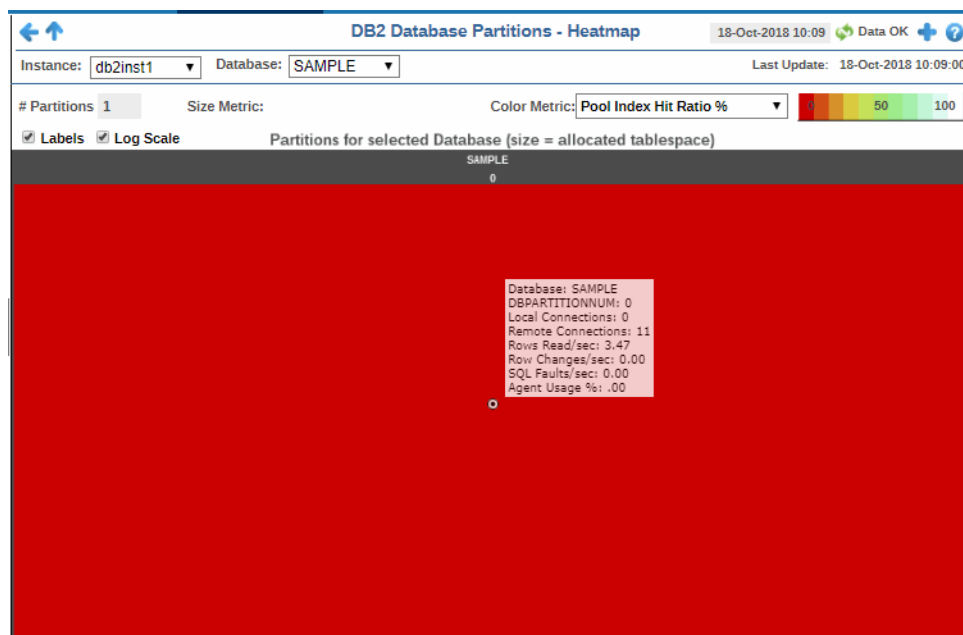
Each rectangle in the heatmap represents a different partition, where the rectangle color indicates the most critical alert state for items associated with that partition, and the rectangle size represents the tablespace allocation size for the partition.

Select an **Instance** and a **Database**. Use the **Metric** drop-down menu to view **Rows Read Per Second**, **Rollbacks Per Second** and **Dirty Steel Triggers/K-Trans**, among many others.

Each metric has its own color gradient bar legend that maps values to colors. By default, the **Commits/sec** metric is shown, which is the number of commits per second for the partition. Values range from **0** to the maximum number in the heatmap, as indicated in the color gradient bar:

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

Use the **Labels** check-box ☒ to include or exclude labels in the heatmap, use the **Log Scale** check-box ☒ to apply log scale and mouse over a rectangle to see additional metrics. Click a rectangle to see performance metrics for the database in which the partition resides in the “Database Summary” display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.


Fields and Data:


- Labels** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display. Each rectangle in the heatmap represents a different partition on the selected database. For additional details about the data, refer to vendor documentation.

Commits/sec

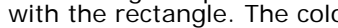
The number of application commits per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


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


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


 Green indicates that no metrics have exceeded their alert thresholds.

Pool Index Hit Ratio %


The average response time, in milliseconds, for items associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2ResponseTimeHigh** alert. The middle value in the gradient bar indicates the middle value of the range.


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


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

 Green indicates that no metrics have exceeded their alert thresholds.

SQL Failures/sec


The number of SQL faults per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


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


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

 Green indicates that no metrics have exceeded their alert thresholds.

Rows Read /sec


The number of rows read per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


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


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

 Green indicates that no metrics have exceeded their alert thresholds.

Rows Selected /sec

The number of rows selected per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

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

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

 Green indicates that no metrics have exceeded their alert thresholds.

Rows Changed /sec	<p>The number of rows changed per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
SQL Select Stmts/sec	<p>The number of SQL statements selected per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Rollbacks/sec	<p>The number of rollbacks per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
SQL Update/Del/Ins/Stmts /sec	<p>The number of SQL updates, deletions, insertions and statements per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.

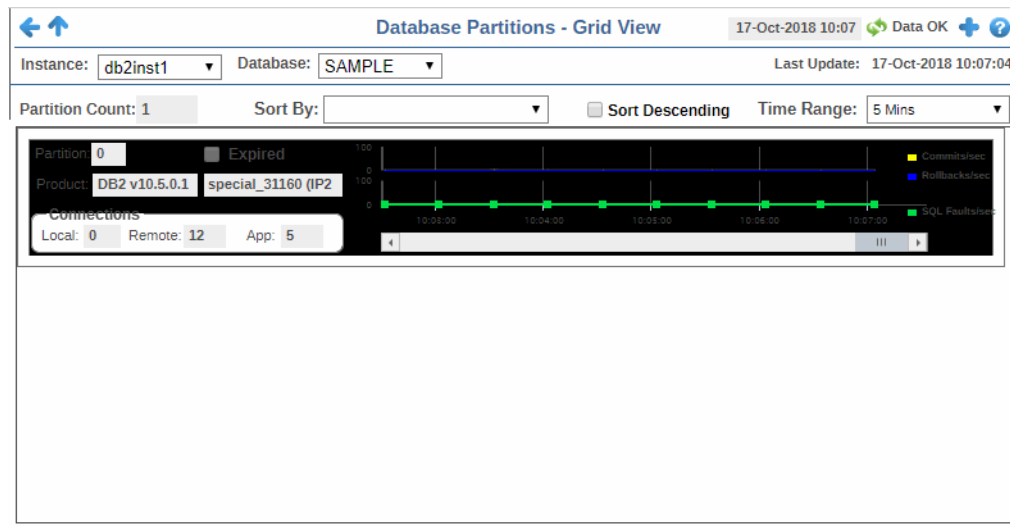
Partition Grid

This display provides a grid view of a subset of data in the [“Partition Table”](#) display.

Each grid object is a different partition on the selected database. Each has a trend graph that traces the **Commits per second**, **Rollbacks per second** and **SQL Faults per second** for the partition. Choose a **Time Range** for the trend graph to trace, or select **All Data** to include all available data in the trace. Scroll forward and backward in the trend graph and click **Reset** to return to the original state.


Show **Expired** databases first using **Sort By: Expired**. Or sort by partition name. Toggle **Sort Descending** to order the grid objects.

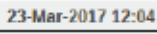
Investigate performance metrics details for a database in the “Database Summary” display by clicking on the (left or right side of the) grid object.



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.

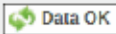
 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **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.

Instance:	Select an instance.						
Database:	Select a database.						
Partition Count:	The number of partitions in the display.						
Sort By:	Orders the grid objects as follows: <ul style="list-style-type: none"> • Partition #: Partition number. • Expired: Partitions that are in an expired state. 						
Sort Descending	Toggle on to order the grid objects.						
Time Range	Choose a time range for the trend graph to trace, from 2 minutes to 7 days, or choose All Data to include all available data. The selected time range applies to all grid objects in the display.						
Grid Objects	Each grid object is a different partition and values describe the partition. <table border="0"> <tr> <td>Partition</td><td>The partition number.</td></tr> <tr> <td>Expired</td><td>When checked, performance data has not been received within the time specified by your administrator for the Expire Time. If your administrator has also set the Delete Time, this grid object will be deleted if no data is received within the time specified for deletion.</td></tr> <tr> <td>Product</td><td>The name of the software and version.</td></tr> </table>	Partition	The partition number.	Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this grid object will be deleted if no data is received within the time specified for deletion.	Product	The name of the software and version.
Partition	The partition number.						
Expired	When checked, performance data has not been received within the time specified by your administrator for the Expire Time . If your administrator has also set the Delete Time , this grid object will be deleted if no data is received within the time specified for deletion.						
Product	The name of the software and version.						

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.

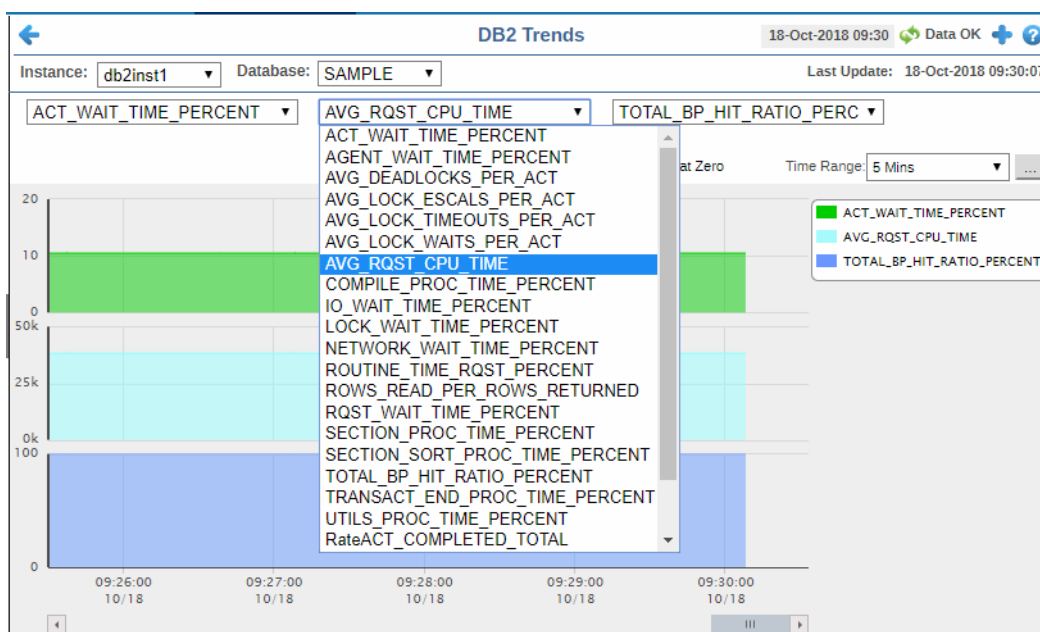
 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

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.

Trends


View IBM DB2 database performance and utilization metrics in a trend graph. Select an **Instance**, a **Database** and up to three IBM DB2 metrics. Metric options include (among others) **ACT_WAIT_CPU_TIME**, **AVG_LOCK_TIMEOUTS**, **ROWS_READ_PER_ROWS_RETURNED** and **UTILS_PROC_TIME_PERCENT**.

Choose a **Time Range** or click [...] to specify your own. Mouse-over to see additional details.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.

 **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.


Log Scale


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




Base at Zero

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

Time Range

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

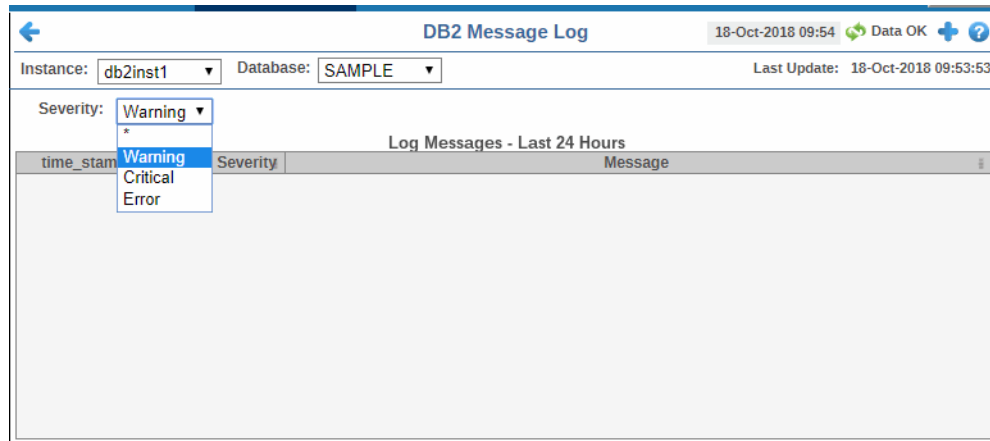


To change the time range click the Open Calendar button , choose the date and time, then click **OK**. Or, enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss** (for example, Aug 21, 2011 12:24 PM) and click **Apply**. Use the Navigation Arrows   to move forward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

DB2 Log


Read database log messages generated in the previous twenty-four hours. Select an **Instance** and a **Database**. Filter the list of log messages using the **Severity** drop-down menu:


- **Warning**: Lists only log messages with a warning severity level.
- **Critical**: Lists only log messages with an alarm severity level.
- **Error**: Lists only log message errors.
- ***** (asterisk): Lists all log messages.



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.

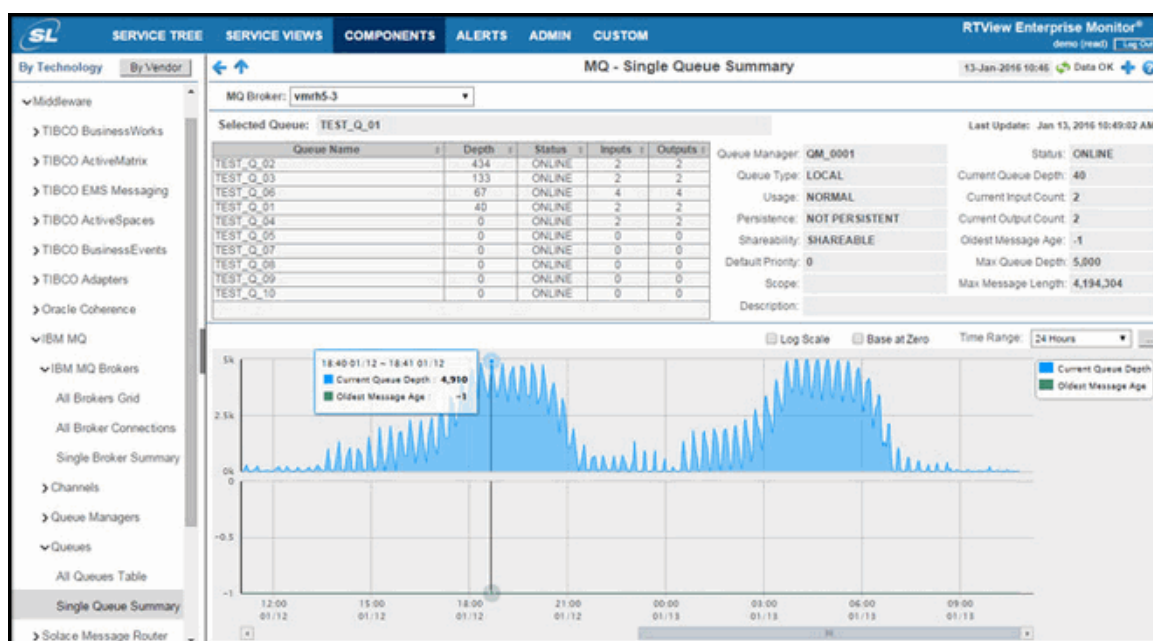
 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 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.

IBM MQ

Gain real-time visibility into the health and performance of IBM® MQ objects including brokers, queues, channels and queue managers. The following IBM MQ Monitor Views can be found under **Components** tab > **Middleware** > **IBM MQ**:

- **"IBM MQ Brokers View"**: The displays in this View present performance and utilization metrics for your IBM MQ Brokers.
- **"Channels View"**: The displays in this View present performance and utilization metrics for your IBM MQ Channels.
- **"Queue Managers"**: The displays in this View present performance and utilization metrics for your IBM MQ Queue Managers.
- **"Queues"**: The displays in this View present performance and utilization metrics for your IBM MQ Queues.



Overview

RTView Enterprise Monitor® and the Solution Package for IBM® MQ provide out of box performance and availability monitoring for support teams and IBM MQ administrators. Configuration options enable both consolidated views across the enterprise or views configured for specific support teams. As part of an end to end monitoring solution, users can view IBM MQ performance in the context of an application or service. This provides visibility into how IBM MQ performance is impacting adjacent technologies and the resulting business impact. Typical installations of RTView Enterprise Monitor and its solution packages take only a few hours, while developing custom views for a variety of IT and development roles can be achieved in just days.

Key Features

- Monitor real-time performance for early warning
- Analyze historical performance to differentiate trends and spikes
- Out of the box discovery and monitoring of key metrics
- Powerful diagnostics and correlations for complex performance analysis
- View IBM MQ in an application context for Application Support teams and Operations
- Minimal training, highly configurable by business and technical users

Metrics for IBM MQ

■ All MQ Brokers:

Queue Manager Status, Number of Channels, Number of Queues, Total Queue Depth

■ MQ Broker Summary:

Total Queue Depth

Overall Health State: Queue Manager, Channel, Queue Depth High, Queue Full

■ MQ Broker Connections:

Connection Status, Alert Status, Channel, Model Queue Name, Max Retries, Retry Interval, Wait Interval, Connection

■ All Queues:

Queue Manager, Queue type, Status, Alert State, Outputs, Inputs, Depth, Max Depth, Persistence Settings, Description, Max Message Length, Host, Default Priority, Get Messages, Put Messages, Scope, Shareability, Usage, Connection

■ Prebuilt Displays:

All Brokers Grid, Single Broker Summary, All Brokers detail table

All Channels table, Single Channel Summary, Single Channel Detail

All Queue Managers Detail table

All Queues Table, Single Queue Summary

All Trend Graphs show Historical Data

End-to-End Context for IBM MQ

- Custom flow diagrams help visualize complex applications and IBM MQ's place in that architecture
- Provides an Intuitive View of How IBM MQ Interacts with other Enterprise PaaS Components
- Designed and Developed for Large Scale, Mission Critical Environments

IBM MQ Brokers View

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

Displays in this View are:

- **"All Brokers Grid"**: This display presents a high-level perspective of utilization metrics for each IBM MQ Broker.
- **"All Broker Connections"**: This display presents detailed connection metrics for each IBM MQ Broker.
- **"Single Broker Summary"**: This display presents performance metrics for a single IBM MQ Broker, as well as detailed metrics for its channels and queues.

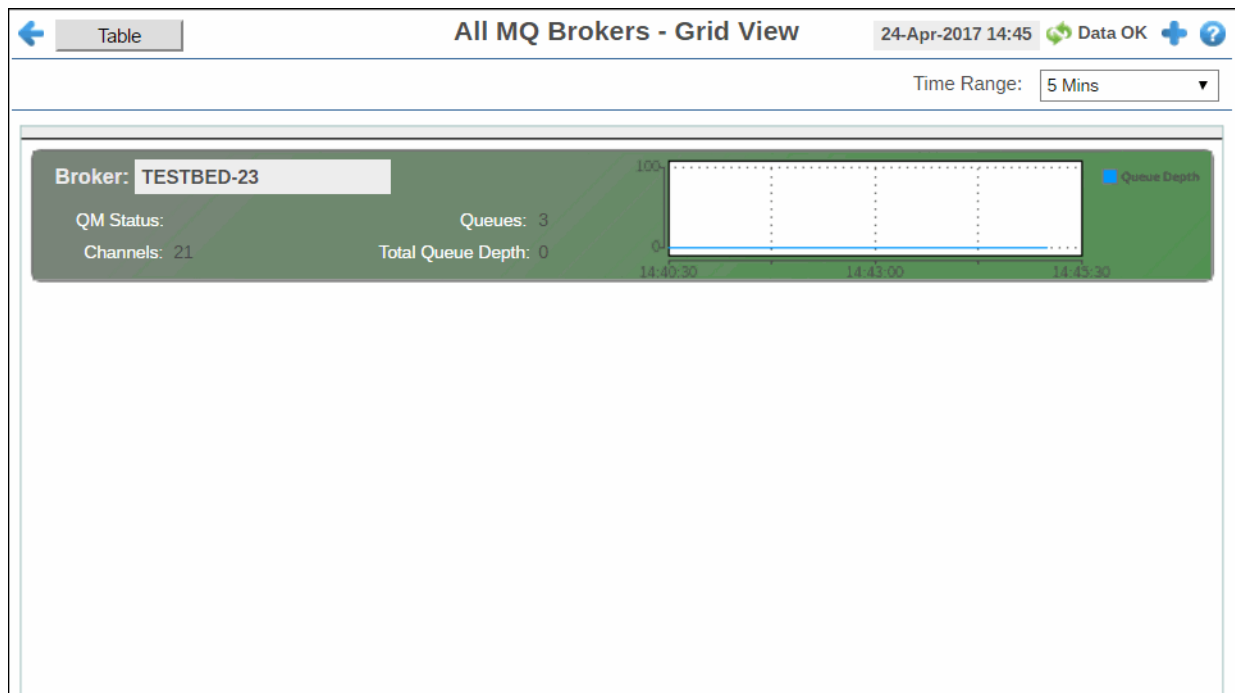
All Brokers Grid

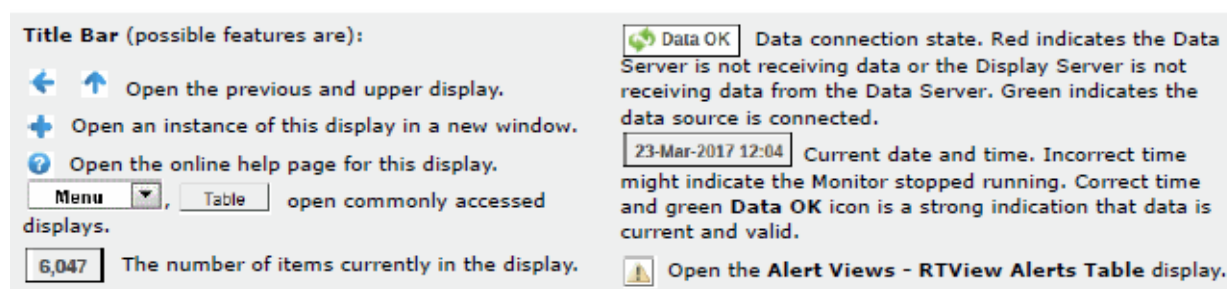
Track current and historical utilization and performance trends of all MQ Brokers in parallel. Use this display to quickly identify hosts with performance issues and verify whether the load is distributed evenly across brokers.

Each grid object is a different MQ Broker. Inactive brokers are shown in dark red, active brokers are shown in green. Metrics include QM status, queue depth and the number of channels per broker.

This display contains data obtained from IBM MQ. For example, MQIACH_HB_INTERVAL, MQIA_MONITORING_CHANNEL and MQIACH_MSG_COMPRESSION. For details, refer to vendor documentation.

Choose a time range to display from the drop-down menu and drill-down and investigate by clicking a broker to view details in the **"Single Broker Summary"** display.





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

Fields and Data:

Broker	The name of the broker.
QM Status:	The status of the Queue Manager on the broker.
Queues	The number of queues on the broker.
Channels	The number of channels on the broker.
Total Queue Depth	The total queue depth.
Trend Graph	Traces the queue depth on the broker.

All Broker Connections

View connection performance details for each MQ Broker in a tabular format, including current status, greatest alert severity and wait time to connect. Each row in the table is a different MQ Broker. Inactive brokers are shown in dark red.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the [“Single Broker Summary”](#) display.

Broker	Host	Port	Connected	Alert	Channel	Model Queue Name
TESTBED-23	192.168.200.73	1417			ADMIN.TEST.SVRCONN	SYSTEM.DEFAULT.MODEL.QUEUE

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.
- 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.

Table

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

- Broker** The name of the broker.
- Host** The name of the host.
- Port** The port number used.
- Connected** The current connection state:
 - Disconnected
 - Connected

Alert	<p>The current alert severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Channel	The name of the channel.
Model Queue Name	Named model queue of the connection.
Max Retries	Maximum number of subsequent connection retry attempts.
Retry Interval	Minimum interval (in seconds) between connection retry attempts.
Wait Interval	Wait interval (in seconds) between attempts to create a connection.
conn	The name of the connection.
time_stamp	The data and time of the last data update.

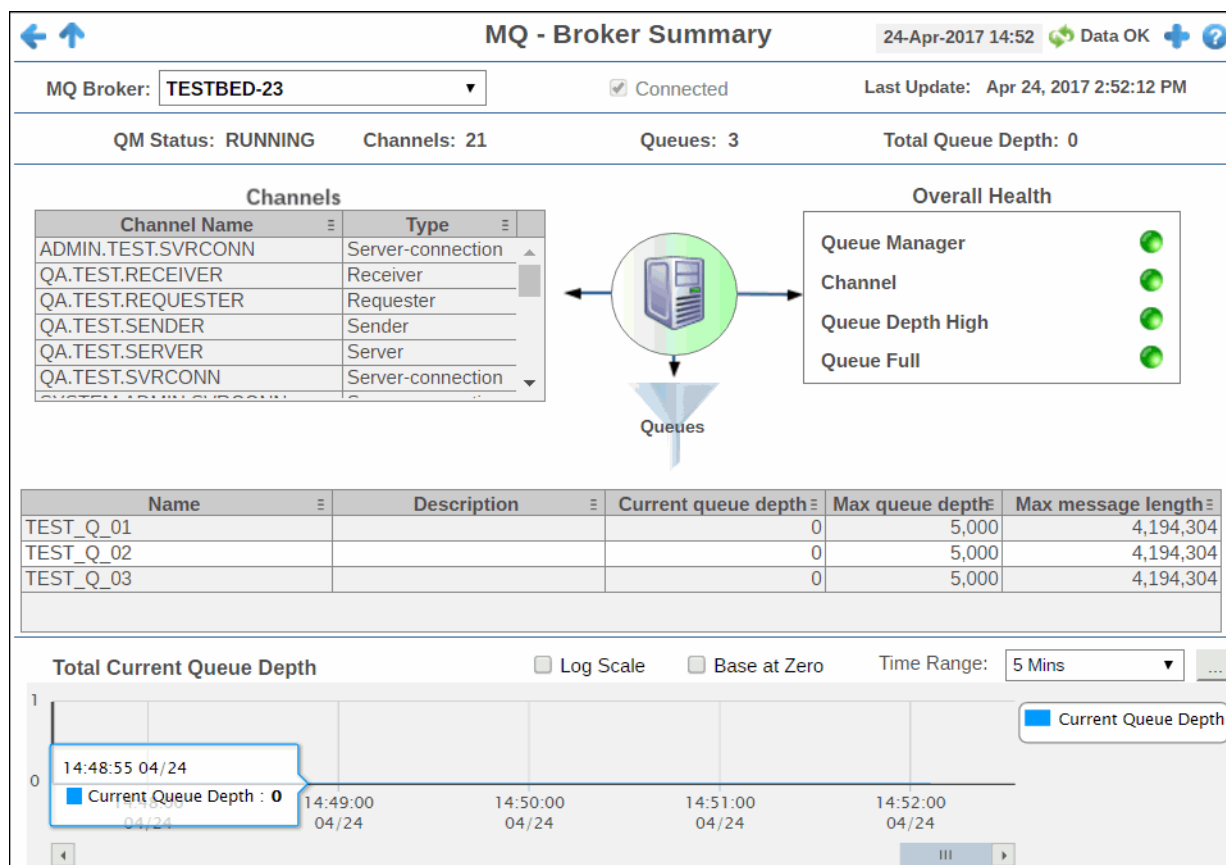
Single Broker Summary

View detailed utilization metrics and health for a single MQ Broker, such as Channel and Queue health and the number of channels per broker.

Use this display to investigate the performance and health of a broker.

Choose a broker from the drop-down menu. Click a row in the **Channels** table to investigate in the [“All Channels Table”](#) display. Click a row in the **Queues** table. to investigate in the [“All Queues Table”](#) display.

The trend graph traces the total **Current Queue Depth** for the selected broker. .



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.




Filter By:

MQ Broker Choose a broker to display.


Fields and Data


All values describe the selected broker except where noted.


- Connected** When checked, the broker is connected.
- Last Update** The data and time of the last data update.
- QM Status** The queue manager status. For example:
RUNNING - The QM is operating properly.



Channels	The current number of channels on the broker.
Queues	The current number of queues on the broker.
Total Queue Depth	The total depth of all queues combined.
Channels	Lists the channels on the broker. Channel Name: The name of the channel. Type: The type of channel (for example, Receiver, Cluster-sender, etc.)
Overall Health	The current alert severity for the broker Queue Manager , Channel , Queue Depth High and Queue Full :  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Queues	Each row is a different queue on the broker. Column values describe the queue. Name: The name of the queue. Description: A textual description of the queue. Current queue depth: The current queue depth. Max queue depth: The maximum queue depth. Max message length: The maximum message length in the queue.
Trend Graph	Traces the Queue Depth for the selected broker.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Base at Zero Select to use zero as the Y axis minimum for all graph traces.

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



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

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

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

Channels View

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

Displays in this View are:

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

All Channels Table

View detailed utilization metrics and parameter settings for MQ Channels on a broker. Metrics include total and delta counts for buffers received/sent. Parameter settings such as MQIACH_KEEP_ALIVE_INTERVAL, MQIACH_HDR_COMPRESSION and MQIACH_MAX_MSG_LENGTH are shown.

Each table row is a different channel. Inactive channels are shown in dark red, active channels are shown in green.

Use this display to quickly identify channels with performance issues and confirm channel configurations.

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

← ↑

MQ - All Channels Detail

24-Apr-2017 17:29 Data OK + ?

MQ Broker: TESTBED-23 ▾

MQ Channels for Selected Broker

Channel Name	Type	Alert	Buffers received	Delta	Buffers sent	Delta	Bytes received
ADMIN.TEST.SVRCONN	Server-connection	🟢	2,087,011	118	2,087,010	118	994,...
QA.TEST.RECEIVER	Receiver	🟢	0	0	0	0	
QA.TEST.REQUESTER	Requester	🟢	0	0	0	0	
QA.TEST.SENDER	Sender	🟢	0	0	0	0	
QA.TEST.SERVER	Server	🟢	0	0	0	0	
QA.TEST.SVRCONN	Server-connection	🟢	0	0	0	0	
SYSTEM.ADMIN.SVRCONN	Server-connection	🟢	0	0	0	0	
SYSTEM.AUTO.RECEIVER	Receiver	🟢	0	0	0	0	
SYSTEM.AUTO.SVRCONN	Server-connection	🟢	0	0	0	0	
SYSTEM.DEF.CLNTCONN	Client connection	🟢	0	0	0	0	
SYSTEM.DEF.CLUSRCVR	Cluster-receiver	🟢	0	0	0	0	
SYSTEM.DEF.CLUSDR	Cluster-sender	🟢	0	0	0	0	
SYSTEM.DEF.RECEIVER	Receiver	🟢	0	0	0	0	
SYSTEM.DEF.REQUESTER	Requester	🟢	0	0	0	0	
SYSTEM.DEF.SENDER	Sender	🟢	0	0	0	0	
SYSTEM.DEF.SERVER	Server	🟢	0	0	0	0	
SYSTEM.DEF.SVRCONN	Server-connection	🟢	0	0	0	0	
SYSTEM.FAKE.SVRCONN	Server-connection	🟢	0	0	0	0	
WALK.TEST.REQUESTER	Requester	🟢	0	0	0	0	
WALKER.ADMIN.SVRCONN	Server-connection	🟢	0	0	0	0	
WALKER.TEST.RECEIVER	Receiver	🟢	0	0	0	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.

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.

MQ Broker Choose a broker to display.

MQ Channels for Selected Broker Table
Each table row is a different channel on the selected broker. Column values describe the channel.

Channel Name The name of the channel.

Type The type of channel.

Alert	The current alert severity:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Buffers Received	The number of buffers received.
Delta	The number of buffers received since the last data update.
Buffers Sent	The number of buffers sent.
Delta	The number of buffers sent since the last data update.
Bytes Received	The number of bytes received.
Delta	The number of bytes received since the last data update.
Bytes Sent	The number of bytes sent.
Delta	The number of bytes sent since the last data update.
Batches Completed	The number of batches completed.
Delta	The number of batches completed since the last data update.
Description	A textual description of the channel.

Vendor Data

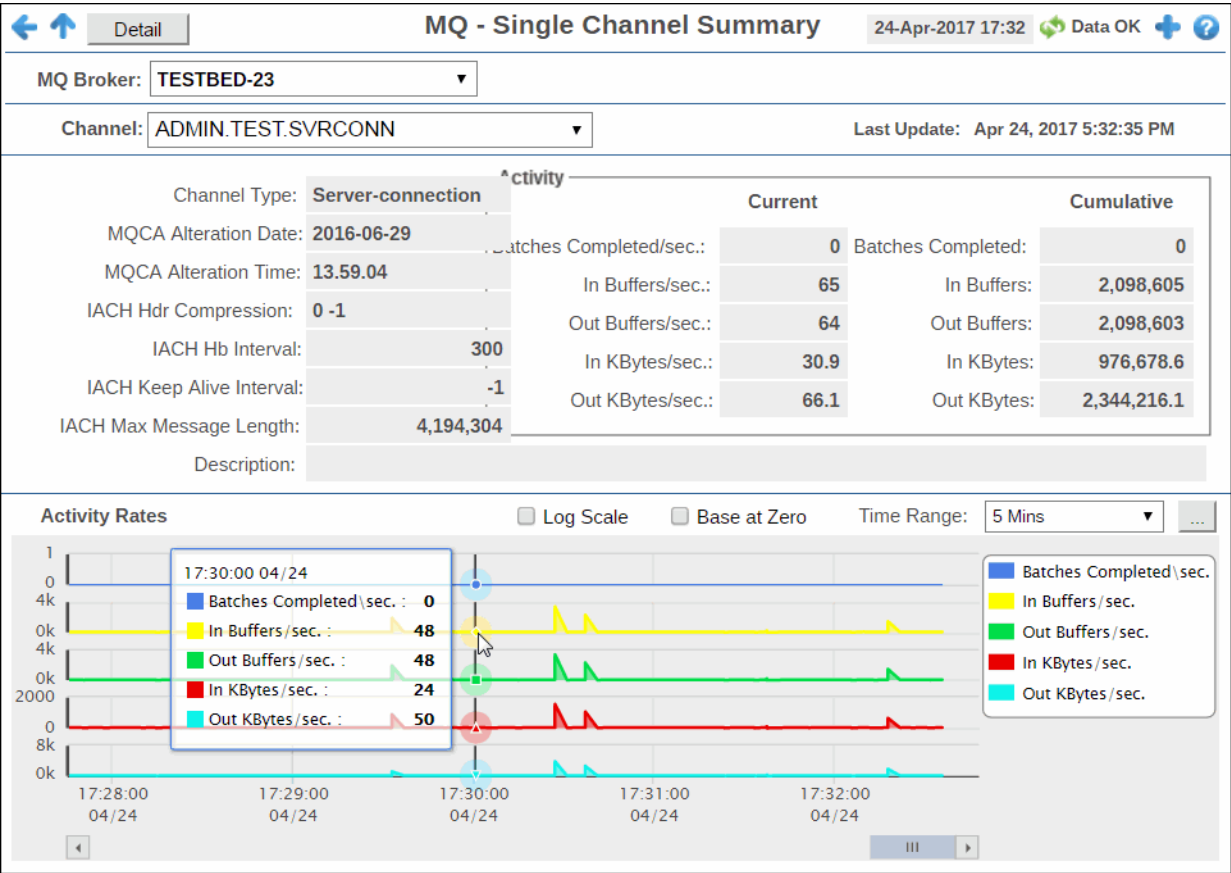
This display contains vendor data. Refer to vendor documentation for details.

Max message length	The maximum length of messages on the channel.
Status	The channel status.
Transmission queue name	The name of the queue that transmits for the channel.
Connection Name	The name of the connection.
time_stamp	The date and time of the last data update.
conn	The name of the connection.

Single Channel Summary

View current and historical activity rates and transmission settings for a single MQ channel. Metrics include buffers received/sent per second and batches completed. Parameter settings such as MQCA Alteration date, IACH Keep Alive Interval and Max Message Length are shown. Use this display to check the health of a channel and its configuration.

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



Title Bar (possible features are):

← ↑

Open the previous and upper display.

+

Open an instance of this display in a new window.

?

Open the online help page for this display.

Menu

Table

open commonly accessed displays.

6,047

The number of items currently in the display.

Data OK

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

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

- MQ Broker

Choose a broker to display.
- Channel

Choose a channel to display.
- Last Update

The date and time of the last data update.
- Channel Type

The type of channel.

Vendor Data

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

Activity Rates Trend Graph

Values describe the selected channel.

	Current	Cumulative
Batches Completed/sec	The current number of batches completed per second.	The total number of batches completed since the channel started.
In Buffers/sec	The current number of buffers received per second.	The total number of buffers received since the channel started.
Out Buffers/sec	The current number of buffers sent per second.	The total number of buffers sent since the channel started.
In KBytes/sec	The current number of kilobytes received per second.	The total number of kilobytes received since the channel started.
Out KBytes/sec	The current number of kilobytes sent per second.	The total number of kilobytes sent since the channel started.

Activity Rates

Traces the following for the selected channel:


- **Batches Completed/sec**: The current number batches completed per second.
- **In Buffers/sec**: The current number of buffers received per second.
- **Out Buffers/sec**: The current number of buffers sent per second.
- **In KBytes/sec**: The current number of kilobytes received per second.
- **Out KBytes/sec**: The current number of kilobytes sent per second.


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


Base at Zero



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

Time Range

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



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

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

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

Single Channel Detail

Get detailed transmission specifications and settings for a single MQ channel. Parameter settings such as CACH Rcv Exit User, IACH Msg Compression and CACH SSL Cipher Spec, as well as DataMQCA Alteration date, IACH Keep Alive Interval and Max Message Length are shown.

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

Summary		MQ - Single Channel Detail		24-Apr-2017 17:34 Data OK	
MQ Broker:	TESTBED-23				
Channel:	ADMIN.TEST.SVRCONN	Last Update: Apr 24, 2017 5:34:14 PM			
Channel Type: Server-connection		MQCA Alteration Date: 2016-06-29			
		MQCA Alteration Time: 13.59.04			
Description:					
IACH Hdr Compression:	0 -1				
IACH Msg Compression:	0 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1				
IACH Hb Interval:	300	IACH MCA User Id:		m	
IACH Keep Alive Interval:	-1	IA Monitoring Channel:		-3	
IACH Max Message Length:	4,194,304				
CACH Rcv Exit User Data:		CACH Rcv Exit Name:			
CACH Sec Exit User Data:		CACH Sec Exit Name:			
CACH Send Exit User Data:		CACH Send Exit Name:			
IACH SSL Client Auth:	0	CACH SSL Cipher Spec:			
IACH Xmit Protocol Type:	2	CACH SSL Peer Name:			

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.

- MQ Broker** Choose a broker to display.
- Channel** Choose a channel to display.
- Last Update** The date and time of the last data update.

Channel Type	The type of channel. For example, Server-connection .
MQCA Alteration Date	The date the MQ CA was last modified.
MQCA Alteration Time	The time the MQ CA was last modified.
Description	The description of the channel definition.
IACH Hdr Compression	The ACH header data compression techniques supported by the channel.
IACH Msg Compression	The ACH message data compression techniques supported by the channel.
IACH Hb Interval	The ACH heartbeat interval setting.
IACH Keep Alive Interval	The ACH keep alive interval setting (the timeout value for the channel).
IACH Max Message Length	The ACH maximum message length setting.
CACH Rcv Exit User Data	The user data that is passed to the receive exit.
CACH Sec Exit User Data	The user data that is passed to the security exit.
CACH Send Exit User Data	The user data that is passed to the send exit.
IACH SSL Client Auth	Denotes whether the channel needs to receive and authenticate an SSL certificate from an SSL client.
IACH Xmit Protocol Type	The transport type used.
IACH MCA User Id	The user ID used by the MCA when attempting to initiate a secure SNA session with a remote MCA.
IA Monitoring Channel	Denotes the attribute used to control the collection of online monitoring data.
CACH Rcv Exit Name	Denotes the name of the user exit program that was run by the channel receive user exit.
CACH Sec Exit Name	Denotes the name of the exit program that was run by the channel security exit.
CACH SSL Cipher Spec	Denotes the single CipherSpec for a TLS or SSL connection.
CACH SSL Peer Name	The Distinguished Name (DN) of the certificate from the peer queue manager or client at the other end of a IBM WebSphere MQ channel.

Queue Managers

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

Displays in this View are:

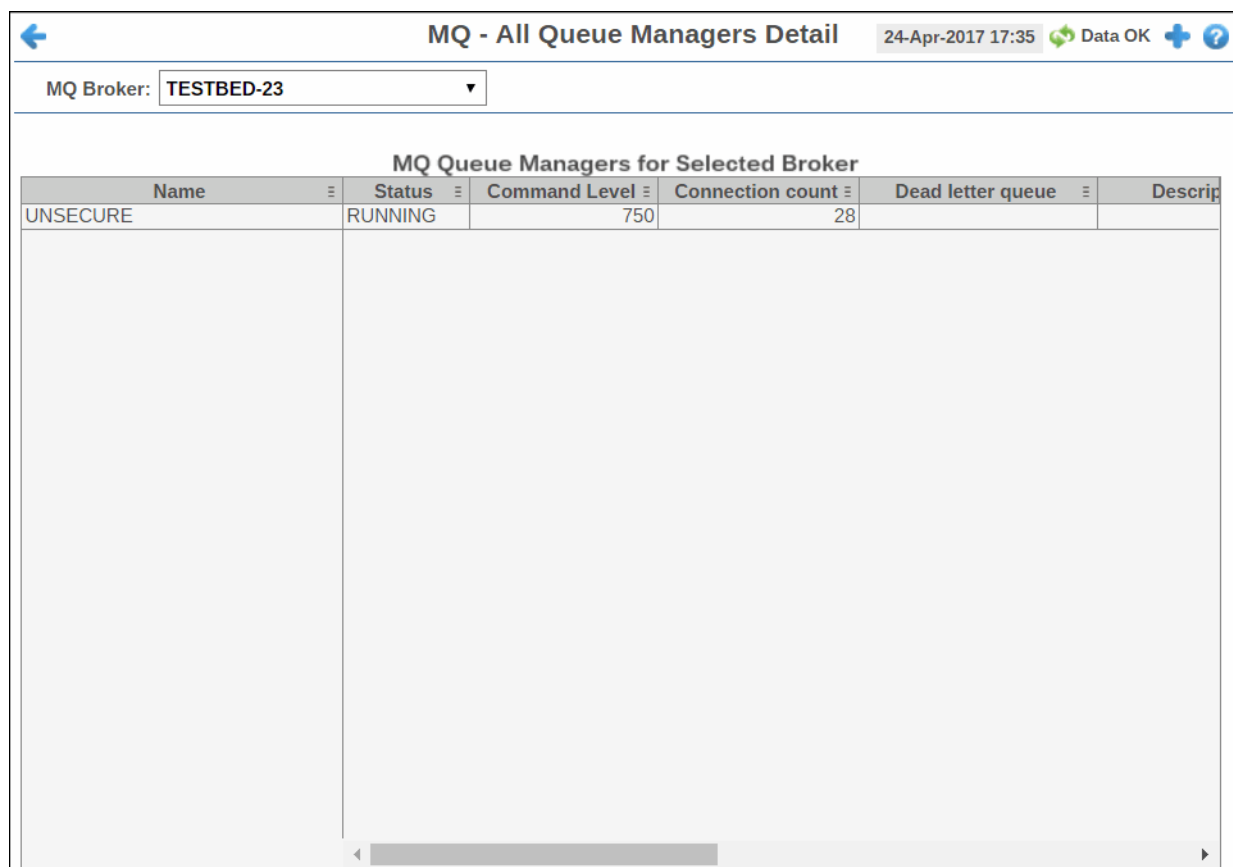
- [“All Queue Managers Table”](#): This display presents a high-level perspective of utilization metrics for each IBM MQ queue managers.

All Queue Managers Table

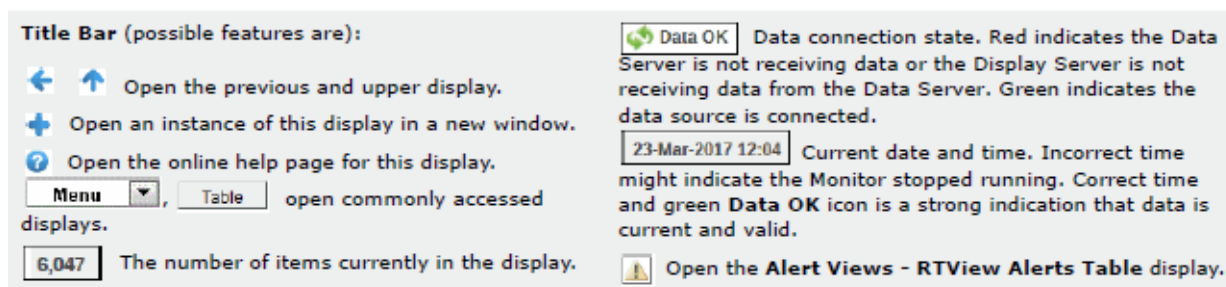
View detailed utilization metrics and parameter settings for all queue managers on a broker. Metrics include Connection Count and Max Message Length. Parameter settings such as Command Level are shown.

Each table row is a different queue manager. Use this display to quickly identify queue managers with performance issues and confirm configurations.

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



Name	Status	Command Level	Connection count	Dead letter queue	Description
UNSECURE	RUNNING	750	28		



MQ Broker Choose a broker to display.

MQ Queues Managers for Selected Broker Table

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

Name	The name of the queue manager.
Status	The queue manager status (for example, Running).
Command Level	The command level.
Connection Count	The number of connections on the queue manager.
Dead Letter Queue	The number of undelivered messages in the dead letter queue.
Description	A textual description of the queue manager.
Max Message Length	The maximum message length sent or received by the queue manager.
Max Priority	The queue manager rank in priority.
Platform	The queue manager platform type.
Host	The host name.
time_stamp	The date and time of the last data update.
Connection	The connection name.

Queues

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




Displays in this View are:

- **"All Queues Table"**: This display lists all IBM MQ queues with detailed performance metrics and configuration information.
- **"Single Queue Summary"**: This display presents detailed performance metrics and configuration information for a single IBM MQ queue.

MQ Broker Choose a broker to display.

MQ Queues for Selected Broker Table

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

Name	The name of the queue.
Queue Manager	The name of the queue manager.
Queue type	The type of queue.
Status	The queue status (for example, Online).
Alert	The current alert severity:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Outputs	The number of outgoing transactions.
Inputs	The number of incoming transactions.
Depth	The current queue depth.
Max Depth	The maximum number of messages allowed on the queue at any one time.
Persistence	Denotes whether or not the queue manager is persistent (PERSISTENT/NOT PERSISTENT).
Description	The description of the purpose of the queue.
Max Message Length	The maximum length of messages.
Host	The IP address of the host.
Default Priority	The default priority value for messages placed on the queue.
Get Messages	Denotes whether or not the queue enabled to get messages (GET ALLOWED/GET NOT ALLOWED).
Put Messages	Denotes whether or not the queue enabled to put messages (PUT ALLOWED/PUT NOT ALLOWED).
Scope	The defined scope setting for the queue.
Shareability	Denotes whether or not the queue is shareable (SHAREABLE/ NOT SHAREABLE).
Usage	The queue usage type (NORMAL/TRANSMISSION).
Connection	The name of the queue connection.
Expired	When checked, this connection is expired due to inactivity.
time_stamp	The date and time of the last data update.

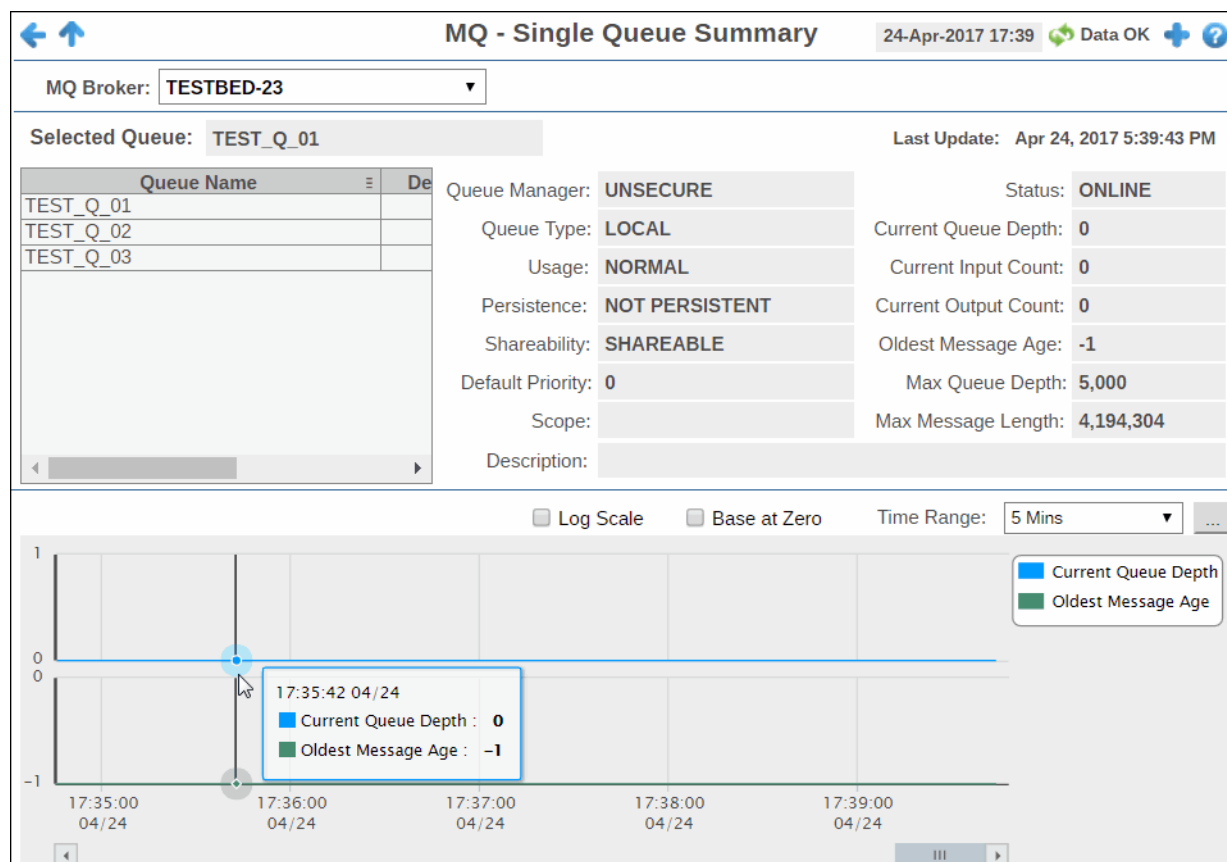
Single Queue Summary

View detailed performance metrics and settings for a single IBM MQ queue. Metrics include input/output counts, queue depth and max message length. Settings such as persistence, shareability, and scope are shown. Trend graph traces current queue depth and the oldest message age for the selected queue.

Choose a broker from the drop-down menu, then choose a queue from the **Queue** table.

Use this display to check the health of a queue and its configuration.

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



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

MQ Broker	Choose a broker to display.										
Selected Queue	Choose a queue from the table to populate the display. The name of the queue selected in the Queue table (below).										
Last Update	The data and time of the last data update.										
Queue Table	Choose a queue to populate the display.										
	<table> <tr> <td>Queue Name</td><td>The queue name.</td></tr> <tr> <td>Depth</td><td>The current queue depth.</td></tr> <tr> <td>Status</td><td>The current queue status.</td></tr> <tr> <td>Inputs</td><td>The current number of incoming transactions for the queue.</td></tr> <tr> <td>Outputs</td><td>The current number of outgoing transactions for the queue.</td></tr> </table>	Queue Name	The queue name.	Depth	The current queue depth.	Status	The current queue status.	Inputs	The current number of incoming transactions for the queue.	Outputs	The current number of outgoing transactions for the queue.
Queue Name	The queue name.										
Depth	The current queue depth.										
Status	The current queue status.										
Inputs	The current number of incoming transactions for the queue.										
Outputs	The current number of outgoing transactions for the queue.										
Queue Manager	The name of the queue manager for the queue.										
Queue Type	The type of queue.										
Usage	The queue usage type (NORMAL/TRANSMISSION).										
Persistence	Denotes whether or not the queue manager is persistent (PERSISTENT/NOT PERSISTENT).										
Shareability	Denotes whether or not the queue manager is shareable (SHAREABLE/NOT SHAREABLE).										
Default Priority	The default priority setting on the queue manager.										
Scope	The defined scope for the queue.										
Description	The description of the queue.										
Status	The status of the queue.										
Current Queue Depth	The current depth of the queue.										
Current Input Count	The number of incoming transactions.										
Current Output Count	The number of outgoing transactions.										
Oldest Message Age	The age of the oldest message in the queue.										
Max Queue Depth	The maximum number of messages allowed on the queue.										

Max Message Length

The maximum message length on the queue.

Trend Graph

Traces the following for the selected queue:

- **Current Queue Depth:** The current depth of the queue.
- **Oldest Message Age:** The age of the oldest message in the queue.


Log Scale

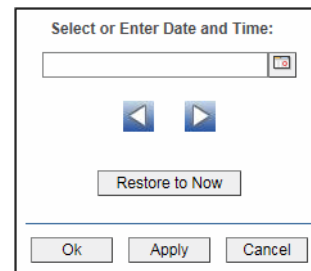
Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale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


Base at Zero

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

Time Range

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



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

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

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

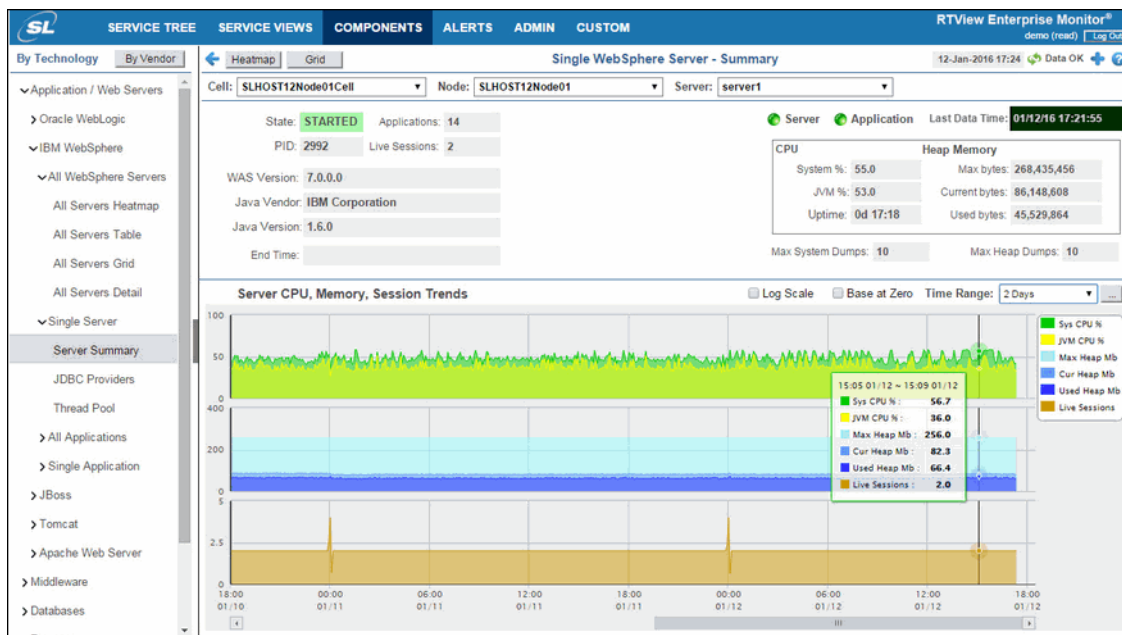
IBM WebSphere

Gain real-time visibility into the health and performance of WebSphere Application Server and deployed applications. The following WebSphere Monitor Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers**> **WebSphere**:

- **"All WebSphere Servers View"**: The displays in this View present high-level performance and utilization metrics for all of your IBM WebSphere servers.
- **"Single Server View"**: The displays in this View present high-level performance and utilization metrics for a single IBM WebSphere server.
- **"All Applications View"**: The displays in this View present high-level performance and utilization metrics for all of your web application sessions.
- **"Single Application View"**: The displays in this View present high-level performance and utilization metrics for a single web application.

Overview

RTView Enterprise Monitor and the Solution Package for IBM® WebSphere provide out of box performance and availability monitoring for support teams and WebSphere administrators. It enables users to ensure effective resource allocation by providing access to a wide variety of current and historical metrics.



Configuration options enable both consolidated views across the enterprise as well as views configured for specific support teams. As part of an end to end monitoring solution, users can view WebSphere performance in the context of an application or service. This provides visibility into how WebSphere performance is impacting adjacent technologies and the resulting business impact. Typical installations of RTView Enterprise Monitor and its solution packages take only a few hours, while developing custom views for a variety of IT and development roles can be achieved in just days.

Key Features

- Monitor real-time performance for early warning
- Analyze historical performance to differentiate trends and spikes
- Out of the box discovery and monitoring of key metrics and resources
- Ensure effective resource allocation
- Powerful diagnostics and correlations for complex performance analysis
- View WebSphere performance in an application context

Metrics for WebSphere Server

- System CPU Usage & Process CPU Usage
- Uptime
- Max Memory
- Heap Size & Max Heap Dumps on Disks
- Java Vendor & Version
- Used Memory, Free Memory & Used Memory Percent
- JVM Memory
- Max Heap, Current Heap, Used Heap
- Live Sessions
- JDBC Providers: Open Count, Created, Pool Size, Used Pool, Use Time
- Thread Pools: Pool Size, Active Count, Growable Indication, Inactivity Timeout & Max size
- Server Applications Sessions/Requests: Number of Sessions, Servlets, Total Requests, Current Requests, Avg. Response Time
- JSP Requests, JSP Response Time, Servlet Requests, Servlet Response Time, EJB Method Calls, EJB Response Time
- Component Detail
- Module Detail Totals for Charts & Tables

End-to-End Context for WebSphere

- Custom flow diagrams help visualize complex applications and WebSphere's place in that architecture
- Provides an Intuitive View of How WebSphere Interacts with other Enterprise PaaS Components
- Designed and Developed for Large Scale, Mission Critical Environments

All WebSphere Servers View

See performance and utilization metrics for all of your IBM WebSphere servers.

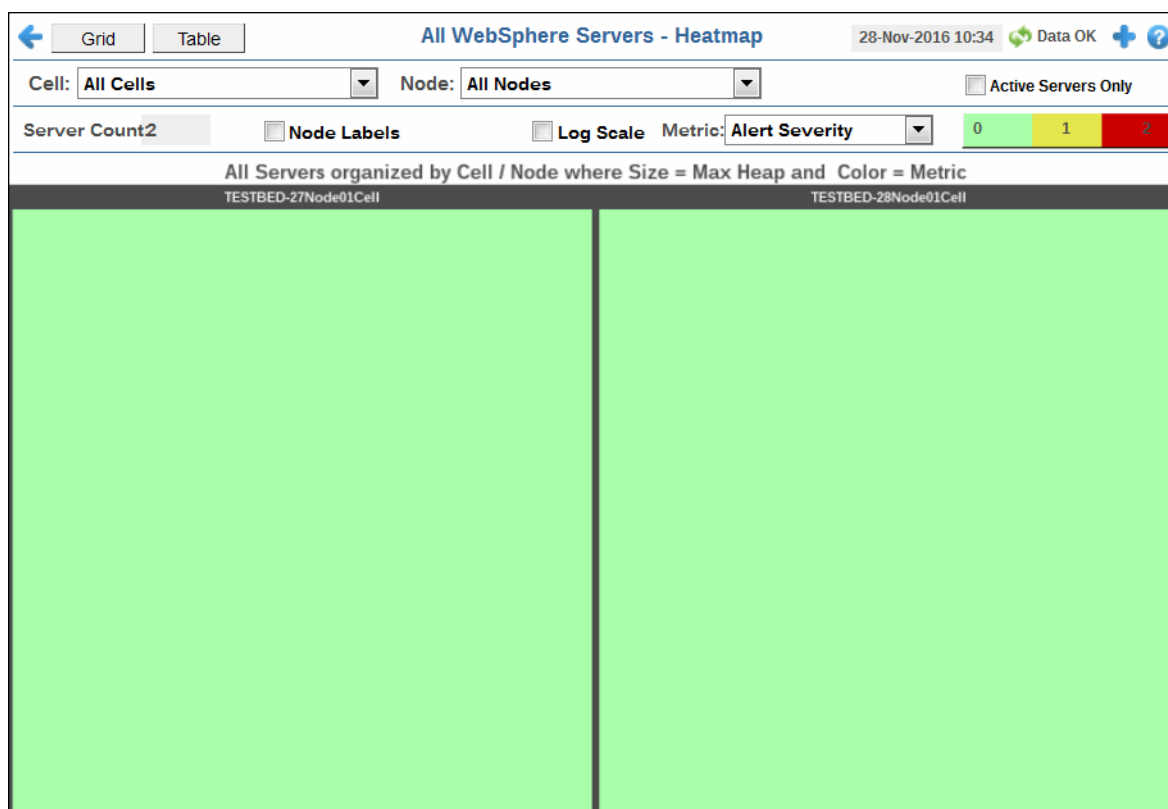
Displays in this View are:

- “All Servers Heatmap”
- “All Servers Table”
- “All Servers Grid”
- “All Servers Detail”

All Servers Heatmap

This heatmap shows the current status and utilization metrics for all WebSphere servers. Choose a cell and node from the drop-down menus. Use this display to see metrics for **Alert Count**, **Live Session Count**, **WAS CPU %**, **Host CPU %** and **Memory Used %**. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle is a different WebSphere server. Use the **Node Labels** check-box ☒ to include or exclude labels in the heatmap, and mouse over a rectangle to see additional metrics for a server. Click a rectangle to drill down to the “[Server Summary](#)” display, which shows additional details about the server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Filter By:

Cell: Choose a cell or **All Cells** to see metrics for.

Node: Choose a node or **All Nodes** to see metrics for.

Fields and Data:

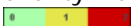
Active Servers Only Choose this check box to only include active servers in the display.


Server Count The number of servers in the display.


Node Labels Select to include node labels in the display.

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


Metric Choose a metric to view in the display.


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

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


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

 Green indicates that no metrics have exceeded their alert thresholds.


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

CPU Used% The percent CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **ProcessCpuLoadPercent**. The middle value in the gradient bar indicates the middle value of the range.

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

V Memory Used% The percent virtual memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SystemCpuLoadPercent**. The middle value in the gradient bar indicates the middle value of the range.

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

Free Memory The total amount of available memory. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount of available memory. The middle value in the gradient bar indicates the average amount.

The **Auto** flag does not have any impact on this metric.

All Servers Table

View WebSphere server utilization details, including memory, CPU and heap size. Choose a cell and node from the drop-down menus. Each row in the table is a different server. The row color for inactive servers is dark red.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the “[Server Summary](#)” display.

Cell	Node	Server	Expired	Alerts	State	PID	LiveCount	CPU
TESTBED-27Node01Cell	TESTBED-27Node01	server1	<input type="checkbox"/>		STARTED	3006	0	
TESTBED-28Node01Cell	TESTBED-28Node01	server1	<input type="checkbox"/>		STARTED	8033	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.
- 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.

Filter By:

Cell: Choose a cell or **All Cells** to see metrics for.

Node: Choose a node or **All Nodes** to see metrics for.

Fields and Data

This display includes:

Active Servers Only Select to only include active servers in the display.

Server Count The number of servers in the display.


Table

Each table row is a different server. Table column values describe the cell on the server.

Cell The name of the cell.

Node The name of the node.

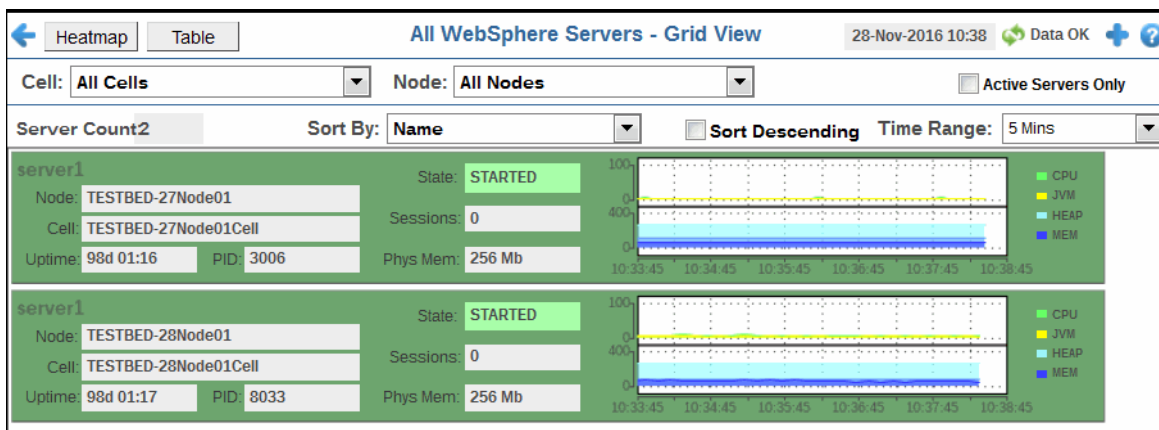
Server The name of the server.

Expired	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
Alert Level	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
State	The WebSphere server current state: <ul style="list-style-type: none"> • STARTED • STOPPED
PID	The WebSphere server process identifier.
LiveCount	The total number of currently active sessions.
CPUUsageSinceLastMeasurement	The amount of CPU usage, in megabytes, since the last data update.
CPUUsageSinceServerStarted	The amount of CPU usage, in megabytes, since the server was started.
ProcessCpuUsage	The amount of process CPU usage, in megabytes, since the server was started.
UpTime	The amount of time, in milliseconds, since the server was started.
maxMemory	The maximum amount of memory used since the server was started.
heapSize	The heap size, in kilobytes.
usedMemory	The amount of used memory, in kilobytes.
freeMemory	The amount of free memory, in kilobytes.
usedMemoryPercent	The amount of used memory, in percent.
HeapSize	The heap size, in megabytes.
UsedMemory	The amount of used memory, in megabytes.
FreeMemory	The amount of free memory, in megabytes.
maxHeapDumpsOnDisk	The maximum amount of heap dumps on disk that have been performed since the last restart.
maxSystemDumpsOnDisk	The maximum amount of system dumps on disk that have been performed since the last restart.
javaVendor	The name of the Java software vendor.
javaVersion	The Java software version.
TIME_STAMP	The date and time of the last data update.

All Servers Grid

View WebSphere server utilization details, including memory, CPU and heap size, in a grid format. Choose a cell and node from the drop-down menus.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the “[Server Summary](#)” display.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

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.


Filter By:

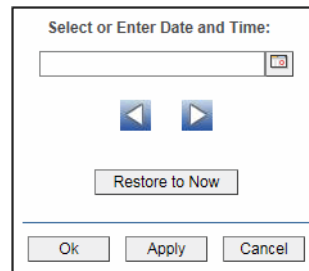
- Cell:** Choose a cell or **All Cells** to see metrics for.
- Node:** Choose a node or **All Nodes** to see metrics for.

Fields and Data


This display includes:



- Active Servers Only** Select to only include active servers in the display.
- Server Count** The number of servers in the display.
- Sort By:** Options are to sort servers in the grid by **Name**, **Live Sessions**, **JVM CPU %**, **Up Time** or **Max Memory**.
- Sort Descending** Select to organize display elements in descending order.

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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right) and a "Restore to Now" button. At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

Node	The name of the node.
Cell	The name of the cell.
UpTime	The amount of time, in milliseconds, since the server was started.
PID	The WebSphere server process identifier.
State	The WebSphere server current state: <ul style="list-style-type: none"> • STARTED • STOPPED
Sessions	The current number of sessions.
Phys Mem	The current amount of disk space, in megabytes.
Trend Chart	<ul style="list-style-type: none"> • CPU Traces the amount of server CPU utilization. • JVM Traces the amount of server CPU utilization. • HEAP Traces the amount of server heap memory utilization. • MEM Traces the amount of server memory utilization.

All Servers Detail

View detailed data for all your WebSphere servers in a tabular format, including system metrics, JVM runtime data and session totals.

Drill-down and investigate by clicking a row in the table to view details for a server in the “Server Summary” display.

← Grid All WebSphere Servers - Detail Tables 28-Nov-2016 10:39 Data OK + ?

Server Info						
node	Server	Expired	state	cell	cellName	cellShortName
TESTBED-27Node01	server1		STARTED	TESTBED-27Node01Cell	TESTBED-27Node01Cell	
TESTBED-28Node01	server1		STARTED	TESTBED-28Node01Cell	TESTBED-28Node01Cell	

System Metrics

node	Server	Expired	CPUUsageSinceLastMeasurement	CPUUsageSinceServerStarted	
TESTBED-27Node01	server1		4	0	TESTBED
TESTBED-28Node01	server1		7	0	TESTBED

JVM Runtime Data

node	Server	Expired	UpTime	ProcessCpuUsage	HeapSize	UsedMemory	heapSize	n
TESTBED-27Node01	server1		8471851	2	112064	65667	114753536	
TESTBED-28Node01	server1		8471913	4	79907	57899	81824768	

Session Totals

node	Server	Expired	LiveCount	TIME_STAMP
TESTBED-27Node01	server1		0	11/28/16 1
TESTBED-28Node01	server1		0	11/28/16 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 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.

Server Info Table

Each table row is a different server. Table column values describe the server.

Node	The name of the node.
Server	The name of the server.
Expired	Red indicates that the server is expired. Green indicates that the server is online.
State	The WebSphere server current state: <ul style="list-style-type: none"> • STARTED • STOPPED
cell	The name of the cell.
cellName	The full name of the cell.

cellShortName	The short name for the cell.
deployedObjects	A list of deployed objects on the server.
eventTypes	A list of events that occurred on the server.
internalClassAccessMode	Describes the internal class access mode. Refer to vendor documentation for details.
j2eeType	The J2EE type.
javaVMs	A list of Java virtual machines.
mbeanIdentifier	The MBean id.
name	Refer to vendor documentation for details.
nodeName	The full name of the node.
nodeShortName	The short name for the node.
pid	The WebSphere server process identifier.
platform	The platform type.
platformName	The name of the platform.
platformVersion	The software version on the platform.
processType	The platform type (e.g. UnManagedProcess).
resources	Refer to vendor documentation for details.
serverVendor	The name of the server vendor.
shortName	Refer to vendor documentation for details.
spec	Refer to vendor documentation for details.
statisticsProvider	Refer to vendor documentation for details.
threadMonitorAdjustmentThreshold	Describes thread monitor settings. Refer to vendor documentation for details.
threadMonitorInterval	Describes thread monitor settings. Refer to vendor documentation for details.
threadMonitorThreshold	Describes thread monitor settings. Refer to vendor documentation for details.
type	Refer to vendor documentation for details.
version	Refer to vendor documentation for details.
Connection	The name of the connection.
TIME_STAMP	The date and time of the last data update.

System Metrics Table



Each table row is a different server. Table column values describe the server.

Node	The name of the node.
Server	The name of the server.
Expired	 Red indicates that the server is expired.  Yellow indicates  Green indicates that the server is online.

CPUUsageSinceLastMeasurement	The amount of CPU usage, in megabytes, since the last data update.
CPUUsageSinceServerStarted	The amount of CPU usage, in megabytes, since the server was started.
cell	The name of the cell.
hasStats	Refer to vendor documentation for details.
mbeanIdentifier	The MBean id.
name	Refer to vendor documentation for details.
platform	Refer to vendor documentation for details.
spec	Refer to vendor documentation for details.
type	Refer to vendor documentation for details.
Connection	The name of the connection.
TIME_STAMP	The date and time of the last data update.

JVM Runtime Data Table




Each table row is a different server. Table column values describe the server.

Node	The name of the node.
Server	The name of the server.
Expired	 Red indicates that the server is expired.  Green indicates that the server is online.
UpTime	The amount of time, in milliseconds, since the server was started.
ProcessCpuUsage	The amount of process CPU usage, in megabytes, since the server was started.
HeapSize	The current heap size, in kilobytes.
UsedMemory	The current amount of memory used, in kilobytes.
heapSize	The current heap size, in kilobytes.
maxMemory	The maximum amount of memory used since the server was started.
freeMemory	The current amount of free memory, in kilobytes.
usedMemory	The current amount of used memory, in kilobytes.
J2EEServer	The name of the J2EE server.
cell	The name of the cell.
hasStats	Refer to vendor documentation for details.
j2eeType	The J2EE type.
Java Vendor	The name of the Java vendor.
Java Version	The Java software version.
maxHeapDumpsOnDisk	The maximum amount of heap dumps on disk that have been performed since the last restart.

maxSystemDumpsOnDisk	The maximum amount of system dumps on disk that have been performed since the last restart.
mbeanIdentifier	The MBean id.
statisticsProvider	Refer to vendor documentation for details.
version	Refer to vendor documentation for details.
Connection	The name of the connection.
spec	Refer to vendor documentation for details.
platform	Refer to vendor documentation for details.
TIME_STAMP	The date and time of the last data update.

Session Totals Table

Each table row is a different server. Table column values describe the server.

Node	The name of the node.
Server	The name of the server.
Expired	 Red indicates that the server is expired.  Yellow indicates  Green indicates that the server is online.
LiveCount	The total number of currently active sessions.
TIME_STAMP	The date and time of the last data update.

Single Server View

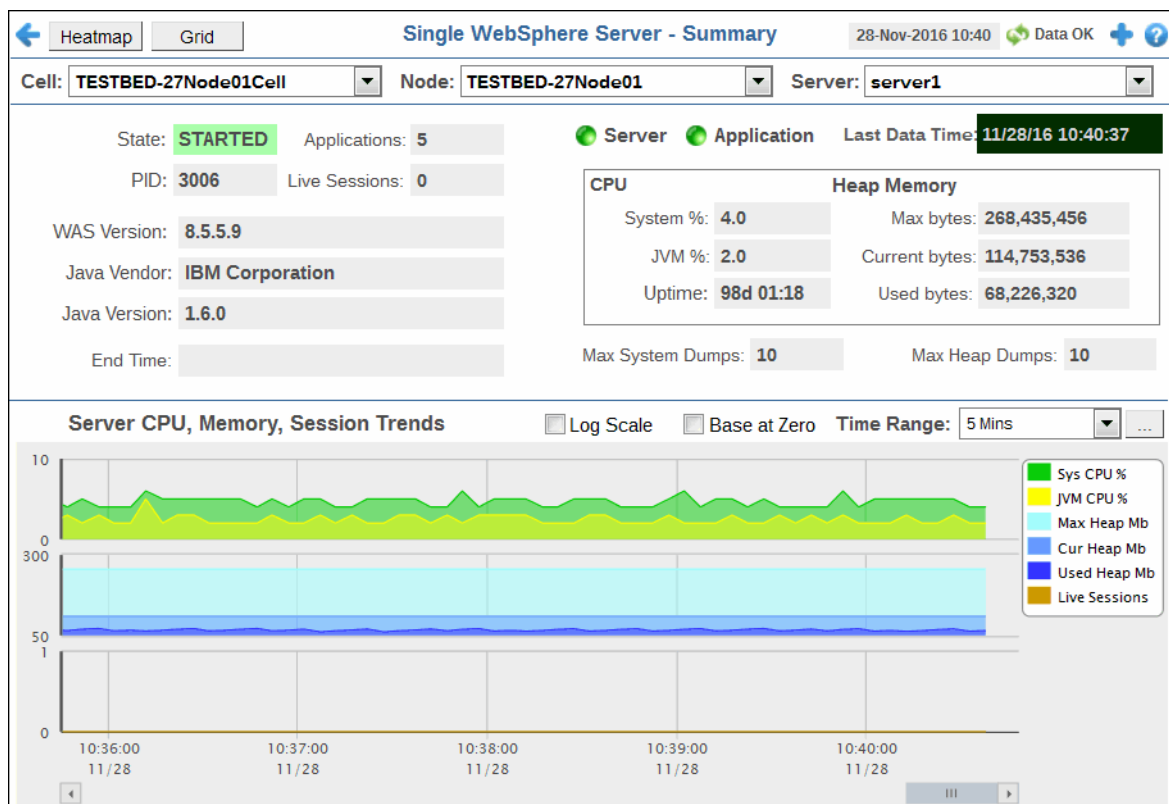
See performance and utilization metrics for a single IBM WebSphere server.

Displays in this View are:

- ["Server Summary"](#)
- ["JDBC Providers"](#)
- ["Thread Pool"](#)

Server Summary

Track current and historical performance of web applications on one server. Choose a cell, node and server from the drop-down menus. Mouse over the trend graph to see metrics.



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.
- 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.







Filter By:

- Cell:** Choose a cell or **All Cells** to see metrics for.
- Node:** Choose a node or **All Nodes** to see metrics for.
- Server:** Choose the server for which you want to show data.

Fields and Data

This display includes:

- State** The WebSphere server current state:
 - STARTED
 - STOPPED

Applications	The number of applications running on the server.
PID	The WebSphere server process identifier.
Live Sessions	The current number of active sessions.
WAS Version	The WebSphere Application Server software version.
Java Vendor	The Java vendor name.
Java Version	The Java software version.
End Time	Refer to vendor documentation for details.
Server	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Application	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Last Data Time	The date and time of the last data update.

CPU

System %	The current amount of CPU used by the system, in percent.
JVM %	The current amount of CPU used by the JVM, in percent.
Uptime	The number of days, hours and minutes since the server was last started.

Heap Memory

Max bytes	The maximum size of heap memory, in bytes.
Current bytes	The current size of heap memory, in bytes.
Used bytes	The size of heap memory being used, in bytes.

Max System Dumps The maximum number of system dumps that are allowed to be performed.

Max Heap Dumps The maximum number of heap dumps that are allowed to be performed.

Server CPU, Memory, Session Trends

The trend graph traces the following for the selected server:

- **Sys CPU%** The percent of system CPU used.
- **JVM CPU%** The percent of JVM CPU used.
- **Max Heap Mb** The maximum amount of heap memory ever used, in megabytes.
- **Cur Heap Mb** The current amount of heap memory available, in megabytes.
- **Used Heap Mb** The current amount of heap memory used, in megabytes.
- **Live Sessions** The current number of active sessions.


Log Scale

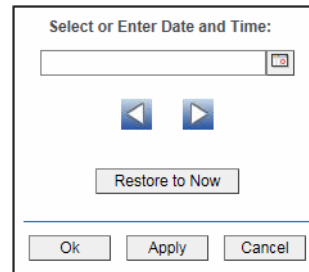
Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale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

Base at Zero


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



Time Range

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



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

JDBC Providers

Track current and historical performance of all JDBC Providers on a server. Choose a cell, node and server from the drop-down menus. Each table row is a different JDBC Provider. Select a row to populate the trend graph with JDBC Provider performance metrics. Mouse over the trend graph to see metrics.



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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

JDBC Providers Table

Each table row is a different JDBC Provider.

- JDBC Provider** The name of the JDBC Provider.

Open	The number of currently open connections.
Closed	The number of currently closed connections.
Created	The number of connections created since the server was last started.
Total	Refer to vendor documentation for details.
PoolSize	The number of connections in the pool.
Used	The number of used connections in the pool.
% Used	The percent of connections used in the pool.
Use Time	The average connection duration, in seconds.
Total	Refer to vendor documentation for details.
WaitTime	The average amount of time to establish a connection, in seconds.
WaitingThreadCount	The current number of threads waiting to establish a connection, in seconds.
description	Describes the JDBC provider.
TIME_STAMP	The date and time of the last data update.


JDBC Provider Trend Graph

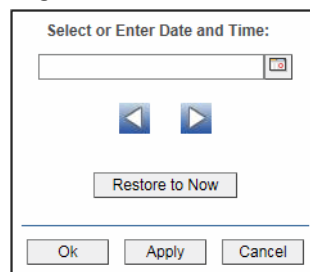
The trend graph traces the following for the JDBC Provider on the selected server:


- **Open Count** The number of currently open connections.
- **Created** The number of connections created.
- **Pool Size** The number of connections in the pool.
- **Used Pool** The number of connections in the pool that are being used.
- **Use Time** The average length of time of connections, in seconds.



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

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

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




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

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

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

Thread Pool

Track current and historical performance of thread pools for web applications on one server. Choose a cell, node and server from the drop-down menus. Each table row is a different thread pool. Select a row to populate the trend graph with thread pool performance metrics. Mouse over the trend graph to see metrics.



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.
- 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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

Thread Pools Table

Each table row is a different thread pool.

- Thread Pool** The name of the thread pool.
- PoolSize** The number of threads in the pool.

ActiveCount	The number of currently active threads.
growable	When checked, the connection pool is growable. Refer to vendor documentation for details.
InactivityTimeout	Refer to vendor documentation for details.
maximumSize	Refer to vendor documentation for details.
minimumSize	Refer to vendor documentation for details.
TIME_STAMP	The date and time of the last data update.

Thread Pool Trend Graph

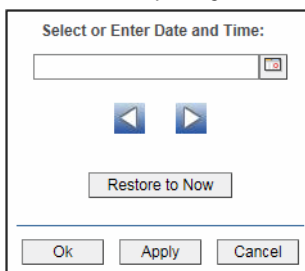
The trend graph traces the following for the thread pool on the selected server:

- **Pool Size** The number of connections in the pool.
- **Active** The number of currently active connections in the pool.


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



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

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



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath these arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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 Applications View

See high-level performance and utilization metrics for all of your web application sessions.

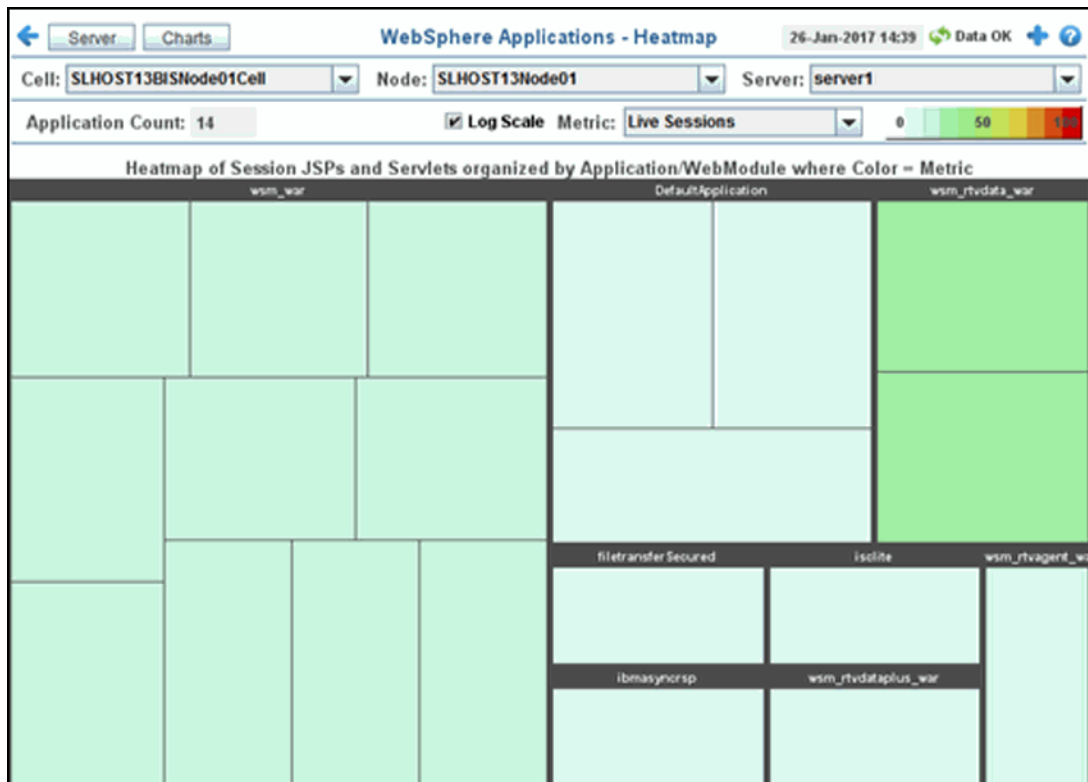
Displays in this View are:

- ["All Sessions Heatmap"](#)
- ["Session Charts By App"](#)
- ["Session Detail By App"](#)
- ["All Applications Detail"](#)

All Sessions Heatmap

This heatmap shows activity metrics for all web application sessions on a selected server. Use this display to see metrics for **Live Sessions**, **Current Requests**, **Total Requests** and **Response Time** for all web application sessions on a server. By default, this display shows the heatmap based on the **Live Sessions** metric.

Each rectangle is a different web application on the server. Choose a **Cell**, **Node** and **Server** from the drop-down menus. Mouse over a rectangle to see additional metrics. Click a rectangle to drill down to the ["Application Summary"](#) display, which shows additional details about the application.



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.
- 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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

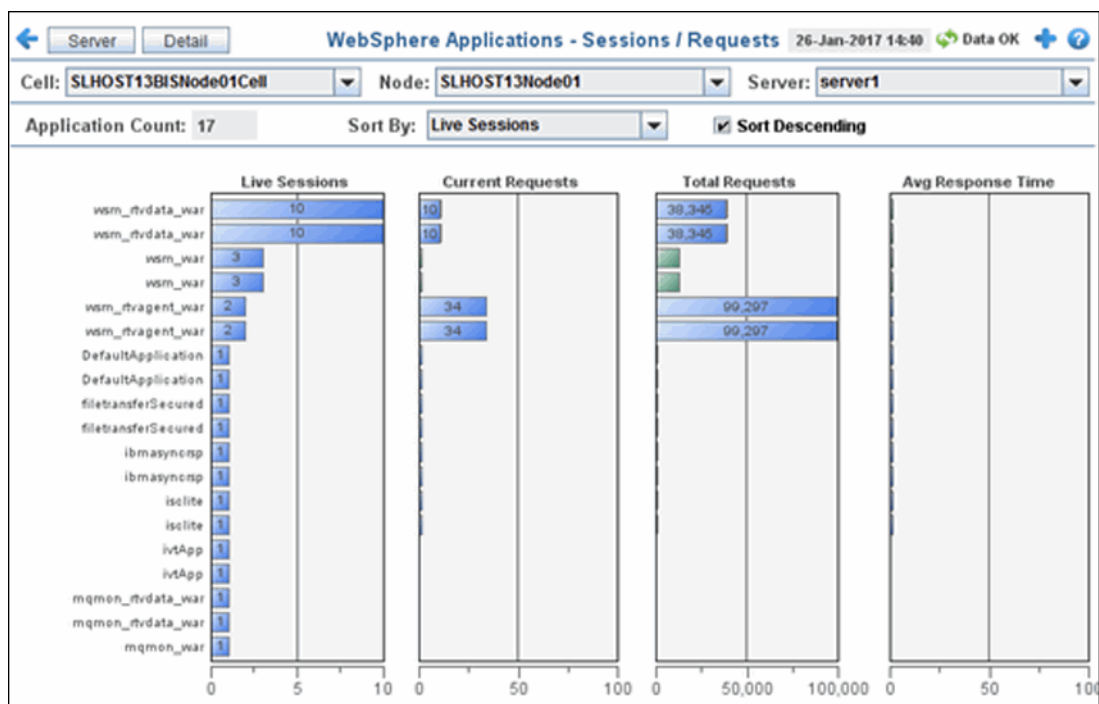
Fields and Data:

Application Count	The number of web applications in the display.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Metric	Choose a metric to view in the display.
Live Sessions	The current number of live sessions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of WasLiveSessionCountHigh . The middle value in the gradient bar indicates the middle value of the range.
Current Requests	The number of current requests. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of WasServletRequestRateHigh and WasJspRequestRateHigh . The middle value in the gradient bar indicates the middle value of the range.
Total Requests	The total number of requests. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of requests. The middle value in the gradient bar indicates the average amount.
Response Time	The average response time. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of WasServletResponseTimeHigh and WasJspResponseTimeHigh . The middle value in the gradient bar indicates the middle value of the range.

Session Charts By App

This display shows activity metrics for all web application sessions on a selected server. Choose a cell, node and server from the drop-down menus. Use this display to see metrics for **Live Sessions**, **Current Servlet Requests**, **Total Servlet Requests** and **Total JSP Requests**, **Current JSP Requests** and average response times for all web application sessions on a server. By default, this display shows the **Live Sessions** metric.

Click on an object to drill down to details.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

Fields and Data:

- Application Count** The number of web applications in the display.
- Sort By:** Choose a value to sort data in the display:
- Name** The name of the item.
 - Live Sessions** The current number of active sessions.
 - Total Servlet Requests** The number of servlet requests since the server was started.

Current Servlet Requests

The current number of servlet requests.

Total JSP Requests

The number of JSP requests since the server was started.

Current JSP Requests

The current number of JSP requests.

Servlet Avg Response Time

The average amount of time for the servlet to respond.

JSP Avg Response Time

The average amount of time for the JSP to respond.

Sort Descending

Select to sort data in descending order.

Session Detail By App

This display shows activity metrics for all web applications on a selected server. Choose a cell, node and server from the drop-down menus. Each row in the table is a different application.


← Server Charts WebSphere Applications - Sessions / Requests 26-Jan-2017 14:41 Data OK + ?

Cell: SLHOST13BISNode01Cell Node: SLHOST13Node01 Server: server1

Application Count: 17

All Applications for Selected Server

Application	Sessions	Servlets	Total Reqs	Current Reqs	Avg Resp Time	JSPs	Total Reqs	Current Req
wsm_rtvdata_war	10	1	38,345	10	0.2	1	870	
wsm_rtvdata_war	10	1	38,345	10	0.2	1	870	
wsm_war	3	0	0	0	0.0	10	12,431	
wsm_war	3	0	0	0	0.0	10	12,431	
wsm_rtvagent_war	2	1	99,297	34	1.1	0	0	
wsm_rtvagent_war	2	1	99,297	34	1.1	0	0	
DefaultApplication	1	3	1	1	1.1	0	0	
DefaultApplication	1	3	1	1	1.1	0	0	
filetransferSecured	1	1	1	1	1.1	0	0	
filetransferSecured	1	1	1	1	1.1	0	0	
ibmasyncrsp	1	1	1	1	1.1	0	0	
ibmasyncrsp	1	1	1	1	1.1	0	0	
isc-lite	1	1	1	1	1.1	0	0	
isc-lite	1	1	1	1	1.1	0	0	
MAApp	1	0	0	0	0.0	0	0	
MAApp	1	0	0	0	0.0	0	0	
mqmon_rtvdata_war	1	0	0	0	0.0	0	0	
mqmon_rtvdata_war	1	0	0	0	0.0	0	0	
mqmon_war	1	0	0	0	0.0	0	0	
mqmon_war	1	0	0	0	0.0	0	0	
PlantsByWebSphere	1	0	0	0	0.0	0	0	
PlantsByWebSphere	1	0	0	0	0.0	0	0	
SamplesGallery	1	0	0	0	0.0	0	0	
SamplesGallery	1	0	0	0	0.0	0	0	
wsm_rtvdataplus_war	1	1	1	1	1.1	0	0	
wsm_rtvdataplus_war	1	1	1	1	1.1	0	0	

**Filter By:**

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

Fields and Data:

- Application Count** The number of web applications in the display.

All Applications Table

Column values describe the application.

- Application** The name of the application.
- Sessions** The number of current sessions.
- Servlets** The number of servlets.
- Total Reqs** The number of requests since the application was started.
- Current Reqs** The number of current requests.
- Avg Resp Time** The average response time, in seconds.
- JSPs** The number of JSPs.
- Total Reqs** The number of requests since the application was started.
- Current Reqs** The number of current requests.
- Total Reqs** The number of requests since the application was started.
- Current Reqs** The number of current requests.
- Avg Resp Time** The average response time, in seconds.

All Applications Detail

This display shows detailed application information and deployment descriptors for all web application sessions on a selected server. Choose a cell, node and server from the drop-down menus. Select an application in the upper table to see the deployment descriptor in the lower portion of the display.

← Server All Applications - Info Detail Table 26-Jan-2017 14:45 Data OK + ?

Cell: SLHOST13BISNode01Cell Node: SLHOST13Node01 Server: server1

Application Count: 17

Application Info for Selected Server - Select to see Deployment Descriptor

Application	implementationVersion	statisticsProvider	eventTypes	version	
DefaultApplication			j2ee.state.starting;j2ee.state.runnin...		W
filetransferSecured			j2ee.state.starting;j2ee.state.runnin...		W
ibmasyncrsp	1 [3]		j2ee.state.starting;j2ee.state.runnin...		W
iscsite	WASX.WSC [rr0823.55]		j2ee.state.starting;j2ee.state.runnin...		W
ivApp	WASX.SERV1 [qq0723.30]		j2ee.state.starting;j2ee.state.runnin...		W
ManagementEJB			j2ee.state.starting;j2ee.state.runnin...		W

Session Manager Data for Selected Server

Application	LiveCount	TIME_STAMP
DefaultApplication	1	01/26/17 14:45:14
DefaultApplication	1	01/26/17 14:45:14
filetransferSecured	1	01/26/17 14:45:14
filetransferSecured	1	01/26/17 14:45:14
ibmasyncrsp	1	01/26/17 14:45:14

Deployment Descriptor:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE application PUBLIC "-//Sun Microsystems, Inc.//DTD J2EE Application 1.3/EN" "http://java.sun.com/dtd/application_1_3.dtd">
<application id="Application_ID">
  <display-name>DefaultApplication.ear</display-name>
  <description>This is the IBM WebSphere Application Server Default Application.</description>
  <module id="WebModule_1">
    <web>
      <web-uri>DefaultWebApplication.war</web-uri>
      <context-root>/</context-root>
    </web>
  </module>
</application>
```

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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

Fields and Data:

- Application Count** The number of web applications in the display.

Application Info Table

Each table row is a different application. Click a row to see data in the **Deployment Descriptor** table.

Application	The name of the application.
implementationVersion	The application version.
statisticsProvider	The name of the statistics provider.
eventTypes	A list of application events.
version	The application version.
modules	A list of application modules.
TIME_STAMP	The date and time of the last data update.

Session Manager Data Table

Each table row is a different application.

Application	The name of the application.
LiveCount	The number of connections for the application.
TIME_STAMP	The date and time of the last data update.

Deployment Descriptor:

Provides details about the application deployment.

Single Application View


See performance and utilization metrics for a single web application.

Displays in this View are:

- ["Application Summary"](#)
- ["Component Detail"](#)
- ["Module Totals - Charts"](#)
- ["Module Totals - Tables"](#)

Application Summary

View performance metrics for a web application on one server. Choose a cell, node, server and application from the drop-down menus.



Title Bar (possible features are):

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

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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.

Fields and Data

This display includes:

- Application Count** The number of applications in the display.

Live Sessions The current number of active sessions.

Module Count The current number of modules.

JSP/Servlet Request Totals

JSPs The current number of JSP requests.

Servlets The current number of servlet requests.

EJB Call Totals

EJBs The current number of EJB (Enterprise JavaBean) requests.


Session/JSP/Servlet/EJB Totals Trend Graph

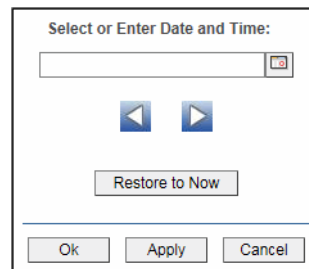
The trend graph traces the following for the selected application:

- **Live Sessions** The current number of active sessions.
- **JSP Requests** The current number of JSP requests.
- **JSP Resp Time** The current average JSP response time, in seconds.
- **Servlet Requests** The current number of servlet requests.
- **Servlet Resp Time** The current average servlet response time, in seconds.
- **EJB Method Calls** The current number of EJB requests.
- **EJB Resp Time** The current average EJB response time, in seconds.


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



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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

Component Detail

View performance metrics for a web application on one server. Choose a cell, node, server and application from the drop-down menus.

WebSphere Application - Component Tables 28-Nov-2016 13:29 Data OK

Cell: TESTBED-27Node01Cell Node: TESTBED-27Node01 Server: server1

Application: DefaultApplication Application Count: 8

Live Sessions

WebModule	LiveCount	NoRoomForNewSessionCour	LifeTime	ExternalRe	ExternalRi	ExternalWr	ExternalWr
DefaultWebApplication.war	0	0	0	0	0	0	0

JSPs

Servlets

Name	Requests	Total Requests	Avg Response Time Rec	Total Response Time	Avg Response Tin
Hello Pervasive Servlet	0	0	0.0	0.0	0.0
Hit Count Servlet	0	0	0.0	0.0	0.0
Snoop Servlet	0	0	0.0	0.0	0.0

EJBs

Name	CreateCount	Method Calls	Total Calls	Ready Count	Response Time	PassiveCour	RemoveCou
------	-------------	--------------	-------------	-------------	---------------	-------------	-----------

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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.

Fields and Data:

- Application Count** The number of applications in the display.

Live Sessions Table

Each table row is a different web module. Table columns describe the module.

WebModule	The name of the web module.
LiveCount	The current number of sessions.
NoRoomForNewSessionCount	The number of times the module had no sessions available.
Life Time	Refer to vendor documentation for details.
ExternalReadSize	Refer to vendor documentation for details.
ExternalWriteSize	Refer to vendor documentation for details.
ExternalWriteTime	Refer to vendor documentation for details.

JSPs Table

Each table row is a different JSP. Table columns describe the JSP.
Refer to vendor documentation for details.

Servlets Table

Each table row is a different servlet. Table columns describe the servlet.

Name	The name of the servlet.
Requests	The current number of requests.
Total Requests	The total number of requests since the servlet was started.
Avg Response Time Recent	The current average amount of time for the servlet to respond.
Total Response Time	The total response time, in seconds, since the servlet was started.
Avg Response Time	The average amount of time for the servlet to respond since the servlet was started.


EJBs Table

Each table row is a different EJB. Table columns describe the EJB.

Name	The name of the EJB.
CreateCount	The number of requests.
MethodCalls	The total number of requests since the servlet was started.
Total Calls	The current average amount of time for the servlet to respond.
Ready Count	Refer to vendor documentation for details.
Response Time	The average amount of time for the EJB to respond.
PassiveCount	Refer to vendor documentation for details.
RemoveCount	Refer to vendor documentation for details.

Module Totals - Charts

View performance metrics for a WAS application on one server. Choose a cell, node, server and application from the drop-down menus.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.

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.

Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.
- Sort By:**
- **Name**
 - **Live Sessions**
 - **Total Servlet Requests**
 - **Current Servlet Requests**
 - **Total JSP Requests**
 - **Current JSP Requests**
 - **Servlet Avg Resp Time**
 - **JSP Avg Resp Time**

Fields and Data:

Sort Descending Select to organize display elements in descending order.

Module Count The number of modules in the display.

Graphs

Refer to vendor documentation for details.

Live Sessions Graph Shows performance metrics for current sessions.

Total Requests Graph Shows performance metrics for total requests.

Current Requests Graph Shows performance metrics for current requests.

Avg Response Time Graph Shows performance metrics for average response time.

Module Totals - Tables

View performance metrics for a web application on one server. Choose a cell, node, server and application from the drop-down menus. Each row in the upper table is a different module for the selected application. Select a row to populate the lower tables.

Single WAS Application - Module Totals 26-Jan-2017 14:51 Data OK

Cell: SLHOST13Node01Cell Node: SLHOST13Node01 Server: server1

Application: isclite

Module Count: 6


All Modules for Selected Application

Module Name	Context Root	Sessions	Servlets	Total Reqs	Current Reqs	Avg Resp Time	JSPs	T
adminredirector.war	admin_host/admin	1	0	0	0	0.0	0	
adminredirector.war	admin_host/admin	1	0	0	0	0.0	0	
iehs.war	admin_host/ibm/help	1	0	0	0	0.0	0	
iehs.war	admin_host/ibm/help	1	0	0	0	0.0	0	
ISCAdminPortlet.war	admin_host/ISCAdminPortlet	1	0	0	0	0.0	0	
ISCAdminPortlet.war	admin_host/ISCAdminPortlet	1	0	0	0	0.0	0	
isclite.war	admin_host/ibm/console	1	1	1	1	1.1	0	
isclite.war	admin_host/ibm/console	1	1	1	1	1.1	0	
wasportlet.war	admin_host/wasportlet	1	0	0	0	0.0	0	
wasportlet.war	admin_host/wasportlet	1	0	0	0	0.0	0	
WIMPortlet.war	admin_host/wim	1	0	0	0	0.0	0	
WIMPortlet.war	admin_host/wim	1	0	0	0	0.0	0	

EJBs for Selected Module

Portlets for Selected Module

isclite#ISCAdminPortlet.war_portlet.ConsoleIdentityPortlet
 isclite#ISCAdminPortlet.war_portlet.ISCProductDetails
 isclite#ISCAdminPortlet.war_portlet.PortletPref
 isclite#ISCAdminPortlet.war_portlet.WelcomePortlet

**Filter By:**

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.

Fields and Data:

- Module Count** The number of modules in the display.

All Modules Table

Each table row is a different module. Column values describe the module. Select a module to see EJBs and Portlets for the module in the lower tables.

- Module Name** The name of the module.
- Context Root** The context root.
- Sessions** The current number of sessions.
- Servlets** The current number of servlets.
- Total Requests** The total number of requests since the servlet was started.
- Current Requests** The current number of requests.
- Avg Response Time** The average amount of time to respond, in seconds.
- JSPs** The current number of JSPs.
- Total Response Time** The total response time, in seconds, since the servlet was started.
- Avg Response Time** The average amount of time for the servlet to respond since the servlet was started.
- Total Requests** The total number of requests since the servlet was started.
- Current Requests** The current number of requests.

**Avg
Response
Time**

The average amount of time to respond, in seconds.

EJBs for Selected Module

List of EJBs for the selected module.

Portlets for Selected Module

List of portlets for the selected module.

CHAPTER 4 RTView DataServer for Infrastructure

The RTView DataServer for Infra provides a way to create connections and modify default configuration settings for the various solution packages and sends collected data to RTView Central, which contains the displays associated with the RTView DataServer for Infra that help you to monitor the health and performance across your components.

RTView Central contains the following Views and their associated displays that will be populated with data collected via the RTView DataServer for Infra:

- "Amazon Web Services"
- "Docker"
- "JBoss"
- "MongoDB"
- "MySQL"
- "MS SQL Server"
- "Node.js"
- "RTView Host Agent"
- "UX"
- "VMware vCenter"

The following displays are also impacted by the settings in the RTView DataServer for Infra:

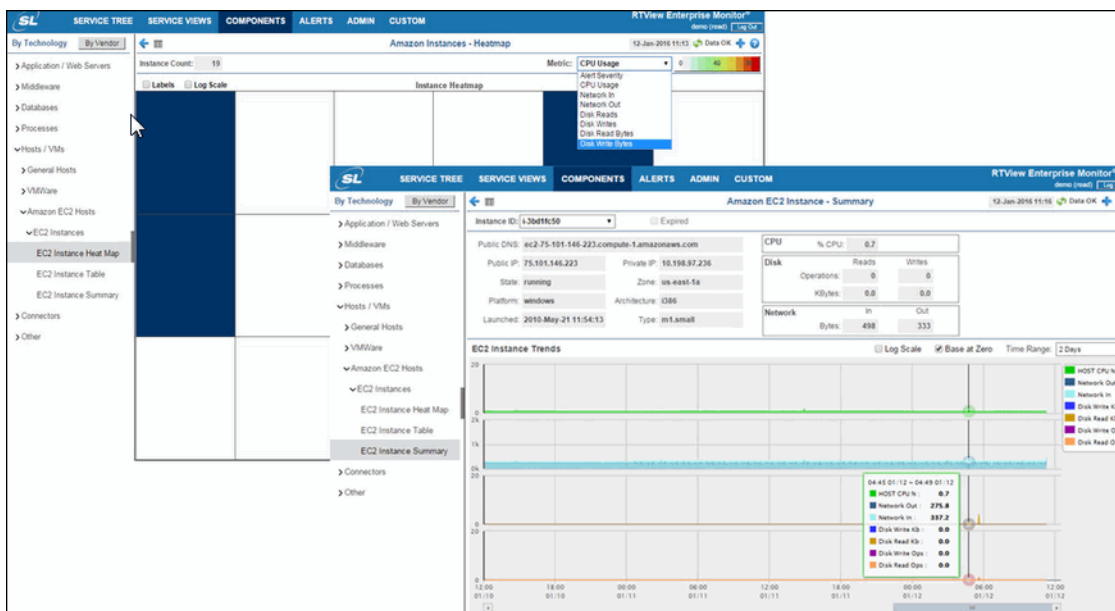
- "JVM "
- "Tomcat "

The RTView *DataCollector* for Infra is also available for use with the RTView DataServer for IBM. RTView DataCollector for Infra is used for collecting and sending data to one or more data servers. The RTView DataCollector for Infra is also useful if you need to distribute data collection.

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

Amazon Web Services

The Solution Package for Amazon Web Services provides a high level Amazon Instance Heatmap for a complete view of your AWS infrastructure with drill down views to individual AWS instances.



Using the RTView Historian, Amazon AWS metrics are persisted to a database for trend analysis. Historical trends are then used to help define alert thresholds against Amazon AWS data which, when correlated with alerts from other application components through RTView Enterprise Monitor's alert management system, can help users identify the source of performance problems more quickly.

With the Solution Package for Amazon Web Services, you are able to drill down from a high level alert on a business service or application into the supporting Amazon AWS infrastructure to determine what is causing the alert and take corrective action. This service-centric approach makes it easy for application support teams to prioritize incidents based on the impact to the business.

The following Solution Package for Amazon Web Services Views (and associated displays) can be found under **Components tab > Hosts/VMs > Amazon EC2 Hosts**. For additional details, see vendor documentation.

This section contains the following:

- "EC2 Instances"

EC2 Instances

Displays in this View are:

- "Amazon EC2 Instance Heatmap"
- "Amazon EC2 Instance Table"
- "Amazon EC2 Instance Summary"


Amazon EC2 Instance Heatmap

View the most critical alert states associated with your Amazon EC2 instances. Use this display to quickly identify instances with critical alerts. Compare heap usage, disk reads and writes and network throughput rates across all monitored instances.

Each rectangle in the heatmap represents an Amazon EC2 instance. The rectangle color indicates the most critical alert state associated with the instance for the selected **Metric**.

Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics, including disk reads and writes, CPU utilization and network in/out rates. By default, this display shows **Alert Severity**.

Use the **Labels** check-box ☒ to include or exclude labels in the heatmap. Click a rectangle to drill-down and view instance details in the ["Amazon EC2 Instance Summary"](#) display.














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.
- 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.

Fields and Data:

Instance Count:	The total number of instances currently shown in the display.
Labels:	Select to show labels in the display.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The maximum level of alerts in the rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
CPU Usage	The percent (%) CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
Network In	The number of incoming bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
Network Out	The number of outgoing bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
Disk Reads	The number of completed disk reads. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
Disk Writes	The number of completed disk writes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
Disk Read Bytes	The amount of disk reads, in bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
Disk Write Bytes	The amount of disk writes, in bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.

Amazon EC2 Instance Table

View detailed utilization data for all your Amazon EC2 instances in a tabular format. Use this display to see all available data for this View.

Each row in the table is a different Amazon EC2 instance. Use the **Show:** drop-down menu to only show instances that are **running** or **stopped**.

Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the ["Amazon EC2 Instance Summary"](#) display.

Instance	Public DNS Name	State	% CPU	Disk Reads (bytes)	Disk Reads (Ops)	Disk V (byt)
i-094dec		stopped				
i-1a0a9e		stopped				
i-224dbe40		stopped				
i-2a31cdb3		stopped				
i-38405553		stopped				
i-3bd1fc50	ec2-75-101-146-223.compute-1.amazonaws.com	running	1.00	0	0	
i-41fa0fd8		stopped				
i-4a5ead28	ec2-54-144-4-125.compute-1.amazonaws.com	running	1.13	0	0	
i-570baf82		running	.40	0	0	
i-67562034		stopped				
i-6a3dd2f3		stopped				
i-6a838401	ec2-184-73-176-2.compute-1.amazonaws.com	running	98.67	0	0	
i-97280b20		stopped				
i-9a712903		running	.03	0	0	
i-9cac5205		stopped				
i-ac348fff		stopped				
i-aceeebc7	ec2-54-221-170-144.compute-1.amazonaws.com	running	1.27	0	0	
i-b6a6492f		stopped				
i-b7e3284d		running	4.57	0	0	
i-bbb3c1d1		stopped				
i-c5fad2ae	ec2-184-73-131-119.compute-1.amazonaws.com	running	1.13	0	0	
i-cfa13aah	ec2-54-235-4-76.compute-1.amazonaws.com	running	1.26	0	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.

Instance Count: The number of instances in the table.

Filter By:
The display might include these filtering options:

Show: Choose to show **All** instances, **running** or **stopped** instances.

All Instances Table:
Each row in the table is a different instance.

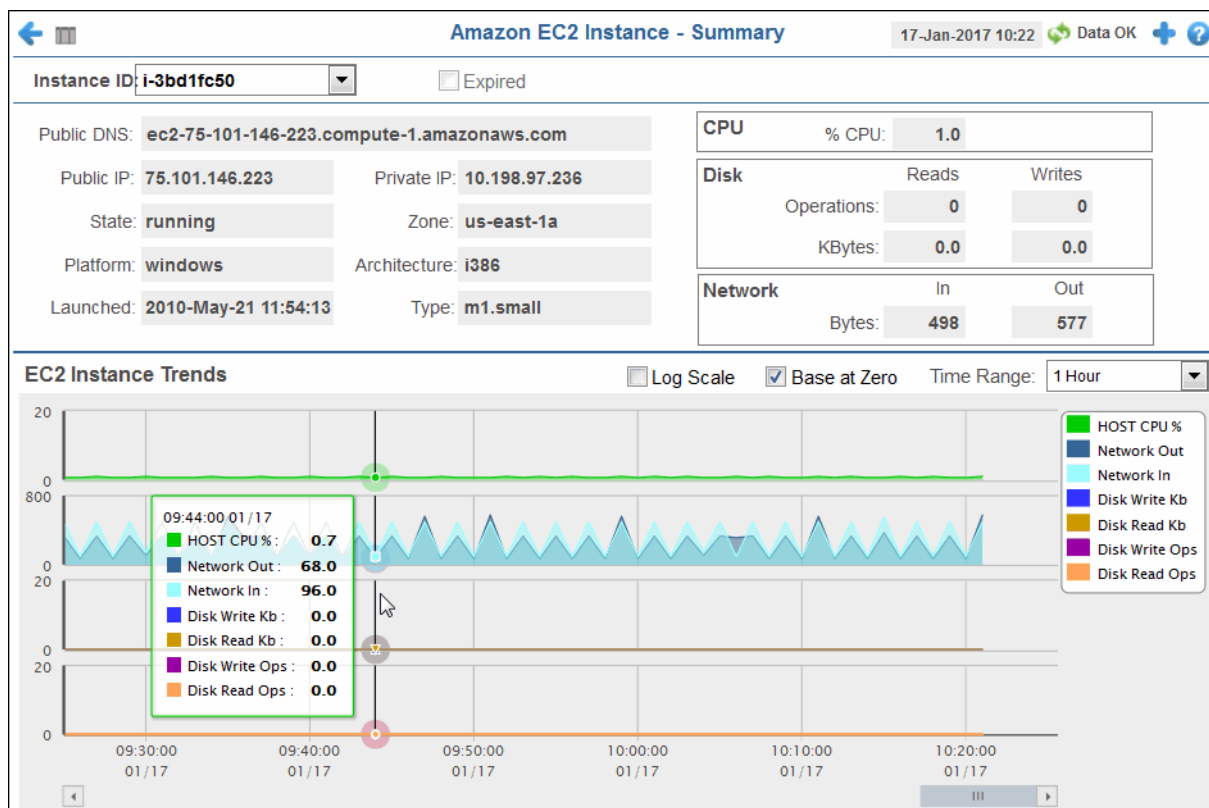
Instance The name of the instance.

Alert Severity	<p>The maximum level of alerts in the row. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics exceeded their alert thresholds.
Public DNS Name	The public domain name of the instance.
State	The instance state (running or stopped).
%CPU	The percent CPU used.
Disk Reads (bytes)	The amount of disk reads, in bytes.
Disk Reads (Ops)	The number of disk reads (count).
Disk Writes (bytes)	The amount of disk writes, in bytes.
Disk Writes (Ops)	The number of disk writes (count).
Network In	The number of incoming bytes.
Network Out	The number of outgoing bytes.
Instance Type	The instance type (e.g. m1.small).
Private IP	The instance private IP address.
Public IP	The instance public IP address.
Platform	The instance operating system (e.g. windows).
Architecture	The instance architecture (e.g. i386).
Image ID	<p>The unique identifier for the image.</p> <p>For details about Amazon EC2 data, refer to vendor documentation.</p>
Root Device Name	<p>The name of the root device.</p> <p>For details about Amazon EC2 data, refer to vendor documentation.</p>
Root Device Type	<p>The type of root device.</p> <p>For details about Amazon EC2 data, refer to vendor documentation.</p>
Availability Zone	<p>The id for the availability zone (e.g. us-east-1a).</p> <p>For details about Amazon EC2 data, refer to vendor documentation.</p>
Group	For details, see vendor documentation.
Tenancy	For details about Amazon EC2 data, refer to vendor documentation.
Tags	For details, see vendor documentation.
State Transition Reason	For details, see vendor documentation.
LaunchTime	The date and time the instance was started.

Timestamp	The date and time the data was last updated.
Expired	When checked, data has not been received from this instance in the specified amount of time. The instance will be removed from the in the specified amount of time. The default setting is 60 seconds.

Amazon EC2 Instance Summary

This display provides detailed utilization metrics for a single Amazon EC2 instance. Use this display to investigate performance details and trends for an instance.



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 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.

Filter By:

The display might include these filtering options:

Instance ID: Choose an instance to show data for in the display.

Fields and Data:

Data describes the selected host except where noted.

	Expired	When checked, data has not been received from this instance in the specified amount of time. The instance will be removed from the in the specified amount of time. The default setting is 60 seconds.	
	Public IP	The instance public IP address.	
	State	The instance state (running or stopped).	
	Platform	The instance operating system (e.g. windows).	
	Launched	The date and time the instance was started.	
	Private IP	The instance private IP address.	
	Zone	The date and time the instance was started.	
	Architecture	The instance architecture (e.g. i386).	
	Type	The instance type (e.g. m1.small).	
CPU	%CPU	The percent CPU used.	
Disk	Operations	Reads	The number of disk reads (count).
		Writes	The number of disk writes (count).
	KBytes	Reads	The amount of disk reads, in kilobytes.
		Writes	The amount of disk writes, in kilobytes.
Network	Bytes In	The number of incoming bytes.	
	Bytes Out	The number of outgoing bytes.	

EC2 Instance Trends

Traces metrics for the selected instance.


- **Host CPU%:** The amount of CPU used, in percent.
- **Network Out:** The number of outgoing bytes.
- **Network In:** The number of incoming bytes.
- **Disk Write Kb:** The amount of disk writes, in kilobytes.
- **Disk Read Kb:** The amount of disk reads, in kilobytes.
- **Disk Write Ops:** The number of disk writes (count).
- **Disk Read Ops:** The number of disk reads (count).


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


Base at Zero



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

Time Range

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



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

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

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

Docker

The Solution Package for Docker is an easy to configure and use monitoring system that gives you extensive visibility into the health and performance of your Docker Engines, Docker Containers, and the applications that rely on them.

The following Docker Views can be found under **Components** tab > **Processes** > **Docker Engines**:

- **"Engine View"**: The displays in this View allow you to view the current and historical metrics for all engines in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single engine.
- **"Container View"**: The displays in this View allow you to view the current and historical metrics for all containers in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single container.

Engine View

These displays provide detailed data for all engines or for a particular engine. Displays in this View are:

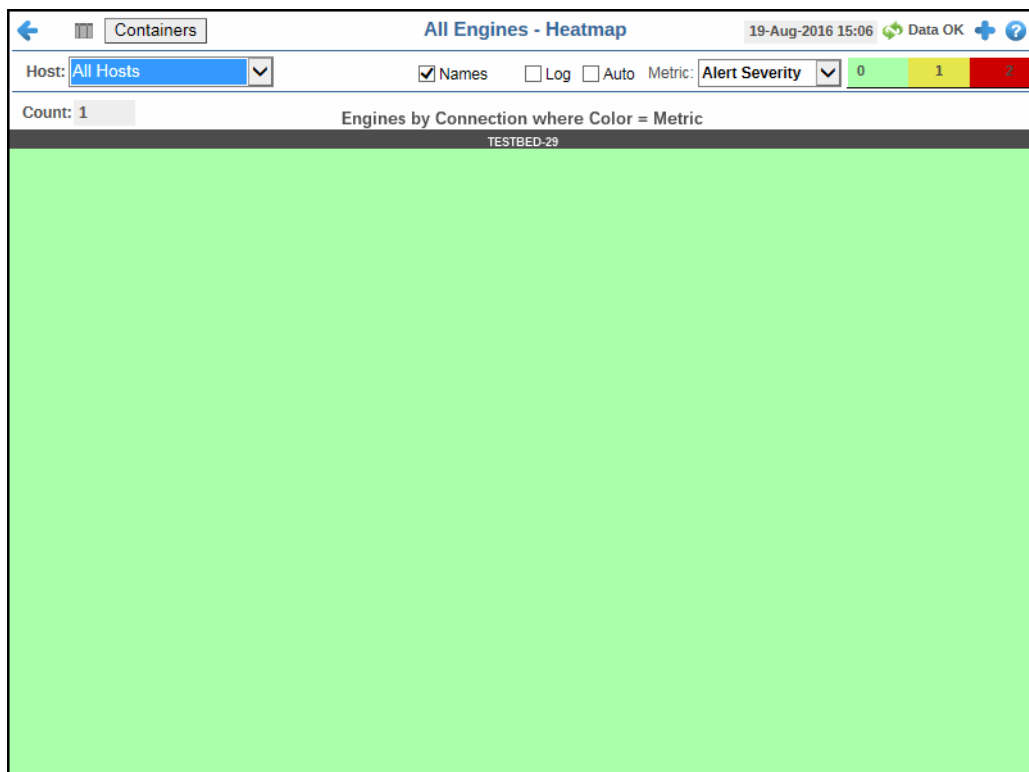
- **"Engines Heatmap"**: A heatmap view of all engines and their associated metrics.
- **"Engines Table"**: A tabular view of your engines and their associated metrics.
- **"Engine Summary"**: Provides additional details and a way to view trending data for a single engine.


Engines Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your engines for each available metric. You can view the engines in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, the amount of memory used, the total incoming bytes, and the total outgoing bytes. By default, this display shows the heatmap based on the **Alert Severity** metric.

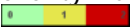





You can use the **Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an engine. Clicking one of the rectangles in the heatmap opens the “[Engine Summary](#)” display, which allows you to see additional details for the selected engine.




Note: When the data for the engine being monitored expires, the color of the rectangle representing that engine in the heatmap automatically changes to a color that is not included in the color gradient bar so that you can easily identify when the data is stale. Expired data could occur for a number of reasons including, but not limited to, the connection to the engine may have been lost, or the engine could have experienced a problem and may no longer be up-and-running.





Fields and Data:

Host	Select the host for which you want to show data in the display.
Count	Lists the total number of engines found using the search parameters.
Names	Select this check box to display the names of the engines at the top of each rectangle in the heatmap.
Log	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. Note: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning unacknowledged alerts in the engine. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
CPU Usage	<p>The percentage of CPU used by the engine. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of DocEngineCpuUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

Memory	<p>The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The Auto option does not impact this metric.</p>
Net Bytes In	<p>The total number of incoming bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of DocEngineNetBytesInHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Net Bytes Out	<p>The total number of outgoing bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of DocEngineNetBytesOutHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

Engines Table

This table provides a view of all of your engines and their associated metric data including host, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected engine in the [“Engine Summary”](#) display

Containers

All Engines - Table

19-Aug-2016 14:44 Data OK

Host: All Hosts

Count: 1

Host	Alert Level	Alert Count	CPU Usage	Memory Available (KB)	Memory Usage (KB)	Memory WS (KB)	Memory RSS (KB)	Memory Limited	Net Bytes In avg
TESTBED-29		0	11.39	3,782,232	3,373,604	1,602,908	58,564	<input checked="" type="checkbox"/>	81,20

Title Bar (possible features are):

← ↑

Open the previous and upper display.

+

Open an instance of this display in a new window.

?

Open the online help page for this display.

Menu

Table

open commonly accessed displays.

6,047

The number of items currently in the display.

Data OK

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

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: The **Containers** button takes you to “Containers Table”.

Fields and Data:

- Host




Select the name of the host (or **All Hosts**) containing the engines for which you want to view data.
- Count

The total number of engines being monitored based on your search criteria.

All Engines Table:


- Host

The name of the host.

Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the host.
CPU Usage	The percentage of CPU used by the engine.
Memory Available (KB)	The amount of memory, in kilobytes, that is available to the engine.
Memory Usage (KB)	The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed.
Memory WS (KB)	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
Memory RSS (KB)	The amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
Memory Limited	When checked, the amount of memory available to the engine is limited.
Net Bytes In avg	The average number of incoming bytes per second.
Net Bytes Out avg	The average number of outgoing bytes per second.
Net Packets In avg	The average number of incoming packets per second.
Net Packets Out avg	The average number of outgoing packets per second.
Docker Version	The Docker software version of the Docker Engine.
Container OS Version	The version of the container's operating system on which the docker engine is running.
Container Kernal Version	The version of the container's Kernal in which the docker engine is running.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Docker > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

Engine Summary

This display allows you to view current as well as trending data for the percentage of CPU used by the engine, memory usage details, and network data details.



Title Bar (possible features are):

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

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

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

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

Note: The **Containers** button takes you to "Containers Table".

Filter By:

Host


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


Container OS Version


The version of the container's operating system on which the docker engine is running.



Fields and Data:

Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Docker > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.	
Alert State	<p>The current alert severity.</p> <ul style="list-style-type: none"> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. ● Green indicates that no metrics have exceeded their alert thresholds. 	
Alert Count	The total number of current alerts.	
Docker Version	The Docker software version of the Docker Engine.	
CPU		
	% CPU	The percentage of CPU used by the engine.
Memory (KB)		
	Usage	The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed.
	WS	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
	RSS	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
Network		
	Packets In/sec	The average number of incoming packets per second..
	Packets Out/sec	The average number of outgoing packets per second.
	Bytes In/sec	The average number of incoming bytes per second.
	Bytes Out/sec	The average number of outgoing bytes per second.
Performance Trends Graph	<p>Traces the following:</p> <ul style="list-style-type: none"> % CPU -- traces the percentage of CPU being used on the engine. Memory (KB) -- traces the amount of memory, in kilobytes, used by the engine. Net Bytes In -- traces the average number of incoming bytes per second. Net Bytes Out -- traces the average number of outgoing bytes per second. <p>Log Scale Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p>	

- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



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

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

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

Container View

These displays allow you to view the current and historical metrics for all containers in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single container. Displays in this View are:


- **"Containers Heatmap"**: A color-coded heatmap view of data for all containers for a particular host.
- **"Containers Table"**: A tabular view of data for all containers for a particular host.
- **"Container Summary"**: This display allows you to view current and trending data for a single container for a particular host.

Containers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your containers for each available metric. You can view the containers in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, and the percentage of memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a container. Clicking one of the rectangles in the heatmap opens the “[Container Summary](#)” display, which allows you to see additional details for the selected container.

Note: When the data for the container being monitored expires, the color of the rectangle representing that container in the heatmap automatically changes to a color that is not included in the color gradient bar so that you can easily identify when the data is stale. Expired data could occur for a number of reasons including, but not limited to, the connection to the container may have been lost, or the container could have experienced a problem and may no longer be up-and-running.

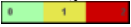








Title Bar (possible features are):

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

Fields and Data:

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

Container	Select the container (or All Containers) for which you want to show data in the heatmap..
Count	Lists the total number of containers (rows) found using the search parameters.
Names	Select this check box to display the names of the containers at the top of each rectangle in the heatmap.
Log	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. Note: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of critical and warning unacknowledged alerts in the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
CPU Usage	<p>The percentage of CPU used by the container. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of DocContainerCpuUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Memory	<p>The current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The Auto option does not impact this metric.</p>

Net Bytes In

The number of incoming bytes per second. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of

DocContainerNetBytesInHigh. The middle value in the gradient bar indicates the middle value of the range.

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

Net Bytes Out

The number of outgoing bytes per second. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of

DocContainerNetBytesOutHigh. The middle value in the gradient bar indicates the middle value of the range.

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

Containers Table

This display allows you to view details in a table format for one container on a particular host, for all containers on a particular host, for a particular container on all hosts, or for all containers on all hosts. You can drill-down and view the details for a particular container in the [“Container Summary”](#) display by clicking on a row in the resulting table.


⬅ ⬆ 🏠
All Containers - Table
19-Aug-2016 14:55 🔄 Data OK + ?

Host: All Hosts Container: All Containers

Count: 3

Host	Container Name	Container ID	Alert Level	Alert Count	CPU Usage	Memory Available (KB)	Memory Usage (KB)	Memc WS (KB)
TESTBED-29	cadvisor-rtview-TB29	4c58c59ae430	🟢	0	0.46	3,782,232	53,704	3
TESTBED-29	cadvisor-rtview-TB3	822a5c6601a8	🟢	0	0.36	3,782,232	24,968	1
TESTBED-29	cadvisor-rtview-TB6	8fac67ccf6d0	🟢	0	0.43	3,782,232	22,168	1

< >



Filter By:

The display includes these filtering options:

Host	Select the host for which you want to show data in the display.
Container	Select the container (or All Containers) for which you want to view data..
Count	Lists the total number of containers (rows) found using the search parameters.


All Containers Table

Host	The name of the host.
Container Name	The name of the container.
Container ID	The absolute container name.
Alert Level	The current alert status. <div> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. ● Green indicates that no metrics have exceeded their alert thresholds. </div>
Alert Count	Total number of alerts for the process.
CPU Usage	The percentage of CPU used by the container.
Memory Available (KB)	The amount of memory, in kilobytes, that is available to the container.
Memory Usage (KB)	Current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.
Memory WS (KB)	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
Memory RSS (KB)	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
Memory Limited	When checked, the amount of memory available to the container is limited. If not checked, then the amount of memory available to the container is unlimited, which means the amount of memory available to the container is the same as the memory available to the engine.
Net Bytes In avg	The average number of incoming bytes per second.
Net Bytes Out avg	The average number of outgoing bytes per second.

Net Packets In avg	The average number of incoming packets per second.
Net Packets Out avg	The average number of outgoing packets per second.
Uptime	The amount of time (in seconds) that the container has been up and running.
Running	When checked, this check box indicates that the container is running.
Status	<p>The current status of the container. Values are:</p> <p>Up - indicates that the container is up and running, and lists the amount of time the container has been up and running (Uptime).</p> <p>Created - indicates that the container has been created but is currently not in use.</p> <p>Exited - indicates that the container has been stopped, and lists the error code as well as the amount of time since the container was stopped.</p>
Starts	<p>The number of times the container (re)started within the time specified (in seconds) in the \$docEventCacheTimeRange field in the conf\rtvapi_dockermon.properties file. The default is 3600 seconds (1 hour). For example, by default, this row column lists the number of times the container has (re)started in the past hour. This number provides a good indication of the stability of the container; the higher the number, the more unstable the container.</p>
Expired	<p>When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Docker > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p>
Timestamp	The date and time the row data was last updated.

Container Summary

This display provides a view of the current and historical metrics for a single container. You can view the current information pertaining to CPU usage percentage, Memory details, Disk read and write details, and network data details in the upper portion of the display. The trend graph in the bottom half of the display traces the current and historical CPU usage, the average memory used, and the number of incoming and outgoing network bytes.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Host** Select the host for which you want to show data in the display.
- Container** Select the container for which you want to show data in the display.
- ID** The absolute container name.


Fields and Data:


- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **Docker** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.



Alert State	<p>The current alert severity.</p> <ul style="list-style-type: none"> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. ● Green indicates that no metrics have exceeded their alert thresholds. 	
Alert Count	The total number of current alerts.	
Running	When checked, this check box indicates that the container is running.	
Status	<p>The current status of the container. Values are:</p> <p>Up - indicates that the container is up and running, and lists the amount of time the container has been up and running (Uptime).</p> <p>Created - indicates that the container has been created but is currently not in use.</p> <p>Exited - indicates that the container has been stopped, and lists the error code as well as the amount of time since the container was stopped.</p>	
CPU		
	% CPU	The percentage of CPU used by the container.
Memory (KB)		
	Usage	The current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.
	WS	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
	RSS	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
Network		
	Packets In/sec	The average number of incoming packets per second.
	Packets Out/sec	The average number of outgoing packets per second.
	Bytes In/sec	The average number of incoming bytes per second.
	Bytes Out/sec	The average number of outgoing bytes per second.
Performance Trends Graph	<p>Traces the following:</p> <ul style="list-style-type: none"> % CPU -- traces percentage of CPU used by the container. Memory (KB) -- traces the current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed. Net Bytes In -- traces the average number of incoming bytes per second. Net Bytes Out -- traces the average number of outgoing bytes per second. 	
	Log Scale	<p>Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p>

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

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



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

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

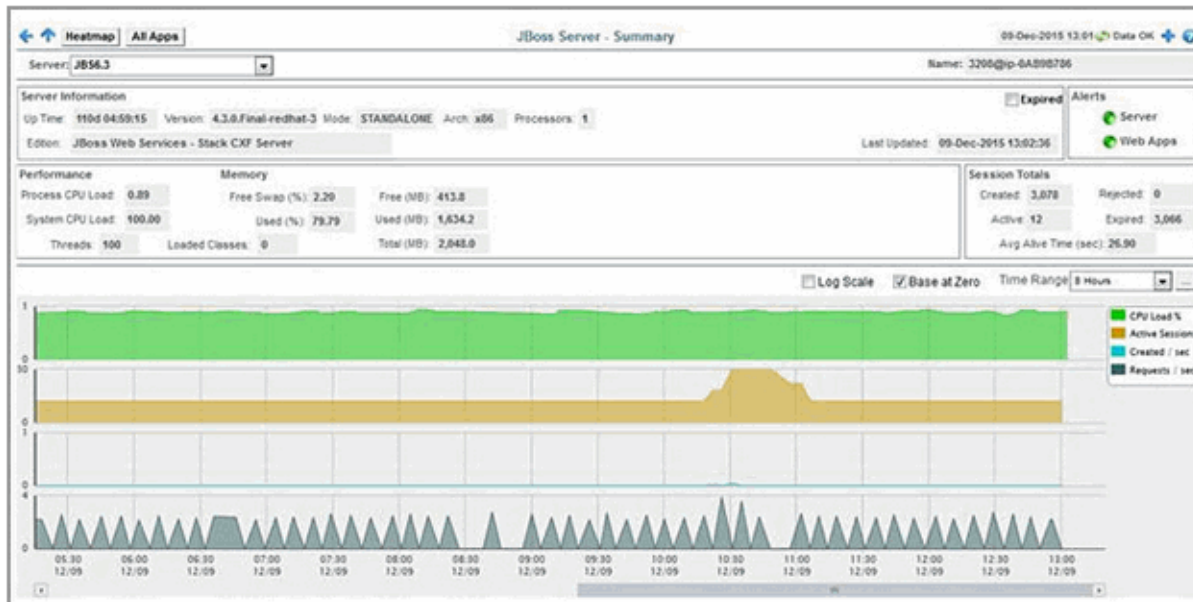
Click **Restore to Now** to reset the time range end point to the current time.

JBoss

The following JBoss Views can be found under **Components** tab > **Application/Web Servers** > **JBoss**:

- **"JBoss Servers View"**: The displays in this View present server performance metrics such as CPU and memory utilization.
- **"JBoss Applications"**: The displays in this View present views of the VPN-level metrics.

With the Solution Package for Red Hat® JBoss® you are able to identify potential problems before they become critical and impact overall application performance. Typical installations of RTView Enterprise Monitor and its solution packages take only a few hours, while developing custom views for a variety of IT and development roles can be achieved in just days.



JBoss Servers View

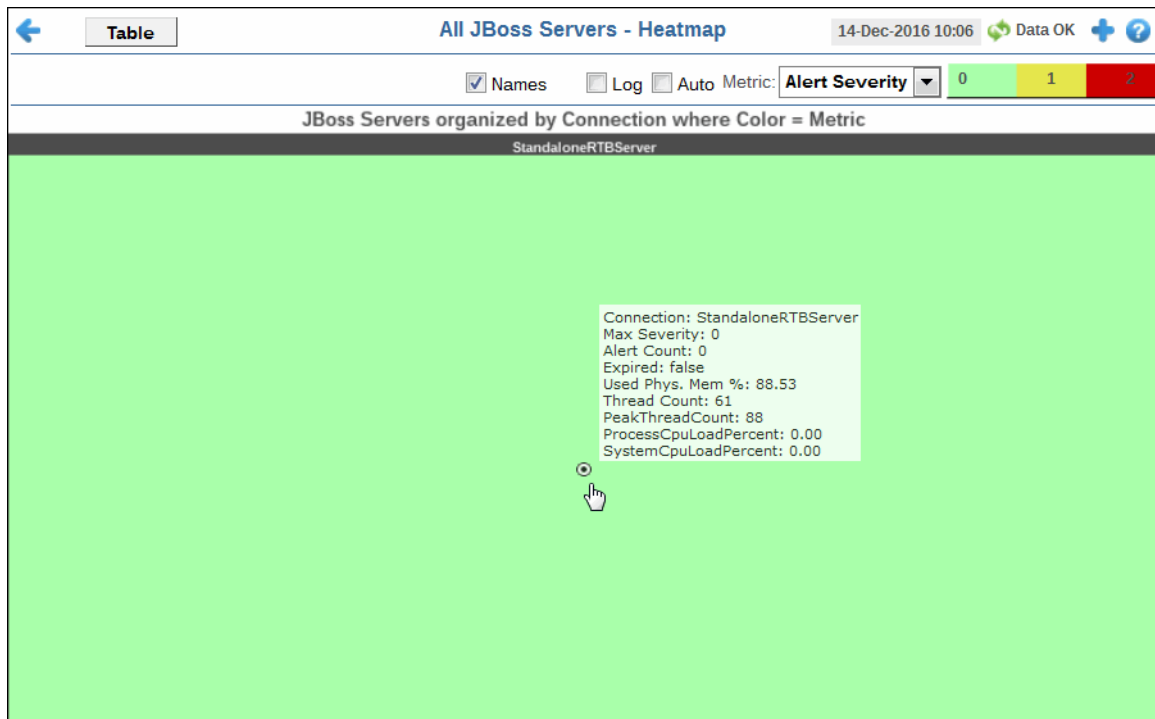
Displays in this View are:

- "All Servers Heatmap"
- "All Servers Table"
- "Server Summary"

All Servers Heatmap

This heatmap shows the current status of connections on all JBoss servers. Use this display to quickly assess the current status of connections using various metrics, including **Alert Count** and **CPU Used %** and **Virtual Memory Used %**. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle is a different JBoss server. Use the **Names** check-box ☒ to include or exclude labels in the heatmap, and mouse over a rectangle to see additional metrics for a server. Click a rectangles to drill down to the “[Server Summary](#)” display, which shows additional details about the server.



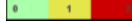







Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data:

- Names** Select to include labels in the heatmap.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning unacknowledged alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
CPU Used%	<p>The percent CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of JbsServerCpuUsedHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
V Memory Used%	<p>The percent virtual memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of JbsServerMemUsedHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Free Memory	<p>The total amount of available memory. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount of available memory. The middle value in the gradient bar indicates the average amount.</p> <p>The Auto flag does not have any impact on this metric.</p>

All Servers Table


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

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the “[Server Summary](#)” display.

Connection	Name	Expired	Alert Level	Process Cpu %	Threads	Peak Threads	Daemon Threads
StandaloneRTBServer	1996@WIN-8-CLONE	<input type="checkbox"/>		0.00	58	88	39

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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.



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

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.



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

Fields and Data

This display includes:

- Connection** The name of the connected server.
- Name** The name of the connection.
- Expired** When checked, data has not been received from this host in the specified amount of time. The host will be removed from the in the specified amount of time. The default setting is **60** seconds.
- Alert Level** The current alert severity.
 - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.
- Process Cpu%** The amount of CPU used by processes, in percent.
- Threads** The total number of currently active threads.
- Peak Threads** The maximum number of active threads.

Daemon Threads

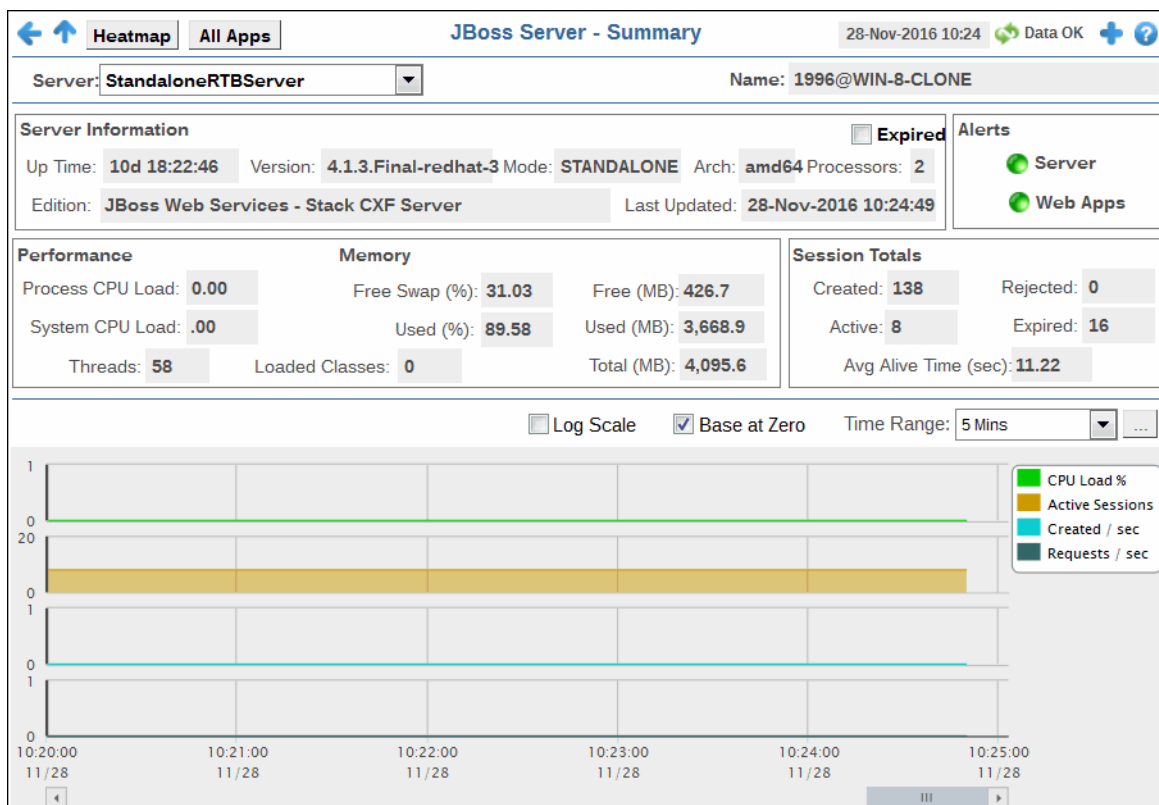
The total number of currently active daemon threads.

Started Threads

The total number of threads started since the server was last started.

Server Summary

Track the performance of one server, see detailed current information as well as historical trends. You can drill down to this display from the “[All Servers Heatmap](#)” and “[All Servers Table](#)” displays.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:




Server Select a server.

Name The name of the connection.

Server Information

Up Time	The amount of time that the server has been up and running.
Version	The version of the operating system.
Mode	The current server mode: <ul style="list-style-type: none"> • STANDALONE • DOMAIN MODE
Arch	The type of server architecture.
Processors	The number of processors on the server.
Expired	When checked, data has not been received from this server in the specified amount of time. The server will be removed from the per the specified amount of time. The default setting is 35 seconds.
Edition	Refers to the vendor edition.
Last Updated	The date and time of the last data update.

Alerts The current alert severity.

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

Performance

Process CPU Load	The amount of process CPU utilization, in megabytes.
System CPU Load	The amount of process CPU utilization, in megabytes.
Threads	The number of active threads.
Load Classes	The number of active load classes.

Memory

Free Swap %	The amount of free swap memory, in percent.
Free MB	The amount of free swap memory, in megabytes.
Used %	The amount of used memory, in percent.
Used MB	The amount of used memory, in megabytes.
Total MB	The memory sum total (Free MB + Used MB), in megabytes.

Session Totals

Created	The total number of sessions created since the server was restarted.
Rejected	The total number of sessions rejected since the server was restarted.
Active	The total number of currently active sessions.
Expired	The total number of currently expired sessions.
Average Alive Time sec	The average amount of time per session, in seconds.

Trend Graph

Traces metrics for the selected server.


Log Scale


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


Base at Zero



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

Time Range

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



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

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

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

CPU Load %

Traces the amount of process CPU utilization, in percent.

Active Sessions

Traces the number of active sessions.

Created/sec

Traces the number of active sessions per second.

Requests/sec

Traces the number of requests per second.

JBoss Applications

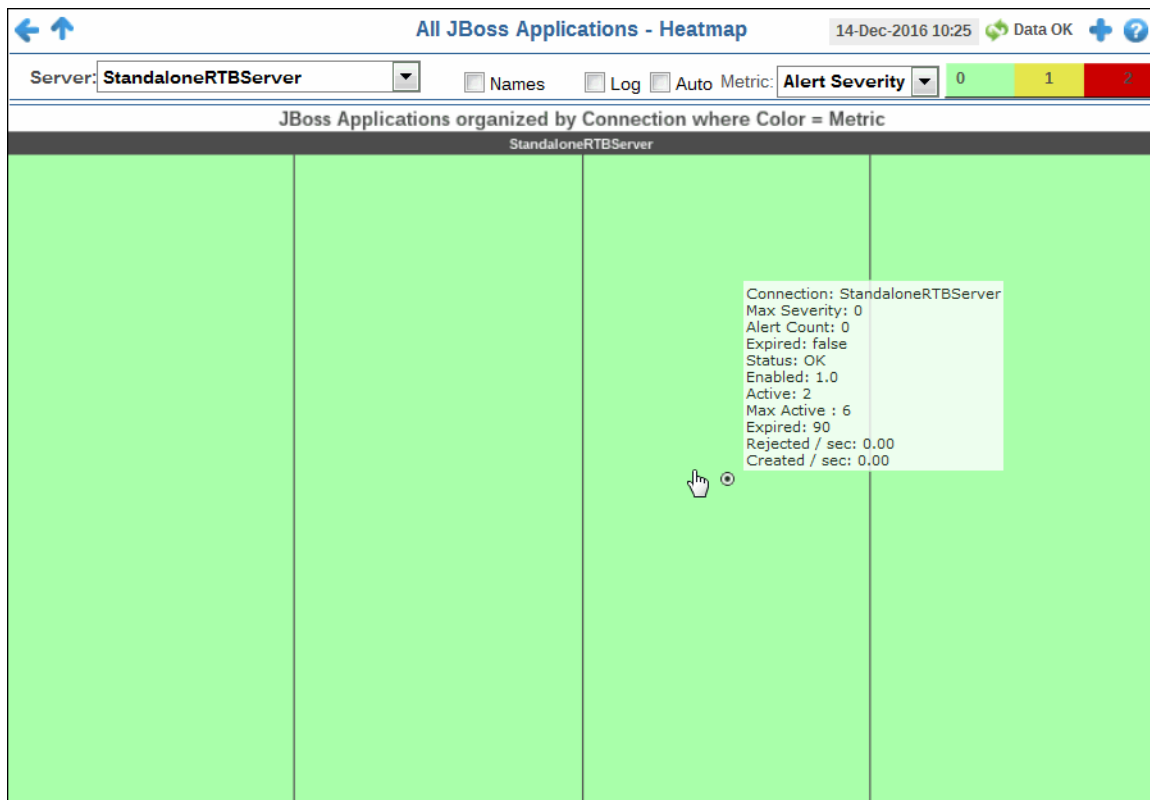
Displays in this View are:

- ["Applications Heatmap"](#)
- ["Applications Summary"](#)

Applications Heatmap

This heatmap shows the current status of all JBoss application connections. Use this to quickly identify the current session metrics for connections on one server or **All Servers**. Select a **Metric** from the drop-down menu, such as **Active Sessions** and **Average Alive Time**. By default, this display shows the **Alert Severity** metric.

Use the **Names** check-box ☒ to include or exclude labels in the heatmap, or mouse over a rectangle to see additional metrics for connections. Clicking one of the rectangles in the heatmap opens the “Applications Summary” display, which allows you to see additional details for the selected server.













Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

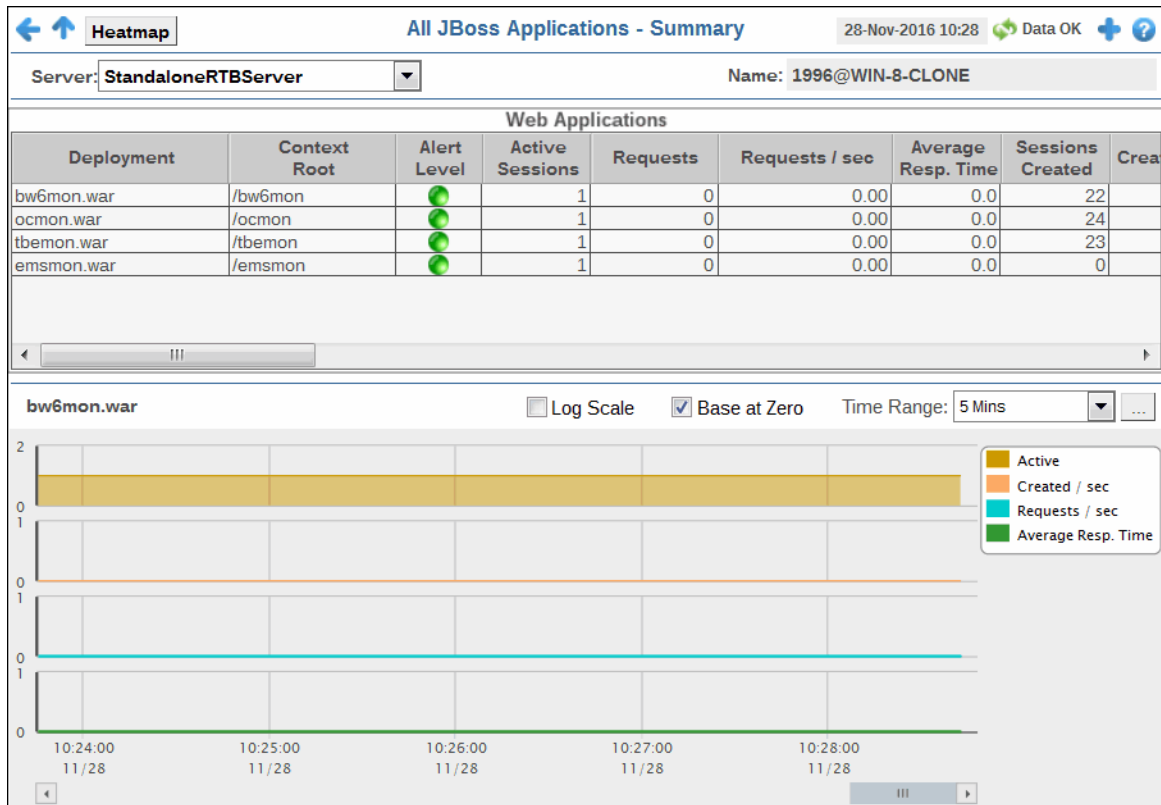
Fields and Data:

- Server** Choose a server to display.
- Names** Select this to include labels in the heatmap.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.

Metric	Choose a metric to view in the display.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning unacknowledged alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Active Sessions	<p>The number of currently active sessions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined JbossDeploymentActiveSessionsHigh alert threshold. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Max Active Count	<p>The total amount of active sessions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount. The middle value in the gradient bar indicates the average amount.</p> <p>The Auto flag does not have any impact on this metric.</p>
Avg Alive Time	<p>The average amount of time a session lasts, in seconds. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount recorded. The middle value in the gradient bar indicates the average amount.</p>
Created/sec	<p>The average number of sessions created per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number recorded. The middle value in the gradient bar indicates the average number.</p>
Rejected/sec	<p>The average number of rejected sessions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number recorded. The middle value in the gradient bar indicates the average number.</p>

Applications Summary

Track current and historical performance of web applications on one server. Select a server from the **Server:** drop-down menu. Each row in the table is a different application on the selected server. Click a table row to populate the trend graphs.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:




Server Choose a server to display.

Name The name of the connection.

Web Application Table


Each table row is a different web application on the selected server.

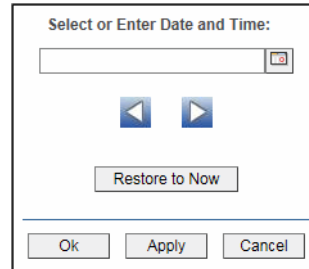
Deployment The name of the **.war** file for the application.


Context Root	The location of the .war file for the application.
Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Active Sessions	The number of currently active sessions for the application.
Requests/sec	The number of requests per second for the application.
Average Resp. Time	The average response time for the application, in seconds.
Sessions Created	The total number of sessions created for the application.
Created/sec	The number of sessions created per second for the application.
Sessions Rejected	The total number of sessions rejected for the application.
Rejected/sec	The number of sessions rejected per second for the application.
Max Active Sessions	The maximum number of simultaneously active sessions counted for the application.
Expired	When checked, the monitor has not received monitoring data for the application in the defined time interval.
Enabled	Indicates whether the application has been enabled for use in the JBoss Server.
Status	Indicates the application status.
Avg. Alive Time	The average amount of time, in seconds,
Max Alive Time	The average amount of time, in seconds,
Expired Sessions	The number of expired sessions for the application.
content	The application content.
Duplicated Session Ids	The number of sessions containing duplicated session IDs.
name	The name of the .war file for the application.
persistent	Indicates whether
Runtime Name	The name of the .war file for the application.
subsystem	The subsystem of the application.
Virtual Host	The name of the virtual machine that hosts the application.
Time Stamp	The date and time of the last data update.
Connection	The connection name.



Trend Graph

Traces metrics for the selected server.

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



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

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

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

- Active** The number of currently active sessions for the application.
- Created/sec** The number of sessions created per second for the application.
- Requests/sec** The number of requests per second for the application.
- Average Resp. Time** The average response time for the application, in seconds.

MongoDB

The following MongoDB Views (and their associated displays) can be found under **Components** tab > **Databases** > **MongoDB**:

- ["Mongo Instance View"](#)
- ["Mongo Database View"](#)
- ["Mongo Collection View"](#)

Mongo Instance View

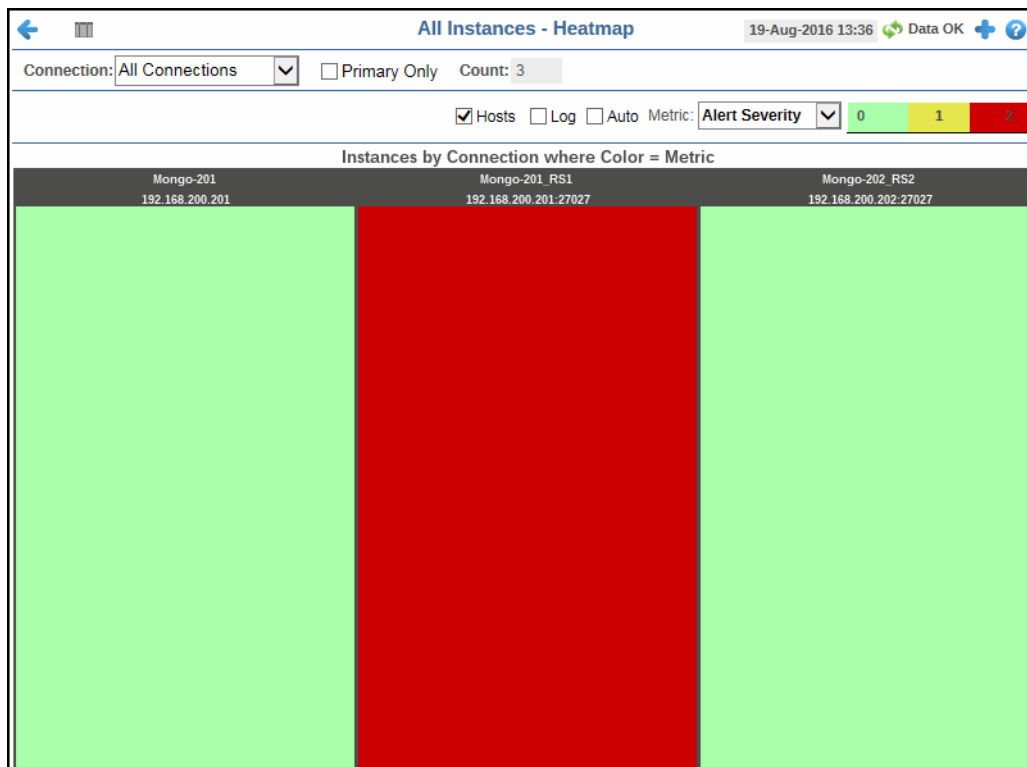
These displays present performance metrics and alert statuses for all MongoDB instances. The first two displays show different views of the same data:


- **"All Instances Heatmap"**: This heatmap shows status and alerts for all MongoDB instances.
- **"All Instances Table"**: This table shows all available utilization metrics for all MongoDB instances.
- **"Single Instance Summary"**: This summary enables you to view available utilization metrics for a single MongoDB instance.

All Instances Heatmap

View status and alerts of all MongoDB Instances. Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **Physical Memory**, **Open Cursors**, **Connections**, or **Databases**.

The heatmap is organized by host, each rectangle representing a connection. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the **"Single Instance Summary"** display and view metrics for a particular connection. You can toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the icon in the upper left-hand corner. Mouse-over rectangles to view more details about host performance and status.

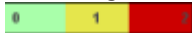









Fields and Data

This display includes:

- Connection** Select the connection from the drop down list for which you want to view data.
- Primary Only** Selecting this check box displays connections in the heatmap that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of active, inactive, and standby connections.
- Hosts** Select this check box to display the IP address of the host for each rectangle.
- Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
- Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
- Metric** Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the instances by host, where each rectangle represents an instance. Mouse-over any rectangle to display the current values of the metrics for the instance. Click on a rectangle to drill-down to the associated ["Single Instance Summary"](#) display for a detailed view of metrics for that particular instance.

Alert Severity	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar , where 2 is the greatest Alert Severity.</p> <p>2 -- Metrics that have exceeded their specified ALARMLEVEL threshold and have an Alert Severity value of 2 are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p>1 -- Metrics that have exceeded their specified WARNINGLEVEL threshold and have an Alert Severity value of 1 are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p>0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of 0 and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
Alert Count	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
Physical Memory	<p>The total amount of physical memory currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount of physical memory in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The Auto option does not impact this metric.</p>
Open Cursors	<p>The total number of open cursors in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of MongoInstanceOpenCursorsHigh, which is 2000. The middle value in the gradient bar indicates the middle value of the range (the default is 1000).</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Connections	<p>The total number of connections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The Auto option does not impact this metric.</p>
Databases	<p>The total number of databases in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of databases in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The Auto option does not impact this metric.</p>

All Instances Table

This display enables you to investigate detailed utilization metrics for all MongoDB Instances. The **All Instances Table** contains all metrics available for instances, including the number of current connections. Each row in the table contains data for a particular connection. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill-down to the “Single Instance Summary” display and view metrics for that particular instance. You can click the icon in the upper left-hand corner to toggle between the commonly accessed **Table** and **Heatmap** displays.

All Instances - Table

19-Aug-2016 13:38Data OK

Connection:

Mongo-201

☐ Primary Only

Count: 1

All Instances Table

Connection	Host	Conn Status	Alert Level	Alert Count	Average Back Flush	Connections	Databases	Designation In Set	
Mongo-201	192.168.200.201			0	0.00	27	2		L

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu

Table

open commonly accessed displays.

6,047

The number of items currently in the display.

Data OK

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

23-Mar-2017 12:04







Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

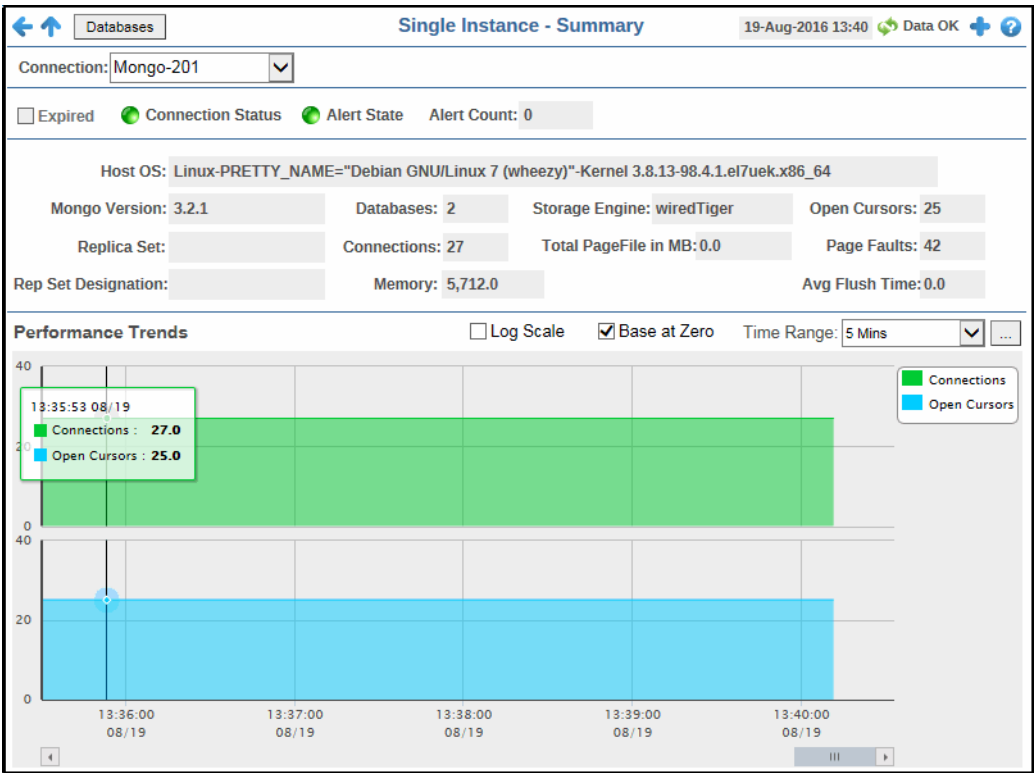
Connection Select the connection for which you want to view data, or select **All Connections** to view data for all connections.

Primary Only	Selecting this check box displays connections in the table that have Designation in Set (within a replica set) defined as Primary , as well as those connections that are not part of a replica set (do not have a defined Designation in Set). Those connections with Designation in Set defined as Secondary will not be displayed.
Count	The total number of connections displayed in the table.
Table	This table shows information for the selected connection(s). Click on a table row to drill-down to the "Single Instance Summary" display and view metrics for that particular server.
Connection	The name of the connection.
Host	The host name returned by MongoDB or the host provided by the user to use for connection if the host is not available.
Conn Status	The connection status of the Connection/Host.  -- The host is not connected.  -- The host is partially connected, which occurs when the connection has succeeded but the credentials given do not allow access to certain metrics.  -- The host is connected.
Alert Level	The current alert level.  -- One or more alerts have exceeded their specified ALARMLEVEL threshold.  -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.  -- No alerts have exceeded an alert threshold.
Alert Count	The total number of alerts for the connection.
Average Back Flush	The average time, in milliseconds, for each flush to disk, calculated by dividing the total number of milliseconds by the total number of flushes. Note: Background flushing information only appears for instances using the This metric only displays when the storage engine is MMAPv1 storage engine.
Connections	The number of connections coming in from the clients to the database server, including the current monitor session.
Databases	The number of databases being hosted by the instance.
Designation In Set	The designation of this member of the replica set (primary/secondary). This column will be empty if no replica set is configured, or set to unknown if there is no connection.
Host OS Version	The version of the operating system used by the host.
How Long As Primary	The amount of time the instance has been a primary instance. This field is only populated for primary instances.
MongoDB Version	The version number of the mongod instance.
Open Cursors	The total number of open cursors for the connection.
Ops Log Lag	The amount of time (in hours:minutes:seconds) in which the secondary instance is behind the primary instance. This field is only populated for secondary instances.
Ops Log Length	The length of the OpsLog collection, in bytes.

Page Faults	The number of page faults for the connection. MongoDB reports its triggered page faults as the total number of page faults in one second.
Physical Memory MB	The total amount of system memory (RAM), in megabytes.
ReplicaSet	The name of the replica set in which the mongod is a part of, if configured. This column will be empty if no replica set is configured, or set to unknown if there is no connection. All hosts in the replica set must have the same set name.
Storage Engine	The name of the current storage engine. The name can be either MMAPv1 or WiredTiger . WiredTiger is the default as of MongoDB version 3.2.
Total Page File MB	The total size of pagefile defined for the connection, in megabytes. This metric only displays when the storage engine is MMAPv1 .
Uptime	The amount of time since the instance was last started, shown in days, hours, and minutes (for example, 1d 23:43).
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > MongoDB > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Time Stamp	The date and time this row of data was last updated.

Single Instance Summary

Track utilization and performance metrics for specific instances.



Title Bar (possible features are):

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

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






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


Open the Alert Views - RTView Alerts Table display.


Fields and Data


This display includes:

- Connection** Select the connection for which you want to view data.
- Expired** This check box becomes automatically checked when the data has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.

Connection Status	<p>The connection status of the Connection/Host.</p> <ul style="list-style-type: none">  -- The host is not connected.  -- The host is partially connected, which occurs when the connection has succeeded but the credentials given do not allow access to certain metrics.  -- The host is connected.
Alert State	<p>The current alert level.</p> <ul style="list-style-type: none">  -- One or more alerts have exceeded their specified ALARMLEVEL threshold.  -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.  -- No alerts have exceeded an alert threshold.
Alert Count	The total number of alerts for the connection.
Host OS	The version of the operating system used by the host.
Mongo Version	The version number of the mongod instance.
Databases	The number of databases being hosted by the instance.
Storage Engine	The name of the current storage engine. The name can be either MMAPv1 or WiredTiger . WiredTiger is the default as of MongoDB version 3.2.
Open Cursors	The total number of open cursors for the connection.
Replica Set	The name of the replica set in which the mongod is a part of, if configured. This column will be empty if no replica set is configured, or set to unknown if there is no connection. All hosts in the replica set must have the same set name.
Connections	The number of connections coming in from the clients to the database server, including the current monitor session.
Total PageFile in MB	The total size of pagefile defined for the connection, in megabytes. This metric only displays when the storage engine is MMAPv1 .
Page Faults	The number of page faults for the connection. MongoDB reports its triggered page faults as the total number of page faults in one second.
Rep Set Designation	The designation of this member of the replica set (primary/secondary). This column will be empty if no replica set is configured, or set to unknown if there is no connection.
Memory	The total amount of system memory (RAM), in megabytes.
Avg Flush Time	<p>The average time, in milliseconds, for each flush to disk, calculated by dividing the total number of milliseconds by the total number of flushes.</p> <p>Note: Background flushing information only appears for instances using the MMAPv1 storage engine.</p>
Performance Trends Graph	<p>Shows connection and open cursor data for the connection.</p> <p>Connections -- Traces the total number of connections coming in from the clients.</p> <p>Open Cursors-- Traces the total number of open cursors on the connection.</p> <p>Log Scale This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.</p>

- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



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

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

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

Mongo Database View


These displays present detailed performance metrics and alert statuses for all databases (in a heatmap or a tabular format) or for an individual database.

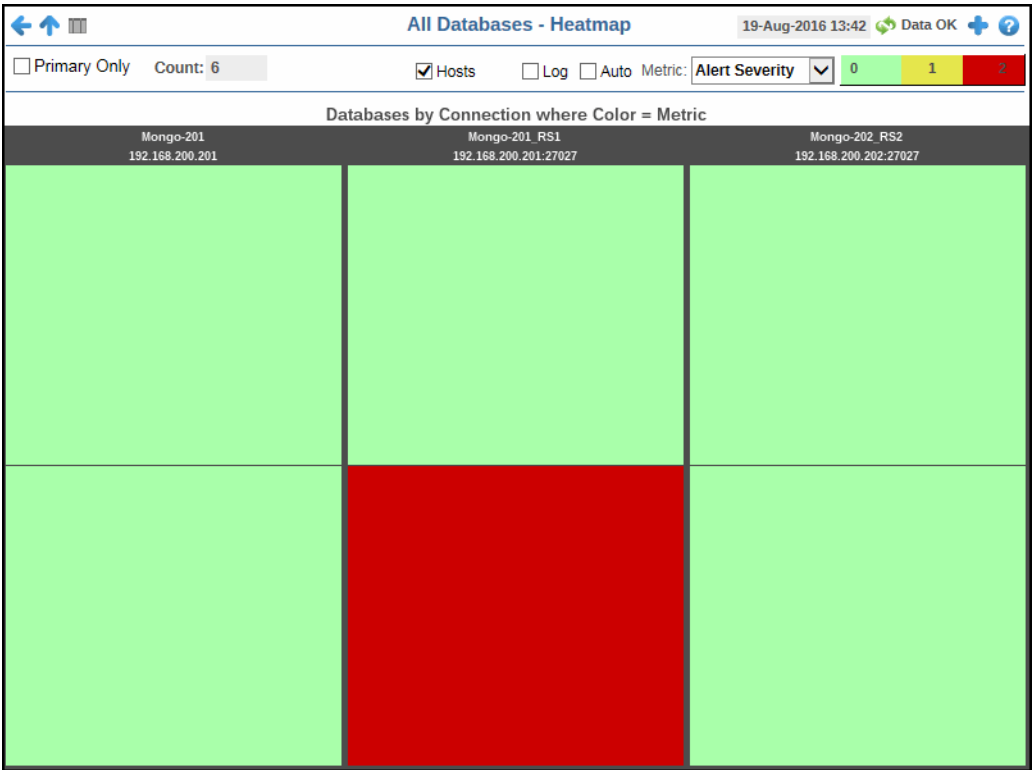
- **"All Databases Heatmap"**: Displays a heatmap view of alert states for all databases.
- **"All Databases Table"**: Displays a tabular view of all databases and their associated metrics for a single connection, or of all databases and their associated metrics for all connections.
- **"Database Summary"**: Displays metrics for a specific database.

Note: No database information will display in the heatmap, table, or summary displays if a connection cannot be established.








All Databases Heatmap


Track utilization and performance metrics for all databases in a heatmap format. Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **NumObjects** (number of objects), or **AvgObjectSize** (average object size).


The heatmap is organized so that each rectangle represents a database associated with a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each database or click on a rectangle to drill-down to the “[Database Summary](#)” display and view metrics for that particular Table database. You can click the table icon  in this display to navigate to the “[All Databases 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.
-   open commonly accessed displays.
-  6,047 The number of items currently in the display.

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

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

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

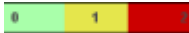

Fields and Data

This display includes:


- Primary Only**

Selecting this check box displays connections in the heatmap that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count**

The total number of active and inactive databases.

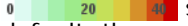
Hosts	Select this check box to display the name/IP address of the host for each rectangle.
Log	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
Auto	When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
Metric	<p>Select the metric driving the heatmap display. The default is Alert Severity. Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the databases by connection, where each rectangle represents a database. Mouse-over any rectangle to display the current values of the metrics for the database. Click on a rectangle to drill-down to the associated "Database Summary" display for a detailed view of metrics for that particular database.</p> <p>Alert Severity The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar , where 2 is the greatest Alert Severity.</p> <p>2 -- Metrics that have exceeded their specified ALARMLEVEL threshold and have an Alert Severity value of 2 are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p>1 -- Metrics that have exceeded their specified WARNINGLEVEL threshold and have an Alert Severity value of 1 are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p>0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of 0 and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p> <p>Alert Count The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

Collections

The total number of collections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum number of collections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

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


Data Size

The total size (in bytes) of the data in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **MongoDatabaseDataSizeHigh**, which is **100,000**. The middle value in the gradient bar indicates the middle value of the range (the default is **50,000**.).



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

All Databases Table

View details for all databases in a single connection, or view details for all databases in all connections.





All Databases - Table

19-Aug-2016 13:44  


Back to All Databases Table

Connection: Mongo-201 ☐ Primary Only Count: 2

All Databases Table								
Connection	Host	Name	Alert Level	Alert Count	Status UpDown	Collections	Data Size	File Size
Mongo-201	192.168.200.201	local		0	<input checked="" type="checkbox"/>	1	7,596	0
Mongo-201	192.168.200.201	meantest		0	<input checked="" type="checkbox"/>	1	390	0

<

>



Fields and Data

This display includes:

Connection	Select the connection for which you want to view data, or select All Connections to view data for all connections.
Primary Only	Selecting this check box displays connections in the table that have Designation in Set (within a replica set) defined as Primary , as well as those connections that are not part of a replica set (do not have a defined Designation in Set). Those connections with Designation in Set defined as Secondary will not be displayed.
Count	The total number of databases displayed in the table.
Table	This table shows information for the selected connection(s). Click on a table row to drill-down to the "Database Summary" display and view metrics for that particular server.
Connection	The name of the connection
Host	The host name returned by MongoDB or the host provided by the user to use for connection if the host is not available.
Name	The name of the database.
Alert Level	The current alert level. <div> ● -- One or more alerts have exceeded their specified ALARMLEVEL threshold. ● -- One or more alerts have exceeded their specified WARNINGLEVEL threshold. ● -- No alerts have exceeded an alert threshold. </div>
Alert Count	The total number of alerts for the database.
Status UpDown	When checked, signifies that the database is up and running.
Collections	The number of collections in the database.
Data Size	The total size, in bytes, of the data held in the database including the padding factor. The Data Size will not decrease when the document size decreases, but will decrease when documents are removed. Note: The scale argument affects this value.
File Size	The total size, in bytes, of the data files in the database. This value includes preallocated space as well as the padding factor, and only reflects the size of the data files in the database and not the size of the namespace file.
Free Space	The total free space remaining on the database (Storage Size minus Data Size).

Index Size

The total size, in bytes, of all indexes created on the database.

Note: The **scale** argument affects this value.

Storage Size

The total amount of space, in bytes, allocated to collections in this database for document storage. The Storage Size does not decrease when documents are removed or the size of the documents decrease.

Note: The **scale** argument affects this value.

Expired

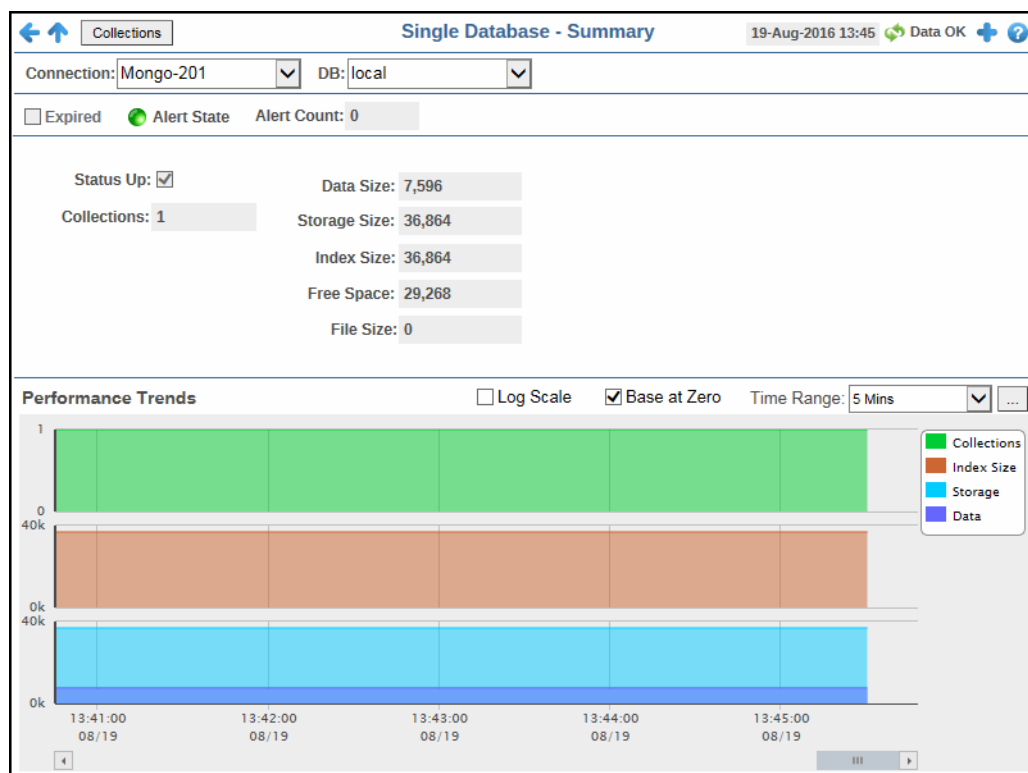
When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **MongoDB** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.


time_stamp

The date and time the data in this row was last updated.

Database Summary

View all available utilization and performance data for a specific database.





Fields and Data

This display includes:

Connection	Select the connection for which you want to view data.
DB	Select the database for which you want to view data.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > MongoDB > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Alert State	The current alert level. <div> ● -- One or more alerts have exceeded their specified ALARMLEVEL threshold. ● -- One or more alerts have exceeded their specified WARNINGLEVEL threshold. ● -- No alerts have exceeded an alert threshold. </div>
Alert Count	The total number of alerts for the database.
Status Up	When checked, signifies that the database is up and running.
Collections	The total number of collections in the database.
Data Size	The total size, in bytes, of the data held in the database including the padding factor. The Data Size will not decrease when the document size decreases, but will decrease when documents are removed. Note: The scale argument affects this value.
Storage Size	The total amount of space, in bytes, allocated to collections in this database for document storage. The Storage Size does not decrease when documents are removed or the size of the documents decrease. Note: The scale argument affects this value.
Index Size	The total size, in bytes, of all indexes created on the database. Note: The scale argument affects this value.
Free Space	The total free space remaining on the database (Storage Size minus Data Size).
File Size	The total size, in bytes, of the data files in the database. This value includes preallocated space as well as the padding factor, and only reflects the size of the data files in the database and not the size of the namespace file.

Performance Trends Graph

Shows connection and open cursor data for the connection.

Collections -- Traces the total number of collections in the database.

Index Size -- Traces the total size of indexes created on the database.


Storage -- Traces the total amount of space allocated to collections in the database.


Data -- Traces the total size of the data held in the database.

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

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

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



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

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

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


Mongo Collection View

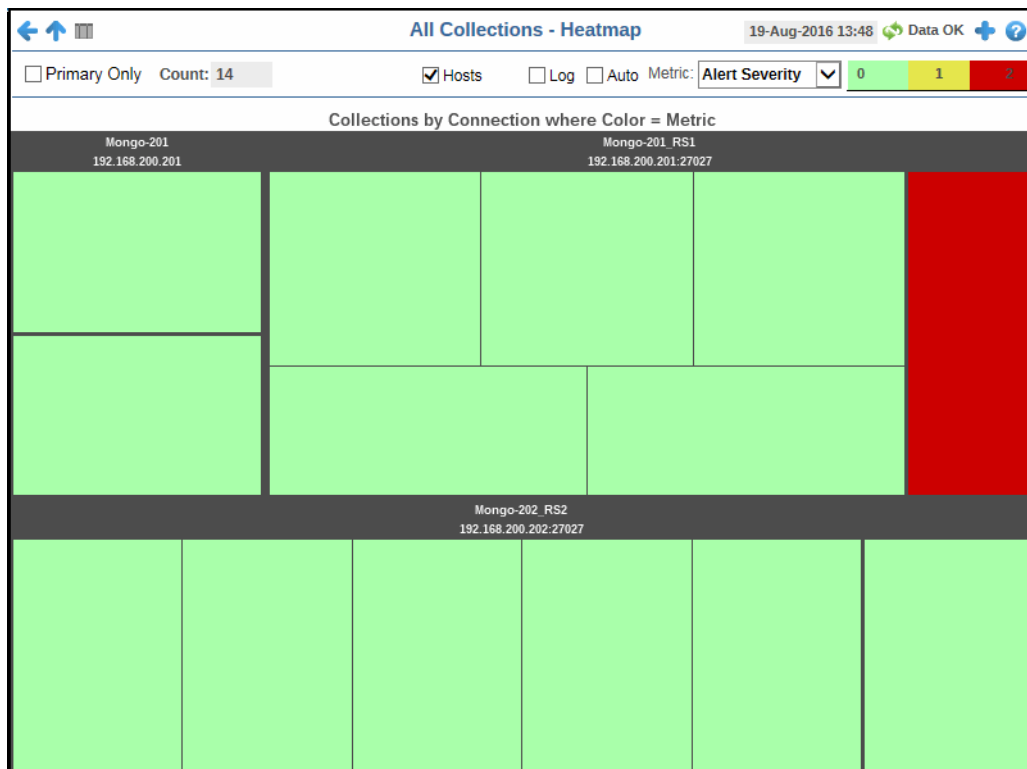
These displays present several views of performance metrics for collections. You can view heatmap or tabular views of all collections that exist in the connections in the [“All Collections Heatmap”](#) and [“All Collections Table”](#) displays, or you can view all details for a specific collection contained in a particular database in the [“Collection Summary”](#) display.


- [“All Collections Heatmap”](#): A heatmap representation that allows you to view performance and utilization metrics for all collections that exist in each of your connections.
- [“All Collections Table”](#): A tabular view that allows you to view performance and utilization metrics for all collections in a particular database, or for all collections on all databases.
- [“Collection Summary”](#): Shows detailed performance and utilization metrics and trends for a specified collection on a particular database.

All Collections Heatmap

This display provides a heatmap view of the status and alerts of all collections within each connection. Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **NumObjects** (number of objects), or **AvgObjectSize** (average object size).

The heatmap is organized so that each rectangle represents a collection contained within a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each collection or click on a rectangle to drill-down to the [“Collection Summary”](#) display and view metrics for that particular collection. You can click the table icon  in this display to navigate to the [“All Collections Table”](#) display.






Fields and Data

This display includes:

- Primary Only** Selecting this check box displays connections in the heatmap that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of collections.
- Hosts** Select this check box to display the names of the hosts in the heatmap.
- Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
- Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting **Auto** helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
- Metric** Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by connection, where each rectangle represents a collection. Mouse-over any rectangle to display the current values of the metrics for the collection. Click on a rectangle to drill-down to the associated ["Collection Summary"](#) display for a detailed view of metrics for that particular collection.

Alert Severity

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


2 -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

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


0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

Alert Count

The total number of alarm and warning alerts in a given item (index) associated with the rectangle.


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

NumObjects

The total number of objects or documents in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **MongoCollectionNumObjectsHigh**, which is **2000**. The middle value in the gradient bar indicates the middle value of the range (the default is **1000**).

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

AvgObjectSize

The average size (in bytes) of an object in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

All Collections Table

Track performance and utilization metrics for all collections on a single database, or for all connections on all databases.

All Collections - Table

19-Aug-2016 13:51 Data OK

Connection: Mongo-201 DB: local ☐ Primary Only Count: 1

Connection	Host	Database	Collection	Alert Level	Alert Count	Number of Objects	Average Object Size	
Mongo-201	192.168.200.201	local	startup_log		0	6	1,266	

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu

,

Table

 open commonly accessed displays.

6,047

 The number of items currently in the display.

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

23-Mar-2017 12:04

 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Fields and Data




This display includes:

- Connection

Select the connection for which you want to view collection data.
- DB


Select the database for which you want to view collection data, or select **All Databases** to view all collections for all databases.
- Primary Only

Selecting this check box displays connections in the table that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.

Count	The total number of collections found for the selected database(s).
All Collections Table	This table describes all topics on the selected server. Click a row to view metrics for a single topic in the "Collection Summary" display.
Connection	The name of the connection.
Host	The name of the host.
Database	The name of the database.
Collection	The name of the collection.
Alert Level	<p>The current alert level.</p> <p> -- One or more alerts have exceeded their specified ALARMLEVEL threshold.</p> <p> -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.</p> <p> -- No alerts have exceeded an alert threshold.</p>
Alert Count	The total number of alerts for the database.
Number of Objects	The total number of objects or documents in the collection.
Average Object Size	The average size, in bytes, of the objects in the collection.
Indexes	The total number of indexes in the collection.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > MongoDB > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
time_stamp	The date and time this row of data was last updated.

Collection Summary

Track performance and utilization metrics for a single collection on a single database.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

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


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


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


Fields and Data

This display includes:

- Connection** Select the connection for which you want to view collection data.
- DB** Select the database for which you want to view collection data, or select **All Databases** to view all collections for all databases.
- Collection** Select the connection for which you want to view data.

Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > MongoDB > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Alert State	The current alert level. <div> ● -- One or more alerts have exceeded their specified ALARMLEVEL threshold. ● -- One or more alerts have exceeded their specified WARNINGLEVEL threshold. ● -- No alerts have exceeded an alert threshold. </div>
Alert Count	The total number of alerts for the database.
# Objects	The total number of objects in the collection.
Avg Object Size	The average size, in bytes, of the objects in the collection.
Performance Trends Graph	Shows message data for the selected collection. # Objects -- Traces the total number of objects in the collection. Avg Object Size -- Traces the average size of objects in the collection.
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
Base at Zero	When this option is checked, zero is set as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click the  button.



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

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

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

MySQL

The MySQL Databases displays provide extensive visibility into the health and performance of the MySQL database. The following MySQL Database Views (and their associated displays) can be found under **Components** tab > **Databases** > **MySQL Database**:

- [“All MySQL Databases”](#)
- [“Single MySQL Database”](#)


All MySQL Databases

Displays in this View are:

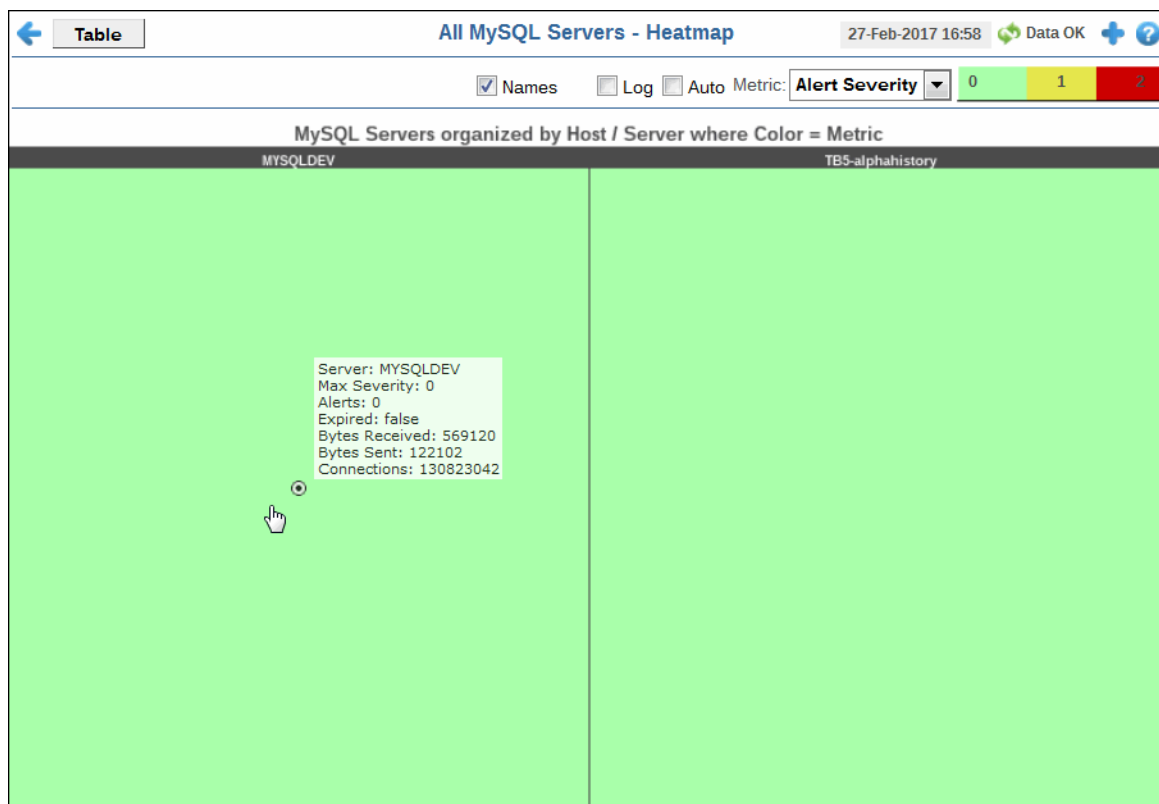
- [“All Servers Heatmap”](#): A heatmap view of all servers and their associated metrics.
- [“All Servers Table”](#): A tabular view of your servers and their associated metrics.

All Servers Heatmap

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



Each rectangle in the heatmap is a different server. Use the **Names** check-box ☒ to include or exclude labels in the heatmap, and mouse over a rectangle to see additional metrics for a server. Click a rectangle to open the “[Server Summary](#)” display and see additional details for the selected server.



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.
- 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.

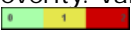
Fields and Data:




- Names** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
- Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.

Metric

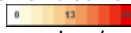
Choose a metric to view in the display. For details about the data, refer to vendor documentation.

Alert Severity


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

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


Alert Count

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

Received

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

Sent

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

Delayed Writes

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







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




Table Locks Waited

The total number of table locks waited. Values range from **0** to the alarm threshold specified for the **MysqlLocksWaited** alert. The middle value in the gradient bar indicates the average count:

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




Slow Queries

The total number of slow queries. Values range from **0** to the alarm threshold specified for the **MysqlSlowQueries**. The middle value in the gradient bar indicates the average count:

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




Slow Launch Threads

The total number of slow launch threads. Values range from **0** to the alarm threshold specified for the **MysqlSlowThreads**. The middle value in the gradient bar indicates the average count:

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

Qcache Low Mem Prunes


The total number of Qcache low memory prunes. Values range from **0** to the alarm threshold specified for the **MysqlQcacheLowMemPrunes**. The middle value in the gradient bar indicates the average count:


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

All Servers Table

This display provides a tabular view of the performance metrics shown in the “[All Servers Heatmap](#)” (alert level, alert count, bytes received, and so forth), as well as additional metrics (such as query information and uptime).

Each table row is a different server. Click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for a server in the “[Server Summary](#)” display.

All MySQL Servers Table									
Heatmap				27-Feb-2017 17:04		Data OK			
Server Name	Expired	Alert Level	Alert Count	Connected	Last Query	Avg Exec Time	Avg Process Time	Bytes Received	Bytes Sent
MYSQLDEV	<input type="checkbox"/>		0	<input checked="" type="checkbox"/>	OK	0.24	0.24	425,250	468



All MySQL Servers Table

Server Name	The name of the server.
Expired	<p>When checked, performance data about the server has not been received within the time specified (in seconds) in the \$mysqlRowExpirationTime field in the conf\rtvapm_mysqlmon.properties file. The \$mysqlRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the server. To view/edit the current values, modify the following lines in the .properties file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS # collector.sl.rtvview.sub=\$mssqlRowExpirationTime:120 collector.sl.rtvview.sub=\$mssqlRowExpirationTimeForDelete:0</pre> <p>In the example above, the Expired check box would be checked after 120 seconds, and the row would never be deleted. If \$mysqlRowExpirationTimeForDelete was set to 3600, then the row would be removed from the table after 3600 seconds.</p>
Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the server.
Connected	When checked, the server is connected.
Last Query	The status of the last query made:
Avg Exec Time	The average amount of execution time, in seconds.
Avg Process Time	The average amount of process time, in seconds.
Bytes Received	The total number of bytes received since the server was last started.
Connections	The total number of connections since the server was last started.
Delayed Writes	The total number of delayed writes.
Queries	The total number of queries.
Query Objects	The total number of query objects.
Slow Queries	The total number of slow queries.

Total Executions	The total number of executions.
Uptime	The amount of time since the server was last started, in seconds.
Concurrent	When checked, the database allows concurrent usage.
Enabled	When checked, the database is enabled for usage.
Timestamp	The data and time of the last data update.

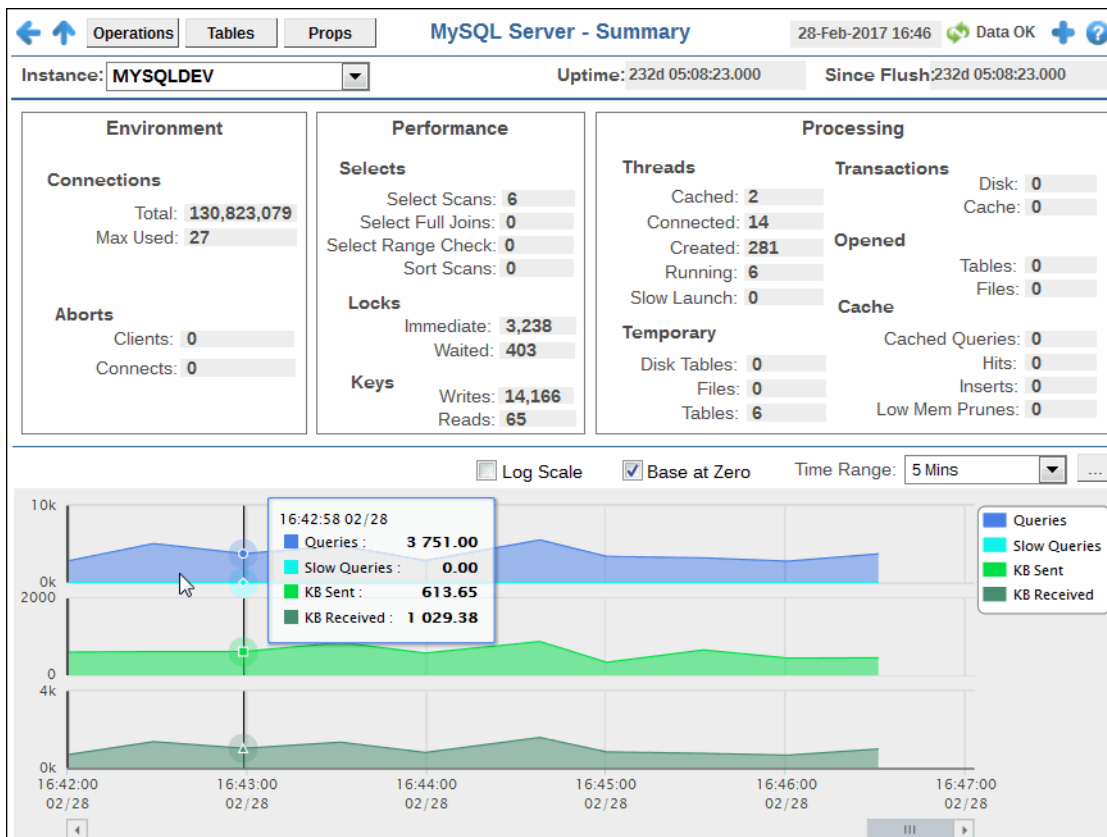
Single MySQL Database

Displays in this View are:

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

Server Summary

View connection, performance and processing details for a single MySQL database server, as well as trending data for the number of kilobytes received and queries. Choose an instance from the **Instance** drop-down menu. Mouse over the trend graph to see performance metrics with time stamps.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

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.

Filter By:

Instance: Select the instance for which you want to show data in the display.

Fields and Data: For details about the data in this display, please refer to vendor documentation.

Uptime The amount of time since the server was last started, in number of days, hours, minutes and seconds.

Since Flush The amount of time since the last flush, in number of days, hours, minutes and seconds.

Performance Trend Graph

Traces the following:

Queries: Traces the amount queries per second.

Slow Queries: Traces the amount of slow queries per second.

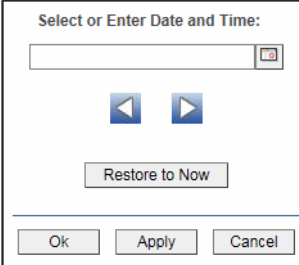
KB Sent: Traces the number of kilobytes sent per second.

KB Received: Traces the number of kilobytes received per second.


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



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

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



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

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

Servers Properties

View properties and property values for a single MySQL database server.

Choose an instance from the **Instance** drop-down menu. Each table row is a different property for the selected instance. Enter a search string in the **Property Filter** field to limit the number of table rows. Click a column header to sort column data in numerical or alphabetical order.

Property	Value
auto_increment_increment	1
auto_increment_offset	1
autocommit	ON
automatic_sp_privileges	ON
back_log	50
basedir	C:\Program Files\MySQL\MySQL Server 5.5\
big_tables	OFF
binlog_cache_size	32768
binlog_direct_non_transactional_updates	OFF
binlog_format	STATEMENT
binlog_stmt_cache_size	32768
bulk_insert_buffer_size	8388608
character_set_client	latin1
character_set_connection	latin1
character_set_database	latin1
character_set_filesystem	binary
character_set_results	latin1
character_set_server	latin1
character_set_system	utf8
character_sets_dir	C:\Program Files\MySQL\MySQL Server 5.5\share\charsets\

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.

Filter By:

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

Fields and Data:


Uptime The amount of time since the server was last started, in number of days, hours, minutes and seconds.

Property Filter: Enter a search string to filter the number of table rows.

Since Flush The amount of time since the last flush, in number of days, hours, minutes and seconds.

Servers Operations

View trending performance data for a single MySQL database server: **Inserts**, **Selects**, **Updates** and **Deletes**. Choose an instance from the **Instance** drop-down menu. Mouse over the trend graph to see performance metrics with time stamps.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Filter By:

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

Fields and Data:

Uptime The amount of time since the server was last started, in number of days, hours, minutes and seconds.

Property Filter: Enter a search string to filter the number of table rows.

Since Flush The amount of time since the last flush, in number of days, hours, minutes and seconds.

**Performance
Trend Graph**

Traces the following:

Inserts: Traces the number of inserts per second.


Selects: Traces the number of selects per second.


Updates: Traces the number of updates per second.


Deletes: Traces the number of deletes per second.



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

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

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



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

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

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


User Tables


View performance and utilization details for cache tables for a single MySQL database server. Each row is a different cache table. Choose an instance from the **Instance** drop-down menu. Click a column header to sort column data in numerical or alphabetical order.

MySQL Server - User Tables							
Instance: MYSQLDEV		Uptime: 231d 05:37:37.000		Since Flush: 231d 05:37:37.000			
Schema	Table	Row Count	Index Size	Data Size	Total Size	Data Free	Engine
alertdefs	alertlevels	0	1,024	0	1,024	0	MyISAM
alertdefs	audit_table	0	1,024	0	1,024	0	MyISAM
rtvhistory	\$bw6_activities_table	515,918	13,483,008	47,221,756	60,704,764	0	MyISAM
rtvhistory	\$bw6_activity_totals_table	56,463	1,383,424	6,107,932	7,491,356	0	MyISAM
rtvhistory	\$bw6_process_totals_app t	9,959	368,640	1,229,296	1,597,936	312,956	MyISAM
rtvhistory	\$bw6_process_totals_appnc	59,718	2,533,376	6,862,184	9,395,560	1,396,252	MyISAM
rtvhistory	\$bw6_process_totals_appsi	9,462	262,144	752,816	1,014,960	0	MyISAM
rtvhistory	\$bw6_process_totals_table	109,461	4,017,152	14,284,080	18,301,232	4,099,164	MyISAM
rtvhistory	\$bw6_processes_table	104,214	2,779,136	8,586,004	11,365,140	0	MyISAM
rtvhistory	bw6_activity_totals	226,128	4,355,072	20,718,016	25,073,088	0	MyISAM
rtvhistory	bw6_appnodes	39,409	764,928	2,597,056	3,361,984	0	MyISAM
rtvhistory	bw6_process_totals	94,395	1,859,584	7,650,588	9,510,172	0	MyISAM
rtvhistory	bw6_process_totals_app	10,979	216,064	777,800	993,864	0	MyISAM
rtvhistory	bw6_process_totals_appnoc	65,919	1,270,784	4,415,924	5,686,708	0	MyISAM
rtvhistory	bw6_process_totals_appsi	65,961	1,274,880	5,211,584	6,486,464	0	MyISAM
rtvhistory	bw6_processes	0	2,048	0	2,048	0	MyISAM
rtvhistory	bw_activities	3,520,325	35,879,936	330,112,152	365,992,088	0	MyISAM
rtvhistory	bw_activity_totals	1,202,835	38,381,568	158,427,548	196,809,116	692,108	MyISAM
rtvhistory	bw_engines	106,159	4,043,776	14,760,112	18,803,888	820,200	MyISAM
rtvhistory	bw_process_totals	78,638	4,087,808	15,453,984	19,541,792	5,266,124	MyISAM
rtvhistory	bw_processes	974,430	39,562,240	198,494,576	238,056,816	47,194,296	MyISAM
rtvhistory	bw_servers	30,982	1,239,040	2,314,796	3,553,836	231,836	MyISAM
rtvhistory	ems_admstats	8,309	158,720	187,194	345,914	12,705	MyISAM
rtvhistory	ems_compedsttotals	270,012	2,754,560	8,640,384	11,394,944	0	MyISAM
rtvhistory	ems_connections	534,561	5,451,776	39,159,128	44,610,904	0	MyISAM
rtvhistory	ems_consumers	2,018,788	20,578,304	87,000,188	117,677,492	0	MyISAM

Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- ⊕ Open an instance of this display in a new window.
- ⓘ Open the online help page for this display.
- Menu Table open commonly accessed displays.

 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.

Filter By:

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

Fields and Data: For details about the data in this display, please refer to vendor documentation.

Uptime The amount of time since the server was last started, in number of days, hours, minutes and seconds.

Property Filter: Enter a search string to filter the number of table rows.

Since Flush The amount of time since the last flush, in number of days, hours, minutes and seconds.

Table

Schema	The name of the database.
Table	The name of the table.
Row Count	The number of rows currently in the table.
Index Size	The size of the table indexes, in bytes.
Data Size	The size of the data stored in the table, in bytes (Total Size - Index Size = Data Size).
Total Size	The total size of the table, in bytes.
Data Free RX	The amount of available space that can be reclaimed to store new data, in bytes.
Engine	The storage engine handling the SQL operations.
Last Updated	The time of the last data update.

MS SQL Server

The Solution Package for Microsoft® SQL Server® includes high level heatmap and tabular displays as well as drilldown views to access real-time and historical performance metrics for each Microsoft SQL Server in your monitored services and applications.

With the Solution Package for Microsoft® SQL Server®, you are able to drill down from a high level alert at a business service or application health level into the supporting database infrastructure, to determine what is causing the alert and to take corrective action. This service-centric approach makes it easy for application support teams and Microsoft DBAs to prioritize incidents based on the impact to the business.

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

- **"All Servers View"**: The displays in this View allow you to view the current and historical metrics for all servers in a heatmap or tabular format.
- **"Single SQL Server View"**: The displays in this View allow you to view the metrics for a particular SQL database server.

All Servers View

These displays provide detailed data for all servers in a heatmap and tabular view. Displays in this View are:

- **"All Servers"**: A tabular view of your servers and their associated metrics.
- **"All Servers Heatmap"**: A heatmap view of all servers and their associated metrics.

All Servers

This table provides a view of all of your servers and their associated metric data including instance, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected server in the **"Server Summary"** display

← All SQL Servers - Table View 28-Oct-2016 09:29 Data OK + ?

Server Name	Instance	Alert Level	Alert Count	Expired	Active Connections	CPU Util	DeltaData Received (KB)	DeltaData Sent (KB)	DeltaIO Errors
WIN13\SQLEXPRESS	SQLEXPRESS		0	<input type="checkbox"/>	32	1	144	192	0
WIN-9J6VR0M8JK3\...	MSSQLS...		0	<input type="checkbox"/>	25	1	56	152	0
QA-SQL2016\MSSQ...	MSSQLS...		0	<input type="checkbox"/>	114	0	64	168	0
WIN17	MSSQLSE...		0	<input type="checkbox"/>	27	0	36	116	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.
- 6,047 The number of items currently in the display.

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




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

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

All SQL Servers Table

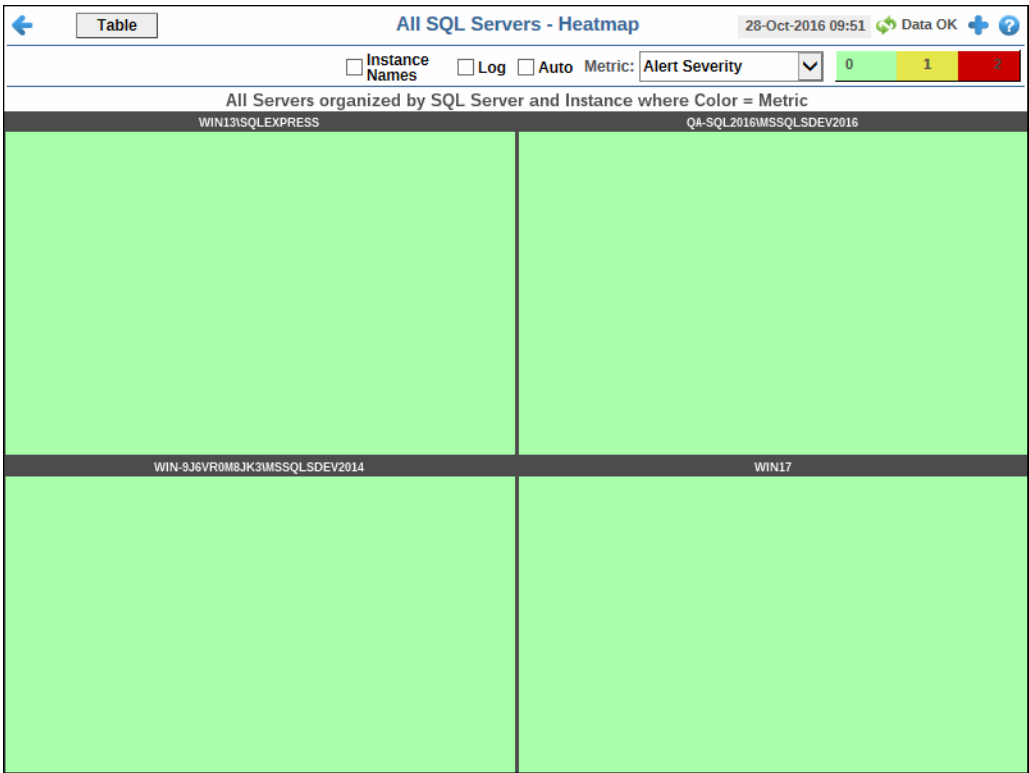
- Server Name** The name of the server.
- Instance** The name of the instance.

Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the host.
Expired	<p>When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Microsoft SQL Server > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p>
Active Connections	The number of currently active connections.
CPU Util	The CPU utilization percentage.*
Delta Data Received (KB)	The increase in the amount of data being received (from the previous polling period to the current polling period), in kilobytes.
Delta Data Sent (KB)	The increase in the amount of data being sent (from the previous polling period to the current polling period), in kilobytes.
Delta IO Errors	The increase in the amount of input/output errors (from the previous polling period to the current polling period).
Delta IO Reads	The increase in the amount of input/output reads operations (from the previous polling period to the current polling period).
Delta IO Writes	The increase in the amount of input/output write operations (from the previous polling period to the current polling period).
Delta Packet Errors	The increase in the amount of packet errors (from the previous polling period to the current polling period).
IO Busy (ms)	The time, in milliseconds, that the system has been busy due to Input/Output operations.*
Memory Used (%)	The percentage of memory used on the server.*
Memory Remaining (%)	The percentage of memory remaining on the server.*
Total DB Size (MB)	The size of the database, in megabytes.*
Server Edition	The version of the server.*
Product Level	The product level of the server.*
Product Version	The product's version number.*

All Servers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your servers for each available metric. You can view the servers in the heatmap based on the following metrics: the current alert severity, the current alert count, and the percentage of CPU used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Instance Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an engine. Clicking one of the rectangles in the heatmap opens the “[Server Summary](#)” display, which allows you to see additional details for the selected server.



Title Bar (possible features are):

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

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







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

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

Fields and Data:

Instance Names

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

Log	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. Note: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of critical and warning unacknowledged alerts in the engine. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
SQL CPU Utilization (%)	<p>The percentage of CPU used by the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of MssqlInstanceSqlCpuUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

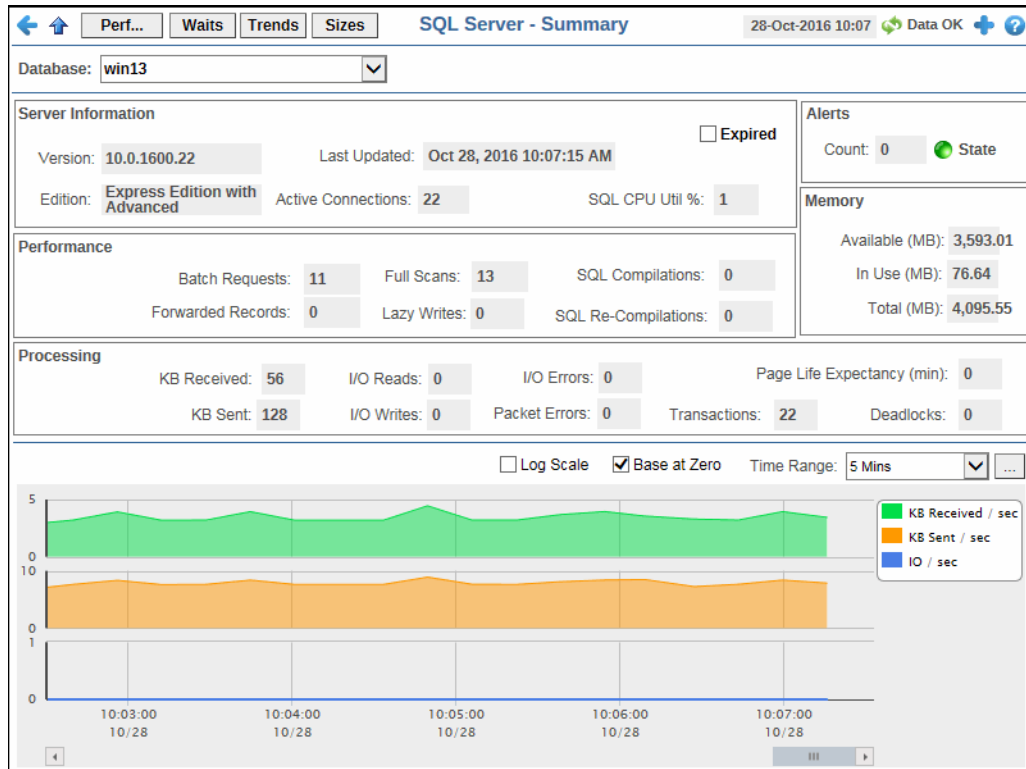
Single SQL Server View


Displays in this View are:

- **"Server Summary"**: Displays performance, processing, alerts, memory, and trend data for a particular database server.
- **"Database Details"**: Displays various database details as well as trending data for the page life expectancy.
- **"Wait Stats"**: Displays server wait time details in a table format for a particular database server.
- **"Performance Trends"**: This display allows you to view performance trend data for a particular SQL database server.
- **"Database/Table Sizes"**: Displays database and table sizes for a particular database server.

Server Summary

This display allows you to view connection and CPU utilization details, memory statistics, various performance and processing metrics, and trending data for the number of kilobytes received and sent as well as input/output details per second for a particular SQL database server.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

Note: The **Perf** button takes you to “Database Details”. The **Waits** button takes you to “Wait Stats”. The **Trends** button takes you to “Performance Trends”. The **Sizes** button takes you to “Database/ Table Sizes”.

Filter By:

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




Fields and Data:

Server Information

Version	The server’s version number
Last Updated	The date and time of the last data update.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Microsoft SQL Server > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Edition	The SQL Server’s edition.*
Active Connections	The number of active connections on the server.*
SQL CPU Util %	The percentage of CPU used by the server.*

Alerts

Count The total number of current alerts.

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

Memory

Available (MB)	The amount of memory currently available, in megabytes.*
In Use (MB)	The amount of memory currently in use, in megabytes.*
Total (MB)	The total amount of memory, in megabytes.*

Processing

KB Received	The number of kilobytes received.*
KB Sent	The number of kilobytes sent.*
I/O Reads	The number of input/output reads.*
I/O Writes	The number of input/output writes.*
I/O Errors	The number of input/output errors.*
Packet Errors	The number of errors involving incoming/outgoing packets.*
Page Life Expectancy (min)	The average number of minutes a page stays in the cache.*
Transactions	The number of transactions.*
Deadlocks	The number of deadlocks.*

Performance Trends Graph

Traces the following:

KB Received/sec -- traces the amount of kilobytes received per second.

KB Sent/sec -- traces the amount of kilobytes sent per second.


I/O /sec-- traces the average number of input/output operations per second.

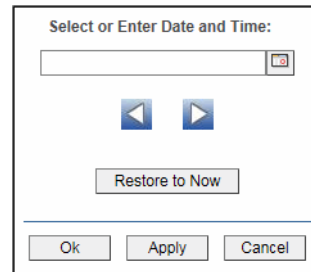
Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in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


Base at Zero



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

Time Range

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



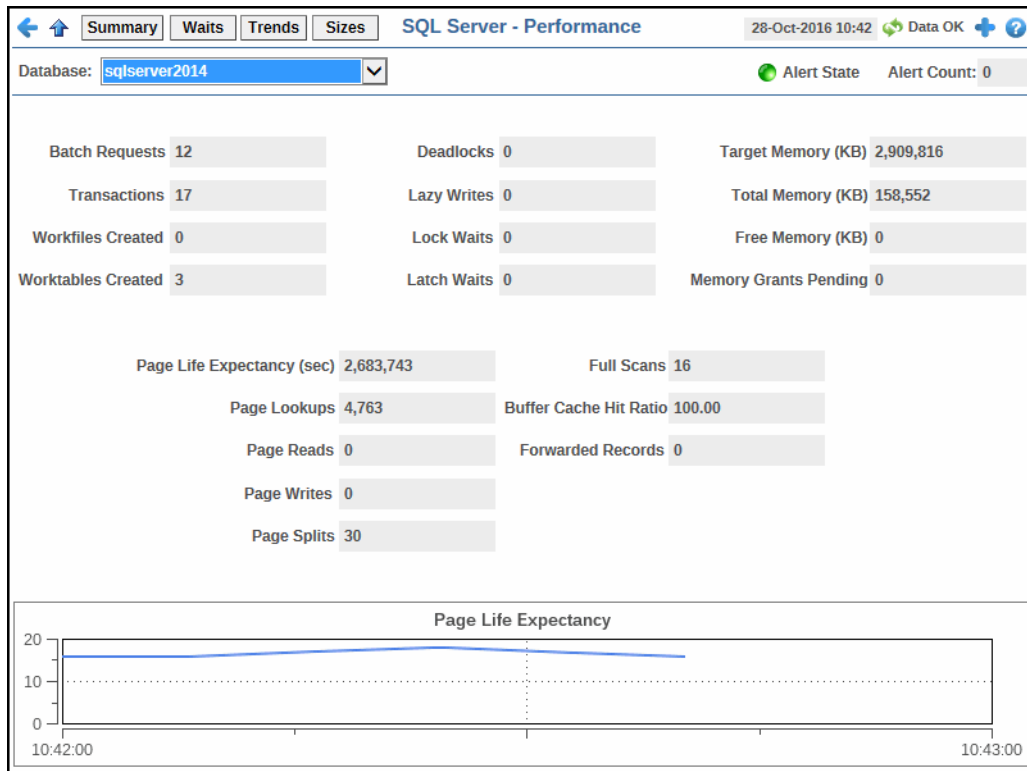
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Database Details

This display allows you to view various database details as well as trending data for the page life expectancy.



Title Bar (possible features are):

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

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

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

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

Note: The **Summary** button takes you to "Server Summary". The **Waits** button takes you to "Wait Stats". The **Trends** button takes you to "Performance Trends". The **Sizes** button takes you to "Database/Table Sizes".

Filter By:

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

Fields and Data:

Alert State	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of current alerts.
Batch Requests	The current number of batch requests.*
Transactions	The current number of transactions.*
Workfiles Created	The number of work files created.*
Worktables Created	The number of worktables created.*
Deadlocks	The current number of deadlocks occurring in the database.*
Lazy Writes	The number of times per second SQL Server relocates dirty pages from buffer pool (memory) to disk.*
Lock Waits	The number of lock requests that required the caller to wait.*
Latch Waits	The number of latch requests that required the caller to wait.*
Target Memory (KB)	The defined target server memory, which is the ideal amount of memory the server can consume, in kilobytes.*
Total Memory (KB)	The total amount of memory the server has committed using the memory manager, in kilobytes.*
Free Memory (KB)	The total amount of free memory, in kilobytes.*
Memory Grants Pending	The current number of processes waiting for a workspace memory grant.*
Page Life Expectancy (sec)	The average number of seconds a page stays in the cache.*
Page Lookups	The number of page lookups.*
Page Reads	The number of database pages read per second.*
Page Writes	The number of database pages written per second.*
Page Splits	The number of page splits that occur as a result of overflowing index pages.*
Full Scans	The number of full database scans.*
Buffer Cache Hit Ratio	The current buffer cache hit ratio, which is the total number of cache hits divided by the total number of cache lookups.*

- Forwarded Records**

The number of records fetched through forward record pointers.*
- Page Life Expectancy Trend Graph**

Traces the average length of time a page stays in the cache.*

Wait Stats

This display allows you to view server wait time details in a table format for a particular database server. You can drill-down and view the details for a particular container in the “[Server Summary](#)” display by clicking on a row in the resulting table.

Summary

Perf...

Trends

Sizes

SQL Server Waits - Table


28-Oct-2016 11:15 Data OK

Database: sqlserver2014

Alert State Alert Count: 0

Wait Types Sorted By Non-Zero Percentage Per Wait Category

Wait Category	Wait Type	Wait (sec)	Signal (sec)	Resource (s...	Percentage	Wait Count	Avg Wait (sec)
I/O		0.05	0.00	0.05	0.01	19.00	0.04
Lock		0.04	0.00	0.04	0.00	2.00	0.03
Non-I/O Page La...		0.00	0.00	0.00	0.00	115.00	0.00
Other		0.00	0.00	0.00	0.00	2.00	0.00
Transaction Log		0.00	0.00	0.00	0.00	2.00	0.00



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

Note: The **Summary** button takes you to "Server Summary". The **Perfs** button takes you to "Database Details". The **Trends** button takes you to "Performance Trends". The **Sizes** button takes you to "Database/Table Sizes".

Filter By:

The display includes these filtering options:

Database	Select the database for which you want to show data in the display.
Alert State	The current alert status. <div> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. </div> <div> ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. </div> <div> ● Green indicates that no metrics have exceeded their alert thresholds. </div>
Alert Count	Total number of alerts for the process.


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


Wait Category	The name of the wait category.*
Wait Type	The name of the wait type.*
Wait (sec)	The average length of the wait time, in seconds.*
Signal (sec)	When the thread is marked as runnable, this field displays the wait time, in seconds, that it takes to get into the running state.*
Resource (sec)	The length of time the thread spent in a suspended state waiting to acquire a resource, in seconds.*
Percentage	The percentage of time the thread spent in a wait state for this wait type.*
Wait Count	The number of lock requests that required the caller to wait.*
Avg Wait (sec)	The average wait time, in seconds.*

Avg Signal (sec)	The average wait signal time, in seconds.*
Avg Resource (sec)	The average length of time taken to acquire a resource, in seconds.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Microsoft SQL Server > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

Performance Trends

This display traces the current and historical percentage of CPU used by the MS SQL Server, the remainder CPU used in other operations, the amount of data sent, the amount of data received, the number of input/output operation reads, the number of input/output operation writes, the number of connections, the number of batch requests, the number of latch waits, and the number of lock waits for a particular SQL database server.





Note: The **Summary** button takes you to [“Server Summary”](#). The **Perf** button takes you to [“Database Details”](#). The **Waits** button takes you to [“Wait Stats”](#). The **Sizes** button takes you to [“Database/Table Sizes”](#).

Filter By:

The display might include these filtering options:

Database	Select the database for which you want to show data in the display.
Alert State	The current alert status. <div> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. </div> <div> ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. </div> <div> ● Green indicates that no metrics have exceeded their alert thresholds. </div>
Alert Count	Total number of alerts for the process.
Last Updated	The date and time the data in the display was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Microsoft SQL Server > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Performance Trends Graph	Traces the following: SQL CPU Used % -- traces percentage of CPU used by the MS SQL Server. CPU Other % -- traces the percentage of CPU used in other operations by the MS SQL Server. Data Sent (KB) -- traces the amount of data sent, in kilobytes. Data Received (KB) -- traces the amount of data received, in kilobytes. IO Reads -- traces the number of input/output operation reads. IO Writes -- traces the number of input/output operation writes. Connections -- traces the number of connections. Batch Requests -- traces the number of batch requests. Latch Waits -- traces the number of latch waits. Lock Waits -- traces the number of lock waits.


Log Scale

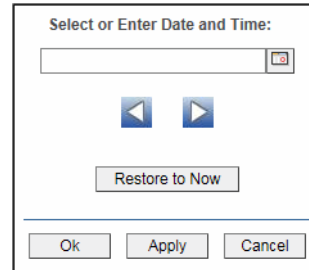
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log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


Base at Zero



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

Time Range

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



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

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

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

Database/Table Sizes

This display provides database and table size data for a particular SQL database server.

SummaryPerf...WaitsTrendsSQL Server - DB Table Sizes28-Oct-2016 13:29Data OKAlert StateAlert Count: 0

Database: sqlserver2014Last Updated: Oct 28, 2016 1:26:36 PMExpired

Database	Expired	Log Size (MB)	Row Size (MB)	State	Total Size (MB)
model	<input type="checkbox"/>	0.75	3.19	ONLINE	3.94
jparker	<input type="checkbox"/>	1	4	ONLINE	5
dcavazos	<input type="checkbox"/>	1	4	ONLINE	5
ReportServer\$MSSQL...	<input type="checkbox"/>	1.06	4.19	ONLINE	5.25
master	<input type="checkbox"/>	2	4	ONLINE	6
tempdb	<input type="checkbox"/>	0.5	8	ONLINE	8.5
ReportServer\$MSSQL...	<input type="checkbox"/>	7.56	5.19	ONLINE	12.75
msdb	<input type="checkbox"/>	19.63	17.13	ONLINE	36.75

Database	Schema	Table	Rows	Total (MB)	Used (MB)	Unused (MB)	Creation Date
jparker	dbo	EMS_TOPICTOTALS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_TOPICSEXT	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_TOPICS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_SERVERINFOEXT	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_SERVERINFO	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_ROUTES	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_ROUTECOUNTS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_QUEUETOTALS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_QUEUESEXT	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_QUEUES	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_PRODUCERS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_QUEUEABLES	0	0	0	0	Sep 27, 2016 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.

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

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

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




Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

Note: The **Summary** button takes you to "Server Summary". The **Perf** button takes you to "Database Details". The **Waits** button takes you to "Wait Stats". The **Trends** button takes you to "Performance Trends".

Filter By:

The display includes these filtering options:

Database	Select the database for which you want to show data in the display.
Alert State	The current alert status.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Total number of alerts for the process.
Last Updated	The date and time the data in the display was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Microsoft SQL Server > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

Database Sizes Table

Database	The name of the database.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Microsoft SQL Server > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Log Size (MB)	The size of the log, in megabytes.*
Row Size (MB)	The row size, in megabytes.*
State	The current state of the database.*
Total Size (MB)	The total size of the database, in megabytes.*

Table Sizes

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

Node.js

The following Node.js Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers**> **Node.js Servers**:

- **"Node/Master View"**: The displays in this View present detailed data for all node instances or for a particular node instance.
- **"Node Request View"**: The displays in this View allow you to view data pertaining to requests for a connection and a host, or view trending request data for a particular URL associated with a connection and a host.
- **"Process View"**: The displays in this View allow you to view the current and historical metrics for all node processes in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single node process.

Node/Master View

These displays provide detailed data for all node instances or for a particular node instance. Displays in this View are:

- **"Node Master Table"**: A tabular view of your connected and recently expired node instances and their associated metrics.
- **"Node Master Summary"**: Provides a way to view trending data for individual node processes.

Node Master Table

This table provides a view of all your connected (and recently expired) node instances and their associated metric data including host, connection, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected node in the **"Node Master Summary"** display.

The screenshot shows a web application window titled "All Node / Master Info - Table". At the top, there is a "Requests" button and a timestamp "19-Aug-2016 10:01" with a "Data OK" status and a help icon. Below the title bar, there are two dropdown menus: "Connection: SL_Cluster2" and "Host: TESTBED-29". A "Count: 1" label is also present. The main content area is a table titled "All Node / Masters Table". The table has the following columns: Connection, Hostname, Alert Level, Alert Count, Process ID, Uptime, CPU %, Request Count, Requests per sec, and Requests Delta. The first row of data shows: SL_Cluster2, TESTBED-29, a green alert level icon, 0, 8562, 17d 02:50, 1.5, 11,201,503, 0.0, and an empty Requests Delta cell. Below the table is a large empty space and a scrollbar at the bottom.

Connection	Hostname	Alert Level	Alert Count	Process ID	Uptime	CPU %	Request Count	Requests per sec	Requests Delta
SL_Cluster2	TESTBED-29		0	8562	17d 02:50	1.5	11,201,503	0.0	

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.
- 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.

Note: The **Requests** button takes you to "Node Requests Table".



Fields and Data:

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

Table:

Each row in the table is a different message router.

- Connection** The name of the connection.

Host Name	The name of the host.
Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of current alerts.
Process ID	The process id for the node instance.
Uptime	The amount of time the process has been running.
CPU %	The percentage of CPU used for the process.
Request Count	The total number of requests on the host.
Requests per sec	The average number of requests per second on the host.
Requests Delta	The total number of requests since the last data update.
Requests Mean Rate	The average number of requests for the server since monitoring was started.
Requests 1 Min Rate	The average number of requests for the last minute.
Requests 5 Min Rate	The average number of requests for the last 5 minutes.
Requests 15 min Rate	The average number of requests for the last 15 minutes.
Expired Workers	The number of expired workers on the host since the last data update.
Arch	The CPU architecture of the operating system on the server. Possible values are x64 , arm , and ia32 .
C-ares	The current version of C-ares running on the host.
Http Parser	The current version of the http parser running on the host.
ICU	The current version of ICU running on the host.
Modules	This number of modules found on the host.
Node Ver	The version of node.js running on the host.
Open SSL	The current version of OpenSSL running on the host.
Platform	The operating system's platform. Possible values, among others, are: darwin , linux , sunos , or win32 .
Release	The operating system's release number.
Type	The name of the operating system. Possible values, among others, are Linux on Linux, Darwin on OS X, and Windows_NT on Windows.
UV	The current version of uv running on the host.
V8	The current version of v8 running on the host.

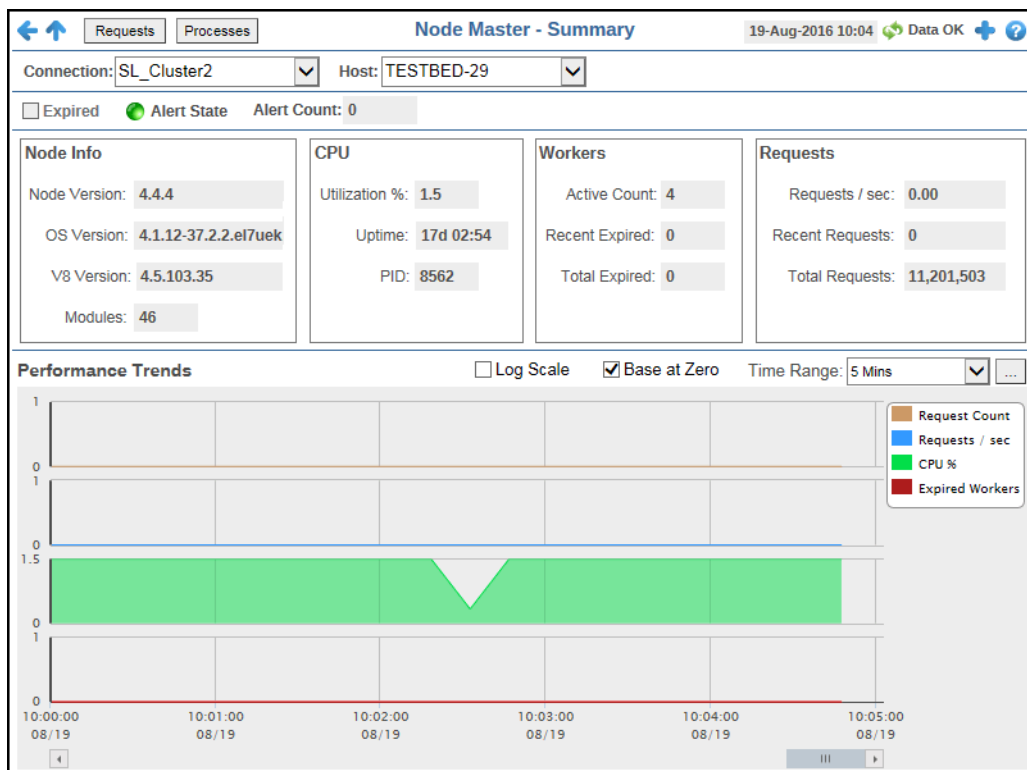
ZLib The current version of **ZLib** running on the host.


Expired When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **Node.js** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

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

Node Master Summary

This display allows you to view current CPU, worker, and request data as well as trending data for the number of requests, the number of requests per second, the percentage of CPU being used, and the number of recently expired workers on a particular host.





Note: The **Requests** button takes you to “[Node Requests Table](#)”. The **Processes** button takes you to “[All \(Node\) Processes Table](#)”.

Filter By:

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

Fields and Data:

- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Node.js** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Alert State** The current alert severity.
 - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.

Alert Count The total number of current alerts.

Node Info

- Node Version** The version of **node.js** running on the host.
- OS Version** The operating system's version number.
- V8 Version** The current version of **v8** running on the host.
- Modules** This number of modules found on the host.

CPU

- Utilization %** The percentage of memory used on the CPU.
- Uptime** The amount of time the process has been running.
- PID** The process id for the node instance.

Workers

- Active Count** The current number of active workers on the host.


Recent Expired	The number of expired workers on the host since the last data update.
Total Expired	The total number of expired workers on the host.

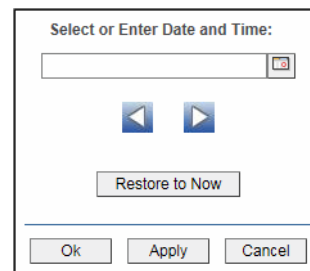
Requests


Requests / sec	The average number of requests per second on the host.
Recent Requests	The total number of requests since the last data update.
Total Requests	The total number of requests on the host.



Performance Trends Graph

Traces the following:

Request Count	-- traces the number of requests on the host.
Requests / sec	-- traces the number of requests/sec on the host.
CPU %	-- traces the percentage of CPU being used on the host.
Expired Workers	-- traces the number of expired workers on the host.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Select to use zero (0) as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar  .



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

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

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

Node Request View

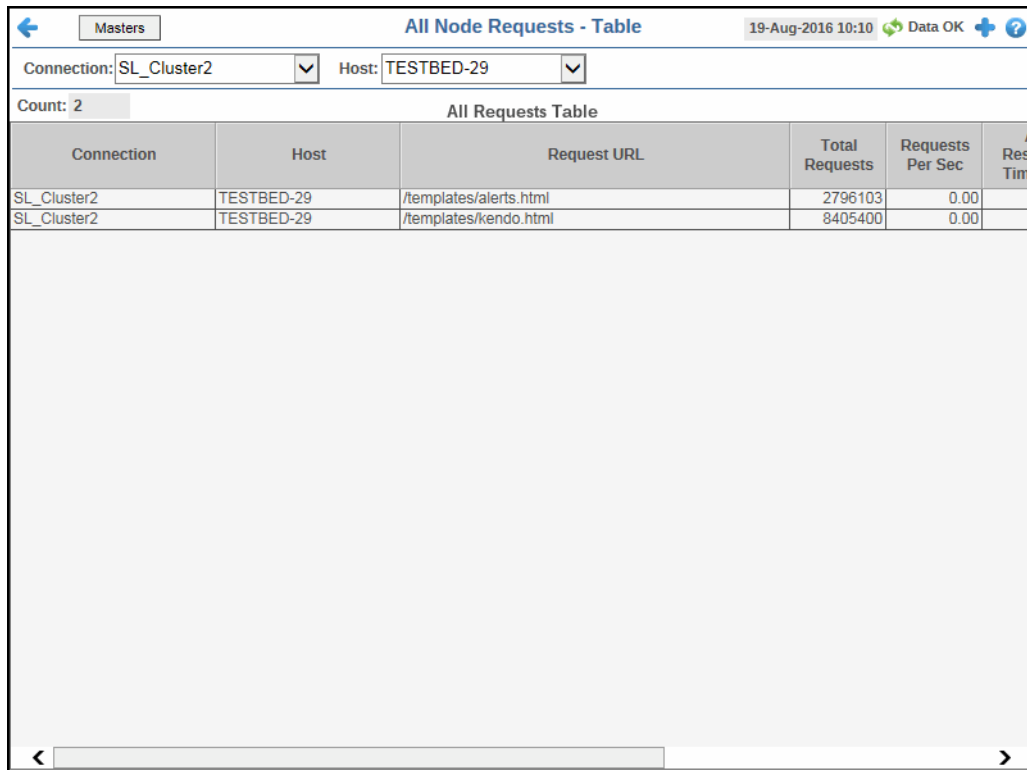
You can view data pertaining to requests for a connection and host, or view trending request data for a particular URL associated with a connection and a host. Displays in this View are:

- **"Node Requests Table"**: A tabular view of request data for one or all hosts on a particular connection.
- **"Node Request Summary"**: Allows you to view trending data (number of requests, number of requests per second, and average response time) for individual URLs by connection and host.


Node Requests Table

This display allows you to view request data for one or all hosts on a particular connection. You can view the request URL, total number of requests, number of requests per second, the average response time, and the number of recent requests for each host.

Drill-down and investigate by clicking a row to view details for the selected host in the **"Node Request Summary"** display.



Connection	Host	Request URL	Total Requests	Requests Per Sec	Avg Res Time
SL_Cluster2	TESTBED-29	/templates/alerts.html	2796103	0.00	
SL_Cluster2	TESTBED-29	/templates/kendo.html	8405400	0.00	



Note: The **Masters** button takes you to "[Node Master Table](#)".

Filter By:

The display might include these filtering options:

Connection	Select the connection for which you want to view data.
Host	Select the host for which you want to view data.

Fields and Data:

Count:	The total number of nodes (rows) in the table.
---------------	--

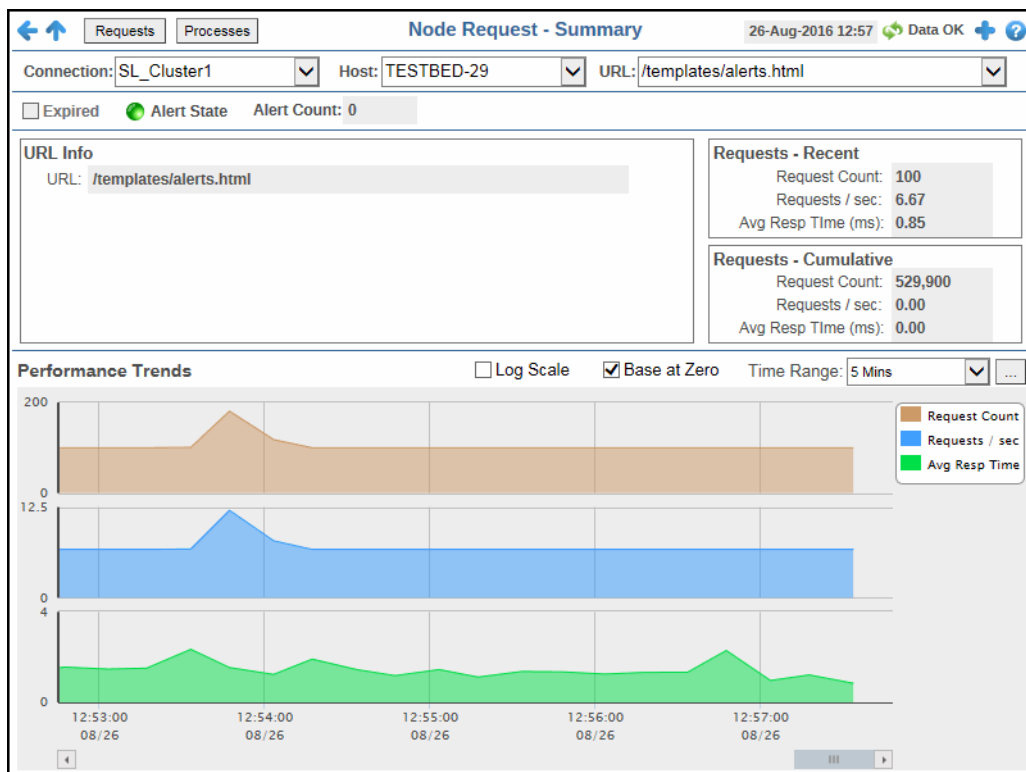
All Requests Table:

Column values describe the node and its associated requests.

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

Node Request Summary

This display allows you to view trending data (number of requests, number of requests per second, and average response time) for individual URLs by connection and host.



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.
- 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.

Note: The **Requests** button takes you to "Node Requests Table". The **Processes** button takes you to "All (Node) Processes Table".




Filter By:

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

Fields and Data:

Expired When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Node.js** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

Alert State The current alert severity.

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

Alert Count The total number of current alerts.

URL Info

URL The URL from which the requests originated.

Requests - Recent

Request Count The total number of requests based on the last query interval.

Requests / sec The rate of requests based on the last query interval.

Avg Resp Time (ms) The average response time (in milliseconds) based on the last query interval.

Requests - Cumulative

Request Count The total number of requests since the server was (re)started.

Requests / sec The rate of requests since the server was (re)started.

Avg Resp Time (ms) The average response time (in milliseconds) since the server was (re)started.

Performance Trends Graph

Traces the following:

Request Count -- traces the total number of requests.

Requests / sec -- traces the rate of requests.


Avg Resp Time-- traces the average response time.

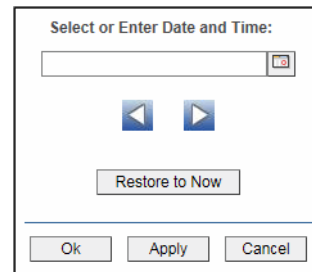
Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in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


Base at Zero



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

Time Range

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



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

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

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

Process View

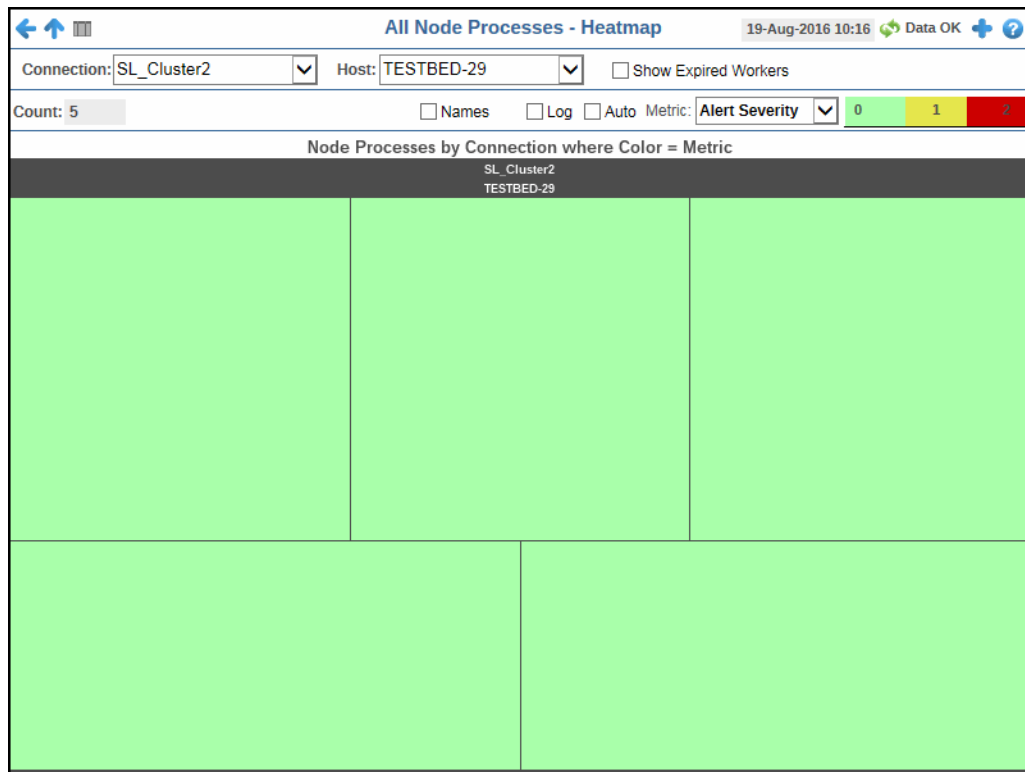
These displays allow you to view the current and historical metrics for all node processes in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single node process. Displays in this View are:

- **"All (Node) Processes Heatmap"**: A color-coded heatmap view of data for all node processes for a particular connection/host combination.
- **"All (Node) Processes Table"**: A tabular view of data for all node processes for a particular connection/host combination.
- **"(Node) Process Summary"**: This display allows you to view current and trending data for a single node process for a particular connection/host combination.

All (Node) Processes Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your node processes for each available metric. You can view the node processes in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, and the percentage of memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a node process. Clicking one of the rectangles in the heatmap opens the **"(Node) Process Summary"** display, which allows you to see additional details for the selected node process.



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.
- 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.

Fields and Data:

- Connection** Select the connection for which you want to show data in the display.
- Host** Select the host for which you want to show data in the display.
- Show Expired Workers** Select this check box to view expired workers in the heatmap.
- Count** Lists the total number of processes (rows) found using the search parameters.
- Names** Select this check box to display the names of the processes at the top of each rectangle in the heatmap.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Auto	<p>Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.</p> <p>Note: Some metrics auto-scale automatically, even when Auto is not selected.</p>
Metric	Choose a metric to view in the display.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning unacknowledged alerts in the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
CPU Used %	<p>The percentage of CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of NodeProcessCpuUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Memory Used %	<p>The total percentage of memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of NodeProcessMemUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

All (Node) Processes Table

This display allows you to view memory, heap memory, and latency data for all processes in a table format. You can drill-down and view the details in the “(Node) Process Summary” display for a specific process by clicking on a row in the resulting table.

← ↑ 🏠 All Node Processes - Table 19-Aug-2016 10:18 Data OK + ?

Connection: SL_Cluster2 Host: TESTBED-29 ☐ Show Expired Workers

Count: 5

All Node Processes Table

Connection	Hostname	Master / Worker	Alert Level	Alert Count	Uptime	CPU %	Process ID	Memory Used (KB)	Memory Used %
SL_Cluster2	TESTBED-29	1	🟢	0	17d 03:07	0.0	8588	29,905	0.8
SL_Cluster2	TESTBED-29	2	🟢	0	17d 03:07	0.0	8593	33,124	0.9
SL_Cluster2	TESTBED-29	3	🟢	0	17d 03:08	0.0	8599	33,042	0.9
SL_Cluster2	TESTBED-29	4	🟢	0	17d 03:08	0.0	8600	29,815	0.8
SL_Cluster2	TESTBED-29	Master	🟢	0	17d 03:07	1.5	8562	40,116	1.0

◀ ▶

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.

Filter By:

The display includes these filtering options:

- Connection** Select the connection for which you want to show data in the display.
- Host** Select the host for which you want to show data in the display.
- Show Expired Workers** Select this check box to view expired workers in the table.
- Count** Lists the total number of processes (rows) found using the search parameters.


Fields and Data:

Connection	The name of the connection.
Hostname	The name of the host.
Master / Worker	Displays whether the process is the Master process or, if the application is clustered, the worker ID.
Alert Level	<p>The current alert status.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Total number of alerts for the process.
Uptime	Lists the amount of time the process has been up and running.
CPU %	A decimal percentage describing how much the process utilizes the CPU.
Process ID	The process ID.
Memory Used (KB)	The used memory as a fraction of total system memory, in kilobytes.
Memory Used %	The percentage of total available memory used.
Memory RSS (KB)	The Resident Set Size, which is the portion of memory held in RAM (as opposed to swap or disk), in kilobytes.
Heap Total (KB)	The total amount of heap memory from which newly created objects will originate, in kilobytes.
Heap Free (KB)	The amount of memory remaining from which newly created objects will originate, in kilobytes.
Heap Used (KB)	The heap memory currently in use, in kilobytes.
Heap Used %	The percentage of heap memory currently being used.
Heap Avail (KB)	The v8 engine's total_available_size value, in kilobytes.
Heap Limit (KB)	The v8 engine's heap_size_limit value, in kilobytes.
Heap Total Executable (KB)	The v8 engine's total_heap_size_executable value, in kilobytes.
Latency p100	The number of microseconds that 100 percent of events were late in the previous 4 seconds.
Latency p99	The number of microseconds that 99 percent of events were late in the previous 4 seconds.
Latency p95	The number of microseconds that 95 percent of events were late in the previous 4 seconds.
Latency p90	The number of microseconds that 90 percent of events were late in the previous 4 seconds.
Latency p50	The number of microseconds that 50 percent of events were late in the previous 4 seconds.

- Lag** The average number of milliseconds a request has to wait in the Node's event queue before being processed. An excess lag means that the process is overloaded.
- time_stamp** The date and time the row data was last updated.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **Node.js** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

(Node) Process Summary

This display provides a view of the current and historical metrics for a single process. You can view the current information pertaining to a particular URL and various request data for the node process in the upper portion of the display. The trend graph in the bottom half of the display contains the current and historical number of requests, the number of requests per second, and the average response time for the node process.



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.

Filter By:

The display might include these filtering options:

Connection	Select the connection for which you want to show data in the display.
Host	Select the host for which you want to show data in the display.
Worker	Select the name of the worker to view. You can select from Master or any of the worker processes created by the Master. Worker processes are defined by numbers: 1 for the first worker process created by the Master , 2 for the second worker process created by the Master , and so on.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > Node.js > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Alert State	The current alert state of the process.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Lists the total number of alerts for the process.
CPU	
Utilization %	A decimal percentage describing how much the process utilizes the CPU.
Uptime	Lists the amount of time the process has been up and running.
PID	The process ID.
Event Loop Latency (microsecs)	
p100	The number of microseconds that 100 percent of events were late in the previous 4 seconds.
p99	The number of microseconds that 99 percent of events were late in the previous 4 seconds.
p95	The number of microseconds that 95 percent of events were late in the previous 4 seconds.
p90	The number of microseconds that 90 percent of events were late in the previous 4 seconds.
p50	The number of microseconds that 50 percent of events were late in the previous 4 seconds.
Memory	
Heap Utilization	The decimal percentage of utilized heap space.
Heap Used	The heap memory currently in use, in kilobytes.
Heap Total	The total amount of memory from which newly created objects can originate, in kilobytes.

Total Mem Used

The used memory as a fraction of total system memory, in kilobytes.

Mem RSS

Resident Set Size, which is the portion of memory held in RAM (as opposed to swap or disk), in kilobytes.

Performance Trends Graph

Traces the following:

CPU %-- traces the CPU utilization percentage.

Total Mem Used-- traces the amount of memory used.

Heap Total-- traces the total amount of available heap memory.

Heap Used-- traces the amount of used heap memory.

p95 -- traces the number of microseconds that 95 percent of events were late in the previous 4 seconds.

p90 -- traces number of microseconds that 90 percent of events were late in the previous 4 seconds.


Log Scale

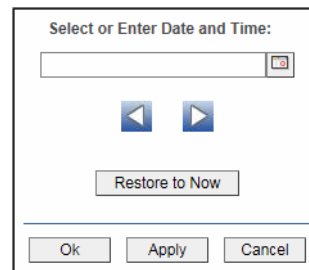
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log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

Base at Zero


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



Time Range

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



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon to its right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

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

RTView Host Agent

The Solution Package for RTView Host Agent monitors the health and performance of your physical servers. These predefined displays allow you to be alerted when hosts reach a critical condition. You can also see their performance impact on the technologies and applications they support. Metrics include CPU, memory and storage utilization, process resource consumption and network traffic load.

the Solution Package for RTView Host Agent is installed onto each host you wish to monitor.

RTView Host Agent displays provide extensive visibility into the health and performance of your hosts. The Solution Package for RTView Host Agent comes with RTView Enterprise Monitor. However, the displays are empty until you configure the Solution Package for RTView Host Agent.

RTView Host Agent Views (and their associated displays) can be found under **Components** tab > **Hosts/VMs** > **General Hosts**:

- ["All Hosts"](#)
- ["Single Host"](#)

All Hosts

These displays present performance data for monitored hosts. Use these displays to examine the state and performance of your hosts. The server displays include summary overviews and detail pages with historical trends.

To see your data in these displays you must install and configure the Solution Package for RTView Host Agent. Displays in this View are:

- ["All Hosts Heatmap"](#)
- ["All Hosts Table"](#)
- ["All Hosts Grid"](#)
- ["All Processes Table"](#)
- ["All Network Table"](#)
- ["All Storage Table"](#)
- ["Host Summary"](#)

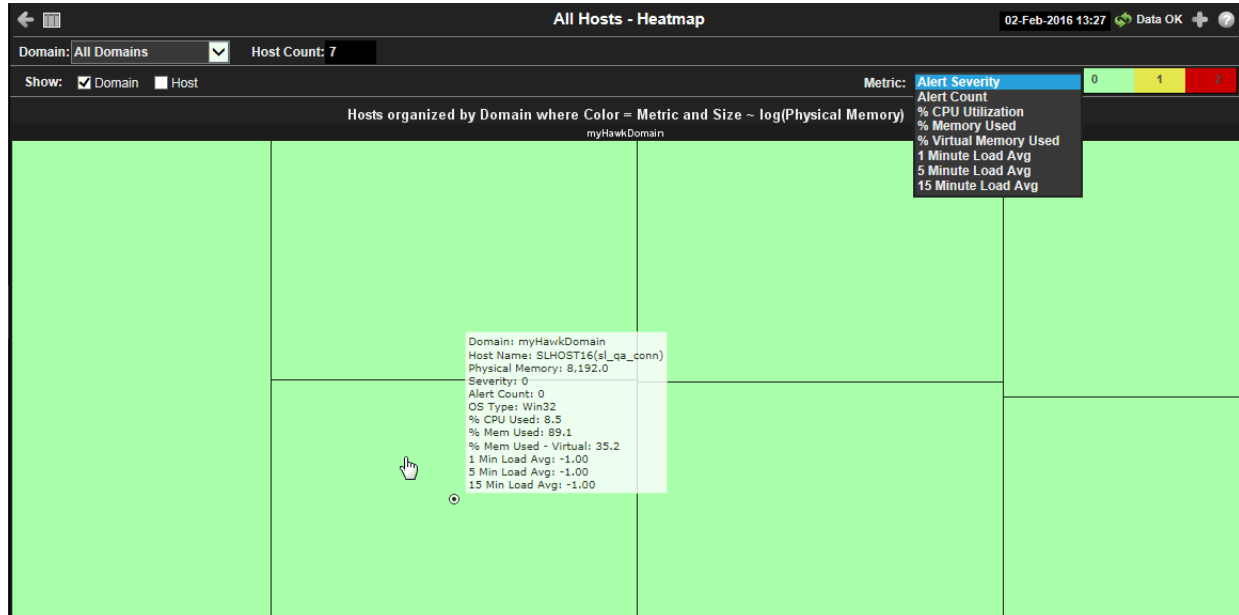
All Hosts Heatmap

View the most critical alert states pertaining to your hosts. Use this display to quickly identify hosts with critical alerts.

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

Choose a domain or **All Domains** from the **Domain** drop-down menu to filter data shown in the display. Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate a host by clicking a rectangle in the heatmap to view details in the **Host Summary** display.



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.
- 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.

Filter By:

The display might include these filtering options:

Domain: Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.









Fields and Data:

Host Count: The total number of hosts currently shown in the display.

Show: **Domain** When selected, includes the Domain name in the display.

Host When selected, includes the Host name in the display.

Metric Choose a metric to view in the display.

Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
% CPU Utilization	<p>The percent of CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
% Memory Used	<p>The percent of memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
% Virtual Memory Used	<p>The percent of virtual memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
1 Minute Load Avg	The average number of processes running over 1 minute.
5 Minute Load Avg	The average number of processes running over 5 minutes.
15 Minute Load Avg	The average number of processes running over 15 minutes.

All Hosts Table

View host utilization data in a tabular format. Use this display to see all available data for this View.

Each row in the table is a different host. Choose a domain or **All Domains** from the **Domain** drop-down menu. Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the **Host Summary** display.

← All Hosts - Table View 02-Feb-2016 13:37 Data OK + ?

Domain: All Domains

Host Count: 7

		Host CPU Stats													
Domain	Host Name	Expired	Severity	Alert Count	Uptime	% CPU User	% CPU System	% CPU Idle	Memory Used	Memory Total	Memory Used %	Swap Used	Swap Total	Swap Used %	Virtual Us
myHawkDomain	SLHOST16(sl_amx)			0	120d 02:24	8.27	-1.00	91.73	7,309	8,192	89.2	1,581	8,192	19.3	
myHawkDomain	SLHOST16(sl_qa_conn)			0	120d 02:21	8.37	-1.00	91.63	7,306	8,192	89.2	1,581	8,192	19.3	
myHawkDomain	SLHOST17(sl_amx)			0	120d 02:17	0.71	-1.00	99.29	4,875	8,192	59.5	180	8,192	2.2	
myHawkDomain	SLHOST21(dev)			0	120d 04:40	3.03	-1.00	96.97	14,339	16,384	87.5	2,975	16,384	18.2	
myHawkDomain	SLHOST22(sl_qa_conn)			0	54d 02:41	0.00	0.00	100.00	2,578	7,824	32.9	0	9,999	0.0	
myHawkDomain	SLHOST5(domain5)			0	0d 13:34	17.19	-1.00	82.81	2,313	4,096	56.5	26	4,096	0.6	
myHawkDomain	SLHOST6(domain6)			0	0d 13:36	1.87	-1.00	98.13	2,137	3,072	69.6	27	3,072	0.9	

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.

Filter By:

The display might include these filtering options:

Domain: Choose a domain to show data for in the display.

Fields and Data:

Host Count: The total number of hosts in the table.





Table:

Each row in the table is a different host.

Domain The domain in which the host resides. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.

Host Name The name of the host.

Expired When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is **60** seconds.

Severity	<p>The maximum level of alerts in the row. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics exceeded their alert thresholds.
Alert Count	The total number of active alerts associated with the host.
Uptime	<p>The amount of time the application has been running, in the following format: 0d 00:00 <days>d <hours>:<minutes>:<seconds> For example: 10d 08:41:38</p>
% CPU Used	The amount of CPU used, in percent.
% CPU System	The amount of CPU used, in percent.
% CPU Idle	The amount of CPU not used, in percent.
Memory Used	The amount of memory, in megabytes, currently used.
Memory Total	The total amount of memory, in megabytes.
Memory Used%	The amount of memory used, in percent.
Swap Used	The amount of swap space, in megabytes, currently used.
Swap Total	The total amount of swap space, in megabytes.
Swap Used %	The amount of swap space used, in percent.
Virtual Mem(ory) Used	The amount of virtual memory currently used, in megabytes.
Virtual Mem(ory) Total	The total amount of virtual memory, in megabytes.
Virtual Mem(ory) Used%	The amount of virtual memory used, in percent.
Load Avg 1 Minute	The average number of processes running over 1 minute.
Load Avg 5 Minute	The average number of processes running over 5 minutes.
Load Avg 15 Minute	The average number of processes running over 15 minutes.
OS Type	The type of operating system (for example, Linux, HP-UX, Windows 2003).
OS Description	The name of the operating system.
OS Version	The operating system version.
CPU Model	The CPU model.
# CPUs	The number of node connections.

Agent Type	The type of agent from which the data was collected: HOSTMON (a SL Host Agent), Hawk , WMI or SNMP .
Agent Class	The specific version of the agent software.
Source	The name of the SL Data Server where the host data was collected.
Timestamp	The date and time the data was last updated.

All Hosts Grid

This grid provides a list view of utilization metrics for all hosts. Use this display to track and view in parallel the general performance of your hosts. Drill down and investigate by clicking a host to view details in the **Host Summary** display.



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.

Filter By:

The display might include these filtering options:

Domain:	Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
Host Count	Displays the number of hosts (including expired hosts) listed in the display.
Time Range:	Choose a time range to show data for in the display. Options are: All Data, 2 Mins, 5 Mins, 20 Mins, 1 Hour, 2 Hours, 4 Hours, 8 Hours, 24 Hours, 2 Days and 7 Days.

Grid


Utilization data shown for hosts in the selected domain.

Host Name	The name of the host.	
OS Type	The name of the operating system.	
Uptime	The amount of time (days, hours, seconds) the operating system has been running.	
Phys Mem	The amount of physical memory used, in megabytes.	
Virtual Mem	The amount of virtual memory used, in megabytes.	
Load Avg	1	The average number of processes running over 1 minute.
	5	The average number of processes running over 5 minutes.
	15	The average number of processes running over 15 minutes.
CPU Usage	The bar graph shows the amount of CPU currently used.	
VMem Usage	The bar graph shows the amount of virtual memory currently used.	
Trend Graphs	CPU	Traces the amount of CPU currently used.
	VM Usage	Traces the amount of virtual memory currently used.
	Rx KB/s	Traces the amount data currently being received per second.
	Tx KB/s	Traces the amount data currently being transmitted per second.

All Processes Table

View host utilization data in a tabular format. Use this display to see all available data for this View. Each row in the table is a different host. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the **Host Summary** display.

All Processes - Table View												02-Feb-2016 13:42		Data OK	
Domain: All Domains		Host: All Hosts													
Process Count: 687				Host Processes											
Domain	Host Name	Expired	PID	User	Process Name	CPU %	Start Time	Memory Used	Memory Resident	Memory Shared	Page Fault				
myHawkDot	SLHOST16(sl_amx)		4	<ACCESS DENIE	System	0.02	01-May-2014 23:18:11	17,056	-1	-1	465,4				
myHawkDot	SLHOST16(sl_amx)		376	NT AUTHORITY\	smss.exe	0.00	01-May-2014 23:18:11	504	-1	-1	1,8				
myHawkDot	SLHOST16(sl_amx)		540	NT AUTHORITY\	csrss.exe	0.00	01-May-2014 23:18:16	2,472	-1	-1	12,089				
myHawkDot	SLHOST16(sl_amx)		628	NT AUTHORITY\	wininit.exe	0.00	01-May-2014 23:18:17	172	-1	-1	1,9				
myHawkDot	SLHOST16(sl_amx)		648	NT AUTHORITY\	csrss.exe	0.00	01-May-2014 23:18:17	216	-1	-1	11,3				
myHawkDot	SLHOST16(sl_amx)		692	NT AUTHORITY\	services.exe	0.01	01-May-2014 23:18:17	5,736	-1	-1	14,404				
myHawkDot	SLHOST16(sl_amx)		708	NT AUTHORITY\	lsass.exe	0.02	01-May-2014 23:18:17	9,576	-1	-1	1,273				
myHawkDot	SLHOST16(sl_amx)		716	NT AUTHORITY\	lsmon.exe	0.00	01-May-2014 23:18:17	3,500	-1	-1	1,030				
myHawkDot	SLHOST16(sl_amx)		800	NT AUTHORITY\	winlogon.exe	0.00	01-May-2014 23:18:17	172	-1	-1	3,6				
myHawkDot	SLHOST16(sl_amx)		864	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:20	3,660	-1	-1	1,496				
myHawkDot	SLHOST16(sl_amx)		416	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:20	4,376	-1	-1	2,872				
myHawkDot	SLHOST16(sl_amx)		472	NT AUTHORITY\	LogonUI.exe	0.00	01-May-2014 23:18:21	2,960	-1	-1	164,7				
myHawkDot	SLHOST16(sl_amx)		640	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:21	13,756	-1	-1	111,69				
myHawkDot	SLHOST16(sl_amx)		548	NT AUTHORITY\	svchost.exe	0.05	01-May-2014 23:18:21	121,608	-1	-1	111,2				
myHawkDot	SLHOST16(sl_amx)		1048	NT AUTHORITY\	svchost.exe	0.28	01-May-2014 23:18:21	26,108	-1	-1	1,605				
myHawkDot	SLHOST16(sl_amx)		1220	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:22	7,336	-1	-1	2,716				
myHawkDot	SLHOST16(sl_amx)		1316	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:22	13,452	-1	-1	4,123				
myHawkDot	SLHOST16(sl_amx)		1548	<ACCESS DENIE	spoolsv.exe	0.00	01-May-2014 23:18:23	3,336	-1	-1	434,0				
myHawkDot	SLHOST16(sl_amx)		1576	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:23	4,268	-1	-1	3,881				
myHawkDot	SLHOST16(sl_amx)		1796	NT AUTHORITY\	HeciServer.exe	0.00	01-May-2014 23:18:24	776	-1	-1	12,6				
myHawkDot	SLHOST16(sl_amx)		1820	NT AUTHORITY\	IProsetMonitor.exe	0.00	01-May-2014 23:18:24	756	-1	-1	10,3				
myHawkDot	SLHOST16(sl_amx)		2700	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:19:05	780	-1	-1	8,8				
myHawkDot	SLHOST16(sl_amx)		684	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:21:06	2,468	-1	-1	2,909				
myHawkDot	SLHOST16(sl_amx)		2944	NT AUTHORITY\	IAStorDataMgrSvc.exe	0.00	01-May-2014 23:21:08	5,836	-1	-1	1,102				
myHawkDot	SLHOST16(sl_amx)		2680	NT AUTHORITY\	jhl_service.exe	0.00	01-May-2014 23:21:19	980	-1	-1	16,6				
myHawkDot	SLHOST16(sl_amx)		4216	NT AUTHORITY\	MSIEXEC.exe	0.00	01-May-2014 23:21:24	1,724	-1	-1	159,1				

**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
- Host:** Choose a host to show data for in the display.

Fields and Data:

- Process Count:** The total number of processes in the table.

Table:

Each row in the table is a different host.

- Domain** The domain in which the host resides.
- Host Name** The name of the host.
- Expired** When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
- PID** The process ID.
- User** The user name.
- Process Name** The name of the process.
- CPU%** The amount of CPU used, in percent.
- Start Time** The host start time, in the following format:
0d 00:00 <days>d <hours>:<minutes>:<seconds>
 For example: **10d 08:41:38**
- Memory Used** The amount of memory currently used, in megabytes.
- Memory Resident** The amount of memory currently used by the process that resides in physical memory and is not paged out. Set to **-1** when the data is not available from an agent. (Hawk does not provide this data.)
- Memory Shared** The amount of physical memory that is shared with other processes. Set to **-1** when the data is not available from an agent. (Hawk does not provide this data.)
- Page Faults** The number of page faults.
- Page Faults /sec** The number of page faults per second.
- Timestamp** The date and time the data was last updated.

All Network Table

View network interface data in a tabular format. Each row in the table is a different network interface card (NIC). Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order.

Host Network Interfaces						
Domain	Host Name	Expired	if Name	Inet Addr	Mask	Flag
QATB	TESTBED-26	<input type="checkbox"/>	lo	127.0.0.1	255.0.0.0	UP LOOPBACK RUNN
QATB	TESTBED-26	<input type="checkbox"/>	enp0s3	192.168.200.76	255.255.255.0	UP BROADCAST RUN
QATB	TESTBED-34	<input type="checkbox"/>	lo	127.0.0.1	255.0.0.0	UP LOOPBACK RUNN
QATB	TESTBED-34	<input type="checkbox"/>	ens32	192.168.200.34	255.255.255.0	UP BROADCAST RUN

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.
- 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.

Filter By:

The display might include these filtering options:

- Domain:** Choose a domain for which to show NIC data. Domain names are specified when your administrator configures your Data Server.
- Host:** Choose a host for which to show NIC data.

Fields and Data:

- Interface Count:** The total number of NICs in the table.

Table:

Each row in the table is a different NIC.

Domain	The domain in which the NIC resides.
Host Name	The name of the NIC in which the network interface resides.
Expired	When checked, data has not been received from this NIC in the specified amount of time. The NIC will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
if Name	The name of the NIC.
Inet Addr	The NIC IP address.
Mask	The NIC subnet mask IP address.
Flags	Descriptive text for NIC flag.
MTU	The largest size packet or frame for the NIC.
Metric	Indicates...
Point To Point	Indicates whether the NIC is a point to point configuration.
Broadcast	Indicates whether the NIC is a broadcast configuration.
rxKBytes	The total number of kilobytes received by the NIC.
rxPackets	The total number of packets received by the NIC.
rxDropped	The total number of received packets that were dropped by the NIC.
rxErrors	The total number of received errors on the NIC.
rxOverruns	The total number of received overruns on the NIC.
rxFrame	The total number of received frames on the NIC.
txKBytes	The total number of kilobytes transmitted by the NIC.
txPackets	The total number of packets transmitted by the NIC.
txDropped	The total number of transmitted packets that were dropped by the NIC.
txErrors	The total number of transmission errors for the NIC.
txOverruns	The total number of transmission overruns for the NIC.
txCollisions	The total number of transmission collisions for the NIC.
txCarrier	The total number of transmission carrier errors for the NIC.
MAC Address	The NIC MAC address.
Rx KB/s	The number of kilobytes received per second.
Tx KB/s	The number of kilobytes transmitted per second.
Rx Packets/s	The number of packets received per second.

Tx Packets/s The number of packets transmitted per second.

Timestamp The date and time the data was last updated.

All Storage Table

View storage data in a tabular format. Each row in the table is a different storage partition. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order.

Domain	Host Name	Expired	File	%	Total	Used	Available	Mount Point	Type
QATB	WIN-8-CLONE	<input type="checkbox"/>	C:\	86.0	59.90	51.09	8.81	C:\	NTFS/c
QATB	WIN-8-CLONE	<input type="checkbox"/>	\\192.168.200.7\...	84.0	452.43	377.54	74.89	Z:\	NTFS/r

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.
- 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.

Filter By:

The display might include these filtering options:

Domain: Choose a domain or **All Domains** to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.

Host: Choose a host or **All Hosts** to show data for in the display.

Fields and Data:

Storage Count: The total number of storage partitions in the table.

Table:

Each row in the table is a different host.

Domain	The domain in which the host resides.
Host Name	The name of the host in which the storage partition resides.
Expired	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
File System	The storage partition location.
% Used	The amount of storage partition used, in percent.
Total Size (GB)	The storage partition size, in gigabytes.
Used (GB)	The amount of storage partition used, in gigabytes.
Available (GB)	The amount of storage partition available, in gigabytes.
Mount Point	The storage partition parent directory.
Type	The file system type.
Timestamp	The date and time the data was last updated.

Single Host


These displays present performance data for a single monitored host. Examine details about the health of your hosts.

To see your data in these displays you must install and configure the Solution Package for RTView Host Agent. Displays in this View are:

- ["Host Summary"](#)

Host Summary

This display provides a detailed view of utilization metrics for a single server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Filter By:

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
- Host:** Choose a host to show data for in the display.
- Expired** When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is **60** seconds.
- Last Update** The time the display was last updated.

Fields and Data:

Data describes the selected host except where noted.

- OS:** The operating system.
- Version:** The operating system version.
- Uptime:** The number of days, hours and minutes since started.
- #CPUs** The number of node connections.


CPU Type:	The type of CPU.	
%CPU	User	The amount of CPU used by the user, in percent.
	System	The amount of CPU used by the system, in percent.
	Idle	The amount of CPU that is not used, in percent.
Physical Memory	Used	The amount of physical memory used, in kilobytes.
	Total(MB)	The amount of physical memory available, in kilobytes.
	%Used	The amount of physical memory used, in percent.
Virtual Memory	Used	The amount of virtual memory used, in kilobytes.
	Total(MB)	The amount of virtual memory available, in kilobytes.
	%Used	The amount of virtual memory used, in percent.
Processes	The number of processes running.	
Load Avg:	1 Min	The average number of processes running over 1 minute.
	5 Min	The average number of processes running over 5 minutes.
	15 Min	The average number of processes running over 15 minutes.
Storage	File System	The amount of storage space used for the file system, in kilobytes.
	Mount Point	The name used by the operating system to mount and provide an entry point to other storage volumes.
	%Used	The amount of storage space used, in percent.
Network	ifName	The name assigned to the network interface by the operating system.
	RxKB/s	The amount of network data received per second, in kilobytes.
	TxKB/s	The amount of network data transmitted per second, in kilobytes.

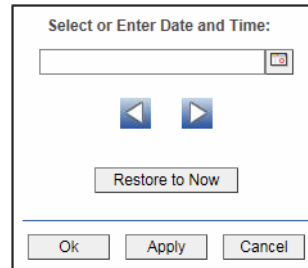
Trend Graphs


Traces metrics for the selected host.



- **CPU% Used:** The amount of CPU used, in percent.
- **Mem Total:** The amount of available memory, in kilobytes.
- **Mem Used:** The amount of memory used, in kilobytes.
- **Net Rx KB/s:** The amount of network data received per second, in kilobytes.
- **Net Tx KB/s:** The amount of network data transmitted per second, in kilobytes.

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

- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



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

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

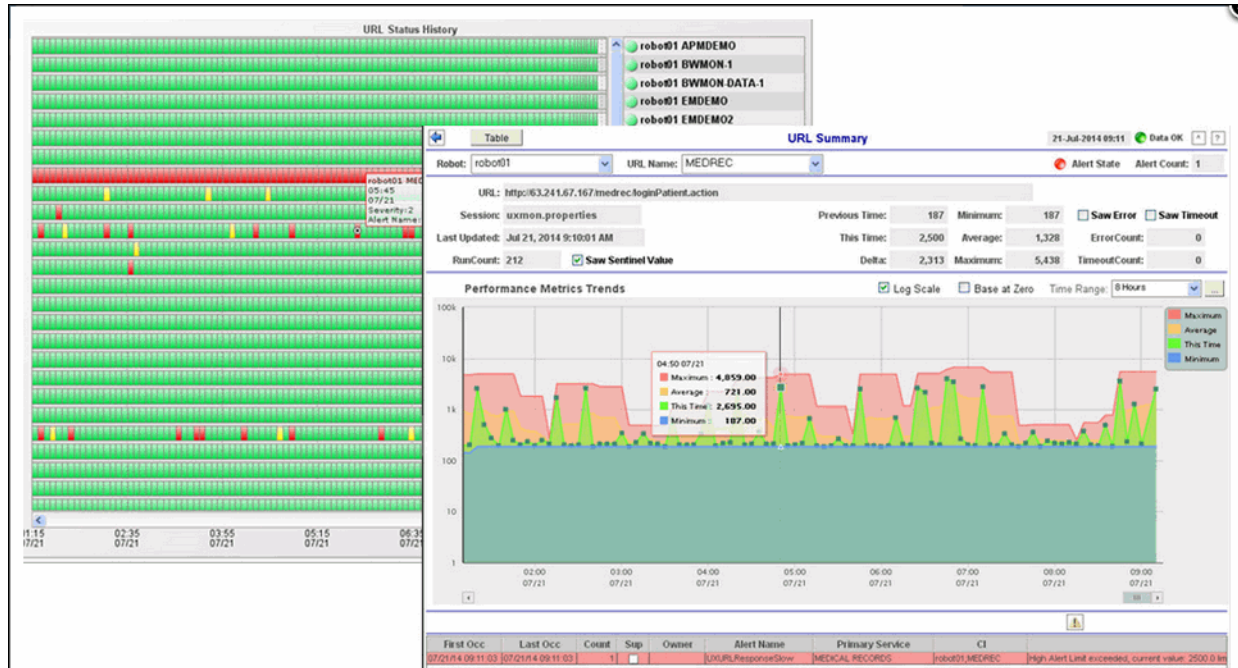
Click **Restore to Now** to reset the time range end point to the current time.

UX

The Solution Package for UX provides end-user monitoring for web applications. It does so by performing simulated transactions as if a real user is accessing a URL. When these simulated transactions are performed, information about the transaction such as the performance of the transaction, whether any errors were produced by the transaction or if the transaction produces invalid data are captured.

By including the Solution Package for UX in the RTView Enterprise Monitor® platform, you can be notified if any web application might be operating in one of these states:

- Unresponsive
- Performing slowly
- Generating errors
- Returning invalid information



The Solution Package for UX provides these features using UX Robots which are a set of robot Java applications. UX Robots read configuration files that designate the URLs to monitor, optional login details and an optional search string to validate. The UX Robots periodically query the URLs at configured interval periods and report the elapsed response time and any errors to the RTView Enterprise Monitor platform. If the search string is not found an error is reported. Preconfigured alerts are available with settings for thresholds to indicate if the response time was slow, the search string was not found, a timeout occurred, a URL error occurred or an error with the UX Robot.

When an error condition is reported, Monitor uses built-in displays that show the historical performance of all configured URLs and the status of the deployed UX Robots. This information can then be used to analyze when the problem occurred, whether performance has been degrading over a period of time, and allows users to correlate other errors reported from the RTView Enterprise Monitor which can indicate why the web application was having problems during that time period.

Displays for UX are:




- "All URLs Table"
- "All URLs Monitor"
- "All URLs Monitor"
- "URL Summary"
- "All Robots Table"
- "All URLs Table"
- "All URLs Table"
- "Robot Summary"

All URLs Table

View the most up-to-date performance data for all URLs under a single Robot or all Robots in a tabular format. Each row in the table is a different URL. Use this display to quickly identify alerts for any URL in your system, get an overview of how the URLs are performing and compare URL performance between UX Robot runs.


Row Color Code:

Tables with colored rows indicate the following:

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

Consider keeping this display open to monitor your URLs in general. For example, you can sort the **Alert** column so that all URLs with at least one Alarm Level (red) alert are in the top rows. Also use this to compare UX Robot performance between runs.

For a historical view of all URLs over time, refer to the **All URLs History** display. For a historical view of a single URL over time, refer to the **URL Summary** display.

Choose a UX Robot from the **Robot** drop-down menu. Enter a search string in the **URL Filter** field to filter data shown in the table. Use the sort  button to order column data. Drill-down and investigate by clicking a row to view details in the **URL Summary** display.

← All URLs - Table View 07-Jul-2016 16:14 Data OK + ?

Robot: All Robots Count: 46

URL Name Filter: Clear ☐ RegEx URL Filter: Clear ☐ RegEx

All URLs Table

Robot	URL Name	Sort	Alert	Run	Previous in ms	This time in ms	Average in ms	Min in ms	Max in ms	Saw Sentinel	Saw Error	Saw Timeout	
robot01	EMSMON-1	1		5	261	343	302	261	364	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	EMSMON-DATA-1	2		5	6	4	4	3	7	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	BWMON-1	3		5	343	325	387	325	550	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	BWMON-DATA-1	4		5	4	4	6	4	12	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	GFMON-1	6		5	1,033	1,085	1,045	1,026	1,085	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	EMDEMO	7		5	15	18	20	15	34	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	EMDEMO2	8		5	15	17	18	15	25	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	GFMON-DATA-1	9		5	11	11	21	11	61	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	MEDREC	10		5	427	130	194	128	427	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	MYOCM-OLD-57	11		5	4	4	5	3	11	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	MISCMON-DATA-DC	12		5	11	11	27	11	93	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	MISCMON_DC	13		5	1,024	1,018	1,030	1,018	1,070	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	OCMON-1	15		5	295	344	349	295	414	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	OCMON-2	16		5	330	352	373	330	405	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	OCMON-DATA-1	17		5	4	4	8	4	26	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	OCMON-DATA-2	18		5	3	3	4	3	7	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	OCMON-DATA-7	19		5	14	18	21	14	40	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	OCMON-7	20		5	21	23	37	20	99	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	SLMON	21		5	95	18	39	17	95	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	SLMON2	22		5	488	419	500	415	766	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	SLMON-ALERT-DATA	23		5	91	10	26	10	91	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	SLMON-CONFIG-DATA	24		5	33	11	21	11	37	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	SLMON4WIKI	25		5	1,098	1,047	1,042	1,017	1,098	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	RTVMGR-DC	26		5	1,147	1,046	1,053	1,018	1,147	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://63.2
robot01	RTVMGR-1	27		5	385	342	380	342	409	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2
robot01	RTVMGR-DATA-1	28		5	4	3	3	3	4	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	http://12.2

Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

Robot Choose a Robot to see metrics for.

URL Name Filter: Enter a (case-sensitive) string to search for.

Clear Clears the **Filter** text entry and filtered search results in the table.

Regex Check to toggle the **Filter** field to accept Regular Expressions for filtering.

URL Filter: Enter a (case-sensitive) string to search for.

- Clear** Clears the **Filter** text entry and filtered search results in the table.
- Regex** Check to toggle the **Filter** field to accept Regular Expressions for filtering.




Fields and Data

This display includes:

- Count** The number of rows currently in the table.
- URL Name Filter** Enter a string to search in the **URL Name** table column, then click **Enter**. Only rows with **URL Name** columns containing the matching search string are shown in the table.
- Clear** Removes entries in the **URL Name Filter** field and filter results in the table.
- Regex** Check to toggle the **Filter** field to accept Regular Expressions for filtering.
- URL Filter** Enter a string to search in the **URL** table column, then click **Enter**. Only rows with **URL** columns containing the matching search string are shown in the table.
- Clear** Removes entries in the **URL Filter** field and filter results in the table.
- Regex** Check to toggle the **Filter** field to accept Regular Expressions for filtering.

All URLs Table

Each row in the table is a different URL. Data in the row columns describe the run for the URL.


- Robot** The name of the UX Robot that is sending these statistics. (For details, see **agentname** in the **uxmon.properties** file.)
- URL Name** The nickname of the URL where the alert data originated.
- Sort** This is the **sortIndex** column in the URL Configuration Line and can be used to by the administrator to define the sort order for URLs in the URL Configuration Line (of the configuration file). This is useful when you do NOT want the sort order defined by the alphabetical sort or the Alert Severity sort. The table sorted in this way is closest to the order in the **uxmon.properties** file."
- Alert** The severity level of any open alert. Values range from **0** to **2**, where **2** is the greatest Severity:
-  One or more alerts exceeded their ALARM LEVEL threshold for the URL.
 -  One or more alerts exceeded their WARNING LEVEL threshold for the URL.
 -  No alert thresholds have been exceeded for the URL.
- Run** The count number of the Robot run in the sequence. The Robot runs in a loop that is controlled by the **repeatType** property. If the Robot is set to repeat then each time it begins with the first URL and the run count is incremented.
- Previous in ms** The amount of time, in milliseconds, for the last completed URL connection. A Robot process can include connecting to one or more URLs, logging on to a web page and performing a search using a specified search string.
- This time in ms** The amount of time, in milliseconds, for the most recently completed URL connection. A Robot process can include connecting to one or more URLs, logging on to a web page and performing a search using a specified search string.

Average in ms	The average amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
Min in ms	The least amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
Max in ms	The most amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
Saw Sentinel	When checked (true), the Robot found the specified search string (the sentinel, which is specified in the configuration file) after contact with the specified URL (and logging in to the page if log in is also specified in the configuration file). Searching for a Search String or Sentinel is optional. You configure this option in the uxmon.properties file by setting searchType=simple or searchType=regular . No search occurs if searchType=none .
Saw Error	When checked (true), the Robot encountered one or more errors while engaged with the URL. Errors can be of many different types ranging from lack of response to the attempt to contact the URL to server error.
Saw Timeout	When checked (true), the Robot encountered a connection timeout while engaged with the URL. Note that your administrator can adjust the amount of time for the timeout.
URL	The fully qualified address for the URL the Robot is testing.
Response End Time	The exact time that the URL finished responding.
Expired	When checked (true), the Robot has not received a response from the URL for the amount of time specified. (The \$uxRowExpirationTime property specifies the time and is set in the rtvapm.uxmon.properties file.)
Time Stamp	The time the last data was delivered.


All URLs Monitor

View historical performance data over time for all URLs in one or all Robots in a status history object. This display also shows the current and historical alert status of the URLs. Each row in the status history object is a different URL. Each column represents a time period. A darker color indicates heavier usage, a lighter color indicates lighter usage.

Use this display to monitor URL performance and determine whether URLs encounter alerts during certain periods of time. Observe utilization trends for your entire system. Analyze load distribution, check for bottlenecks and identify URLs with high usage. You can also answer questions such as, Is the web page using what I expect? Is the system using it across URLs in a uniform scale? If there is an issue, mouse-over the heatmap to see when the issue started, what behavior preceded it, and the name of the resource.

Choose one or **All Robots** from the **Robot** drop-down menu to filter display data. Change the **Time Range** to "zoom in" on the graph and see more detail or "zoom out" from the graph to see larger trends over time. To change the time range click Open Calendar , choose the date and time, then click **OK**.

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



Title Bar (possible features are):

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

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

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

Alert Views - RTView Alerts Table display.

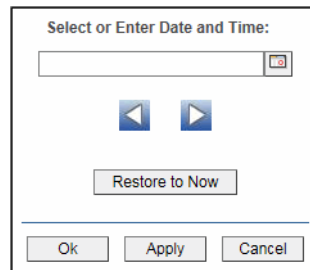
Color Code:

Row color indicates the following:



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

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 for the graph, click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.


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

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

URL Summary

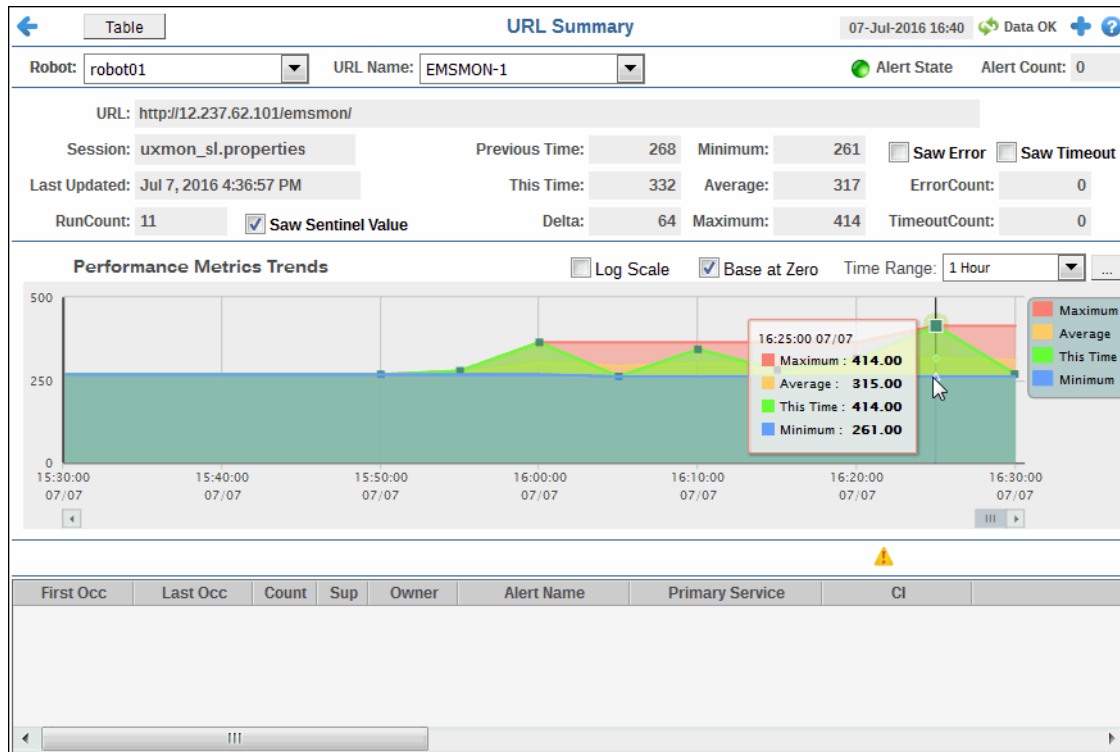
View historical and current performance and alert data, over time, for a single URL in a trend graph. This display shows all the current and historical alerts for the URL. Each trace in the trend graph is a different measurement of the UX Robot results for the URL - the **Maximum**, **Average, This Time** and **Minimum**. This display shows the data for the selected URL that is shown in the **All URLs** display table.

Use this display to monitor the performance of a URL and see details about the alerts it encounters.

Choose a **UX Robot** and a **URL Name** from the drop-down menus to filter display data. Move the bar at the base of the graph to time to see values for specific times. Change the **Time Range** to "zoom in" on the graph and see more detail or "zoom out" from the graph to see larger trends over time. To change the time range click Open Calendar , choose the date and time, then click **OK**.

For a historical view of all URLs over time, refer to the **All URLs History** display.

Use the sort button to order column data.







Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
- 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.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

Fields and Data

This display includes:

- Alert State** Indicates the greatest severity level of all open alerts for this URL:

 -  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold.
 -  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold.
 -  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold.
 -  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** buttons are disabled.

Alert Count The number of open alerts for the selected URL.


URL The URL address.

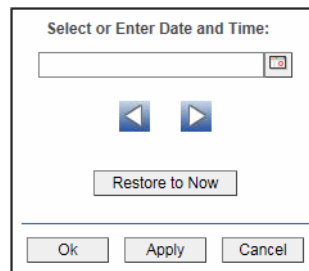
Session	The filename from which the Robot started.
Last Updated	The date and time that the Robot last sent data.
RunCount	The number of runs made by the UX Robot since it started.
Saw Sentinel Value	When checked (true), the Robot found the specified search string (the sentinel, which is specified in the configuration file) after contact with the specified URL (and logging in to the page if log in is also specified in the configuration file). It might also optionally look for a search string called a Sentinel.
Previous Time	The amount of time, in milliseconds, for the last completed Robot run to complete the Robot process specified. A Robot process can include connecting to one or more URLs, logging on to a web page and performing a search using a specified search string.
This Time	The amount of time, in milliseconds, for the most recently completed Robot run to complete the Robot process specified. A Robot process can include connecting to one or more URLs, logging on to a web page and performing a search using a specified search string.
Delta	The time difference, in milliseconds, between the latest and previous Robt runs.
Minimum	The least amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
Average	The average amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
Maximum	The most amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
Saw Error	When checked (true), the Robot encountered one or more errors while engaged with the URL. Errors can be of many different types ranging from lack of response to the attempt to contact the URL to server error.
Saw Timeout	When checked (true), the Robot encountered a connection timeout while engaged with the URL. Note that your administrator can adjust the amount of time for the timeout.
Error Count	Indicates whether the Robot encountered an error: 0 = No error encountered. 1 = Error encountered.
Timeout Count	Indicates whether the Robot encountered a timeout error: 0 = No timeout error encountered. 1 = Error encountered.

Performance Metrics Trends


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



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath these arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

Performance Metrics Trends

Maximum	Traces the most amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
Average	Traces the average amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.
This Time	Traces the amount of time, in milliseconds, for the most recently completed Robot run to complete the Robot process specified. A Robot process can include connecting to one or more URLs, logging on to a web page and performing a search using a specified search string.
Minimum	Traces the least amount of time in this session of Robot runs, in milliseconds, for the URL to complete the process specified by the UX Robot.

Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the display. Each row in the table is a different active alert. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized.

 Opens the **Alerts Table** display in a new window.


First Occ	The date and time the alert first occurred.
Last Occ	The date and time the alert last occurred.
Count	The number of times the alert was generated.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Primary Service	The name of the Service with which the alert is associated.
CI	The CI alert source.
Alert Text	Description of the alert.
AlertClass	An optional alert field which can be used when integrating with other alerting systems.
CompID	An optional alert field which can be used when integrating with other alerting systems.
TicketID	An optional alert field which can be used when integrating with other alerting systems.
TicketGroup	An optional alert field which can be used when integrating with other alerting systems.

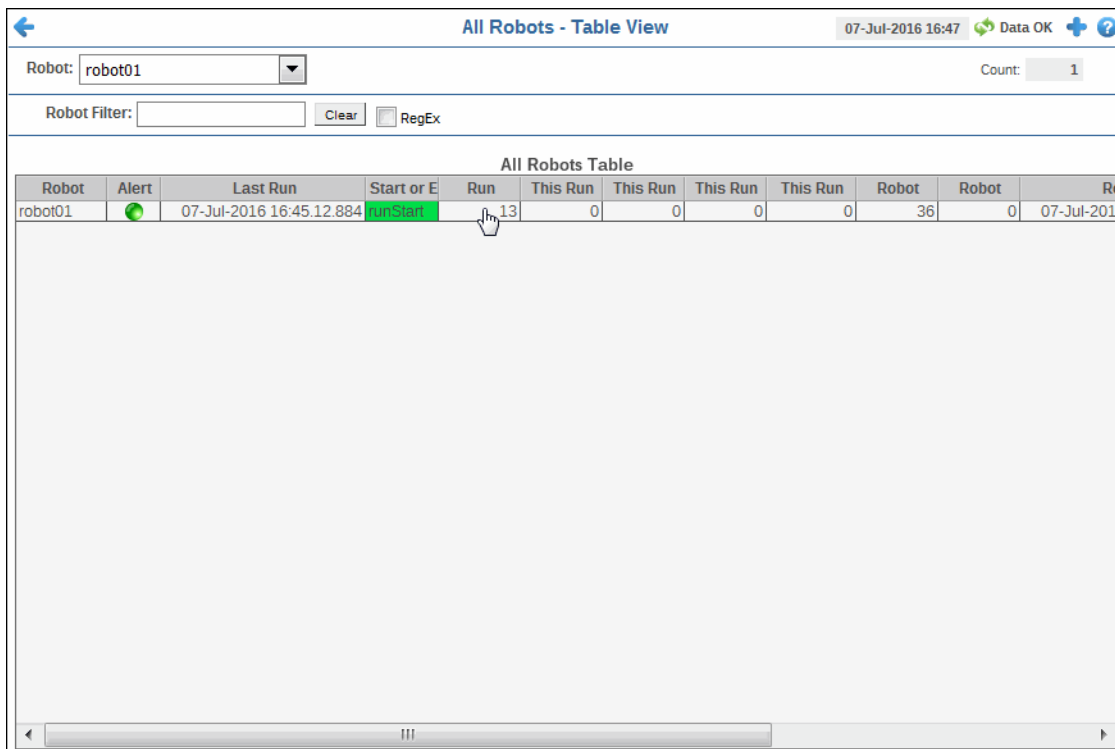
All Robots Table


View the most up-to-date performance data for one or all Robots in a tabular format. Each row in the table is a different Robot. Use this display to quickly identify alerts for any Robot in your system, get an overview of how the Robots are performing and compare Robot performance between UX Robot runs.

Consider keeping this display open to monitor your Robots in general. For example, you can sort the **Alert** column so that all URLs with at least one Alarm Level (red) alert are in the top rows. Also use this to compare Robot performance between runs.








For a historical view of Robots over time, refer to the **All Robots History** display. For a historical view of a single URL over time, refer to the **Robot Summary** display.




Choose a UX Robot from the **Robot** drop-down menu. Enter a search string in the **Robot Filter** field to filter data shown in the table. Use the sort  button to order column data. Drill-down and investigate by clicking a row to view details in the **All Robots Monitor** display.



Robot	Alert	Last Run	Start or End	Run	This Run	This Run	This Run	This Run	Robot	Robot	Re
robot01		07-Jul-2016 16:45:12.884	runStart	13	0	0	0	0	36	0	07-Jul-201




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.
-   open commonly accessed displays.
-  6,047 The number of items currently in the display.

-  Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
-  23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
-  Open the Alert Views - RTView Alerts Table display.

Row Color Code:

Tables with colored rows indicate the following:

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

Filter By:

The display might include these filtering options:

Robot Choose a robot to see metrics for.





Robot Filter:	Enter a (case-sensitive) string to search for.
Clear	Clears the Filter text entry and filtered search results in the table.
Regex	Check to toggle the Filter field to accept Regular Expressions for filtering.

Fields and Data

This display includes:

Count	The number of rows currently in the table.
Regex	Check to toggle the Filter field to accept Regular Expressions for filtering.

All Robots Table


Robot	The name of the UX Robot that is sending these statistics. (For details, see agentname in the uxmon.properties file.)
Alert	Indicates the greatest severity level of all open alerts for this Robot:  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold.  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold.  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. Start or End of Run
Last Run	The time, in Java format, the last run started.
Start or End of Run	Indicates whether a Robot is in progress or paused. This column name is runStart if a run is in progress and runEnd if the last run ended and the UX Robot is pausing before beginning the next run.
Run	The number of the Robot run that was performed last.
This Run Total MS	If the Robot is in progress (runStart), the amount of time, in milliseconds, for all URLs to complete their specified process during this run. If the Robot is paused (runEnd) the value is zero (0). While in the middle of the run, no count is possible.
This Run	The total number of errors seen for all URLs during this run or zero (0) if runStart . A count occurs only after a run completes.
This Run Timeouts	The total number of timeouts the Robot encountered for all URLs during this run. If the Robot is in progress (runStart) the value is zero (0). While in the middle of the run, no count is possible.
This Run Missed Search Sentinels	The total number of Missed Search Sentinels the Robot encountered for all URLs during this run. If the Robot is in progress (runStart) the value is zero (0). While in the middle of the run, no count is possible.
Robot	The total number of errors seen for all URLs for all runs so far this session.
Robot Overall Errors	The total number of timeouts the Robot encountered for all URLs during this session. If the Robot is in progress (runStart) the value is zero (0).

Robot Overall Start Time	The time the Robot session started.
Expired	When checked (true), the Robot encountered a connection timeout while engaged with the URL. Note that your administrator can adjust the amount of time for the timeout (see the \$uxRowExpirationTime property in the <code>rtvapm.uxmon.properties</code> file).
Time Stamp	The time the last data was delivered for either runStart or runEnd .


All Robots Monitor


View historical performance data over time for one or all Robots in a status history object. This display also shows the current and historical alert status of the Robots. Each row in the status history object is a different Robot. Each column represents a time period. A darker color indicates heavier usage, a lighter color indicates lighter usage.

Use this display to monitor Robot performance identify whether Robots encounter alerts during certain periods of time. Observe utilization trends for your entire system. Analyze load distribution, check for bottlenecks and identify URLs with high usage. You can also answer questions such as, Is the web page using what I expect? Is the system using it across URLs in a uniform scale? If there is an issue, mouse-over the heatmap to see when the issue started, what behavior preceded it, and the name of the resource.

Choose one or **All Robots** from the **Robot** drop-down menu to filter display data. Change the **Time Range** to “zoom in” on the graph and see more detail or “zoom out” from the graph to see larger trends over time. To change the time range click Open Calendar , choose the date and time, then click **OK**. Drill-down and investigate by clicking a row to view details in the **Robot Summary** display.

Enter a (case-sensitive) string in the **Robot Filter** to perform search. Click **Clear** to clear the **Robot Filter** string and filtered search results in the table.





Color Code:

Row color indicates the following:

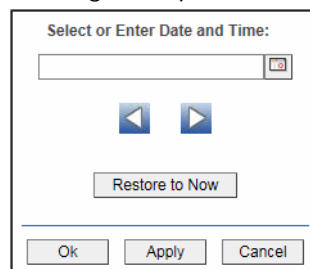
- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the row.

Fields and Data

This display includes:

Base At Zero Select to 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. By default, the time range end point is the current time.








To change the time range for the graph, click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the display. Each row in the table is a different active alert. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized.


 Opens the **Alerts Table** display in a new window.

First Occ	The date and time the alert first occurred.
Last Occ	The date and time the alert last occurred.
Count	The number of times the alert was generated.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Primary Service	The name of the Service with which the alert is associated.
CI	The CI alert source.
Alert Text	Description of the alert.
AlertClass	An optional alert field which can be used when integrating with other alerting systems.
CompID	An optional alert field which can be used when integrating with other alerting systems.
TicketID	An optional alert field which can be used when integrating with other alerting systems.
TicketGroup	An optional alert field which can be used when integrating with other alerting systems.

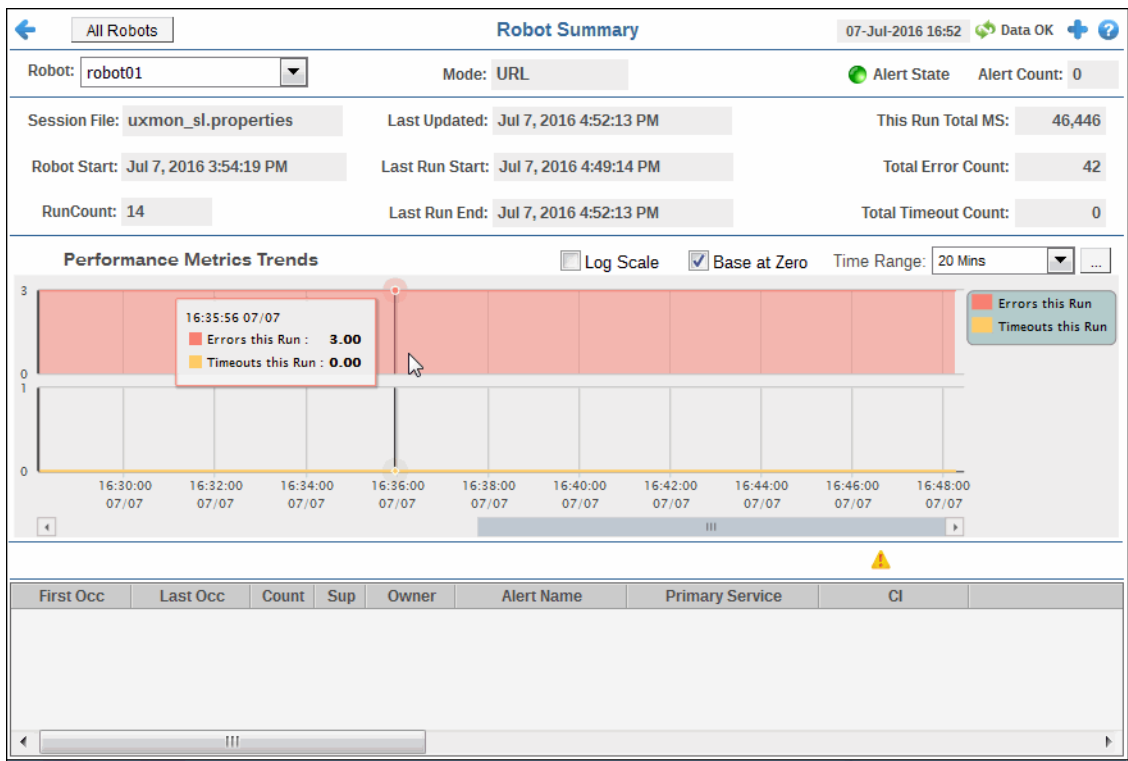
Robot Summary

View historical and current performance and alert data, over time, for a single Robot in a trend graph. This display shows all the current and historical alerts for the Robot. Each trace in the trend graph is a different measurement of the UX Robot results for the Robot - the **Errors this Run** and **Timeouts this Run**. This display shows the data for the selected Robot that is shown in the **All Robots Table** display.

Use this display to monitor the performance of a Robot and see details about the alerts it encounters.

Choose a UX Robot from the **Robot** drop-down menus to filter display data. Move the bar at the base of the graph to time to see values for specific times. Change the **Time Range** to "zoom in" on the graph and see more detail or "zoom out" from the graph to see larger trends over time. To change the time range click Open Calendar , choose the date and time, then click **OK**.

For a historical view of all Robots over time, refer to the **All Robots History** display.



Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table
open commonly accessed displays.

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

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

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

Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:


- Mode:** The mode that the UX Robot operates in. The default is **URL**. (Do not modify this setting.)
- Alert State** Indicates the greatest severity level of all open alerts for this Robot:
 - Red indicates that one or more alerts exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold.
 - Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold.
 - Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** buttons are disabled.
- Alert Count** The number of open alerts for the selected Robot.

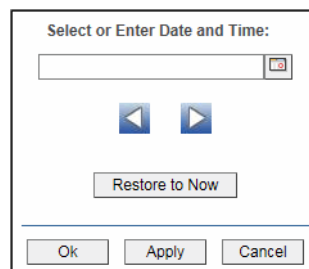
Session File	The name of the properties file that this UX Robot is running from.
Robot Start	The time the this Robot session started.
RunCount	The number of runs made in this Robot session.
Last Updated	The time when the Robot data was last updated.
Last Run Start	The time when the last run started for this UX Robot.
Last Run End	The time when the last run ended for this UX Robot.
This Run Total MS	Total MS in response times for all of the URLs in this run.
Total Error Count	The total number of errors for all of the URLs in this run.
Total Timeout Count	The total number of timeouts for all of the URLs in this run.


Performance Metrics Trends

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

Time Range

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




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





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

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

Errors this Run	The total number of errors this Robot encountered for ALL URLs the Robot is connecting to on the current run.
Timeouts this Run	The total number of timeouts this Robot encountered for ALL URLs the Robot is connecting to on the current run.

Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details**, **Own**, **Suppress**, **Close**, **Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** buttons are disabled.

 Opens the **Alerts Table** display in a new window.

Own	Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
Suppress	Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
Close	Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
Details	Select an alert, right-click and choose Alert/Details to open the Alert Detail window and view alert details. Or, double-click an alert to open the Alert Detail window.
Annotate	Select one or more alerts, right-click and choose Alert/Annotate to open the Set Owner and Comments dialog and enter comments or change alert owner.
Options	Select an alert, right-click and choose Alert/Options to open the Alert Options dialog. This dialog is provided for customizing your own alert options.

First Occ	The date and time the alert first occurred.
Last Occ	The date and time the alert last occurred.
Count	The number of times the alert was generated.
Sup	When checked, the alert has been suppressed by a user.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert.
Primary Service	The name of the Service with which the alert is associated.
CI	The CI alert source.
Alert Text	Description of the alert.
AlertClass	An optional alert field which can be used when integrating with other alerting systems.
CompID	An optional alert field which can be used when integrating with other alerting systems.
TicketID	An optional alert field which can be used when integrating with other alerting systems.
TicketGroup	An optional alert field which can be used when integrating with other alerting systems.

VMware vCenter

VMware vCenter displays enable you to monitor the health and performance of your virtual machines at the cluster level or the machine level, including viewing datastores, network data events and alerts.

The following Views (and associated displays) can be found under **Components tab > Hosts/VMWare**:

- **"Clusters View"**: View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.
- **"Virtual Machines View"**: View current and historical data for your virtual machines.
- **"Datastores View"**: The displays in this View provide a list of datastores on one or all servers, a list of all hosts mounted to a particular datastore, a list of all virtual machines hosted by a particular datastore, or data for a particular datastore.
- **"Networks View"**: View a list of all networks, as well as data associated with the networks, that exist on one server or on all servers.
- **"Events/Alarms View"**: View event data and alarm data for one server or for all servers.

This section includes the following Views:

- **"Clusters View"**: View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.
- **"Hosts View"**: View performance and utilization data for hosts running on one or all clusters, view utilization data for a specific host running virtual machines, view a list of components contained on a selected host, and view physical and virtual network adapters located on a particular host.
- **"Virtual Machines View"**: View current and historical data for your virtual machines.
- **"Datastores View"**: The displays in this View provide a list of datastores on one or all servers, a list of all hosts mounted to a particular datastore, a list of all virtual machines hosted by a particular datastore, or data for a particular datastore.
- **"Networks View"**: View a list of all networks, as well as data associated with the networks, that exist on one server or on all servers.
- **"Events/Alarms View"**: View event data and alarm data for one server or for all servers.

Clusters View

The display available in the View lists all clusters that are configured on a particular server or on all servers.

The display available in this view is:

- **"All Clusters"**: View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.

All Clusters

View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.

← VMware Cluster Compute Resources - Table 22-Mar-2017 10:14 Data OK + ?

Server: All Servers ▾

Cluster Count: 1 Clusters

Server	clustername	Alert Severity	Alert Count	Overall Status	# Hosts	# Effective Hosts	HA Enabled
qavSphere1	SL CORP HA CLUSTER		0		2	2	<input type="checkbox"/>

< >

Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.








Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

Filter By:

The display might include these filtering options:

- Server** Select the server for which you want to view data.
- Cluster Count** The total number of clusters in the selected server(s), which are listed in the **Clusters** table.

Clusters Table

Server	The name of the server.
clustername	The name of the cluster.
Alert Severity	<p>The highest level alert on the cluster.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The number of alerts currently on the cluster.
Overall Status	<p>The general health status of the cluster.*</p> <ul style="list-style-type: none">  Red indicates that the host is experiencing a problem.  Yellow indicates that the host might have a problem.  Grey indicates that the status of the host's health is unknown.  Green indicates that host's status is OK.
# Hosts	Lists the number of hosts on the cluster.*
# Effective Hosts	Lists the number of effective hosts.*
HA Enabled	When checked, this check box signifies that High Availability is enabled on the cluster.*
HA Admission Enabled	When checked, this check box signifies that High Availability strict admission is enabled.*
HA Admission Policy	Lists the High Availability admission policy for the cluster.*
HA Datastore Candidate	Displays the High Availability datastore candidate defined on the cluster.*
HA Host Monitoring	Lists whether or not High Availability host monitoring is enabled on the cluster.*
HA VM Monitoring	Lists whether or not High Availability virtual machine monitoring is enabled on the cluster.*
DRS Enabled	When checked, this check box signifies that DRS (Distributed Resource Scheduler) is enabled.*
DRS Enable Behavior Overrides	When checked, this check box signifies that DRS behavior overrides for individual virtual machines are enabled.*
DRS Default VM Behavior	Lists the cluster-wide default DRS behavior for virtual machines.*
Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Timestamp	The date and time the data was last updated.

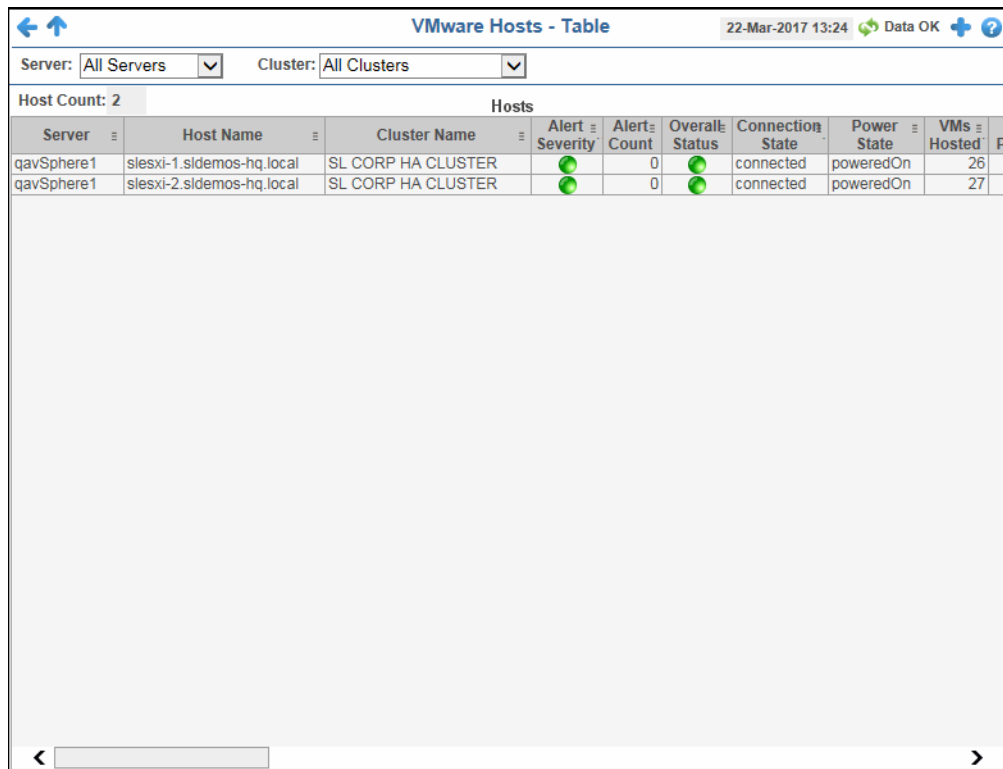
Hosts View

These displays present performance and utilization data for hosts running on one or all clusters, present utilization data for a specific host running virtual machines, list components contained on a selected host, and list physical and virtual network adapters for a particular host. Displays in this View are:


- **"All Hosts"**: A tabular view of the utilization data for all hosts running on one or on all clusters.
- **"Single Host Summary"**: Displays utilization data for a specific host running virtual machines.
- **"Host Health"**: View the components contained on a selected host and the component's associated data.
- **"Host NICs"**: View data for all physical and virtual network adapters (NICs) for a particular host.

All Hosts

View the utilization data for all hosts running on one cluster or on all clusters.



Server	Host Name	Cluster Name	Alert Severity	Alert Count	Overall Status	Connection State	Power State	VMs Hosted
qavSphere1	slesxi-1.sldemos-hq.local	SL CORP HA CLUSTER		0		connected	poweredOn	26
qavSphere1	slesxi-2.sldemos-hq.local	SL CORP HA CLUSTER		0		connected	poweredOn	27



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

Filter By:

The display might include these filtering options:

Server	Select the server for which you want to view data.
Cluster	Select the cluster for which you want to view data.
Host Count	The total number of hosts of the selected cluster(s), which are listed in the table

Hosts Table

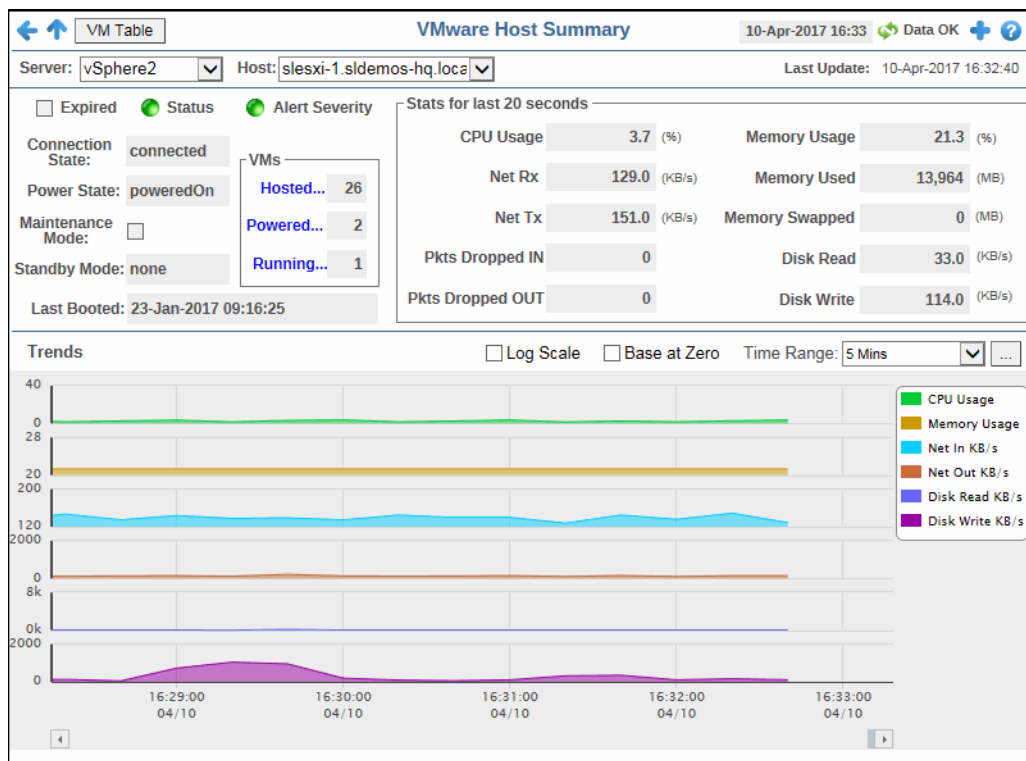
Server	The name of the server.
Host Name	The name of the host.
Cluster Name	The name of the cluster.
Alert Severity	The highest level alert on the host. <div> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. ● Green indicates that no metrics have exceeded their alert thresholds. </div>
Alert Count	The number of alerts currently on the host.
Overall Status	The general health status of the host.* <div> ● Red indicates that the host is experiencing a problem. ● Yellow indicates that the host might have a problem. ● Grey indicates that the status of the host's health is unknown. ● Green indicates that host's status is OK. </div>
Connection State	Lists the status of the connection.*
Power State	Lists whether the host is powered on or powered off.*
VMs Hosted	The number of virtual machines that exist on the host.*


VMs Powered	The number of virtual machines powered on on the host.*
VMs Running	The number of virtual machines running on the host.*
Maintenance Mode	When checked, this check box signifies that the host is in maintenance mode.*
Standby Mode	The host's standby mode.*
CPU % Usage	The percentage of CPU used by the virtual machines.*
Num CPU Cores	The total number of cores on the CPU.*
Num CPU Threads	The total number of threads on the CPU.*
Memory % Usage	The percentage of the host's memory currently in use.*
Memory Used (MB)	The total memory used, in megabytes, on the host.*
Memory Total (MB)	The total amount of memory, in megabytes.*
Swap Used (MB)	The total amount of swap space used by the host, in megabytes.*
Disk Reads (KB/sec)	The amount of data being read from the disk per second, in kilobytes.*
Disk Writes (KB/sec)	The amount of data being written to the disk per second, in kilobytes.*
Net IN (KB/sec)	The amount of network data being received per sec, in kilobytes.*
Net OUT (KB/sec)	The amount of network data being transmitted per sec, in kilobytes.*
% IN Packet Loss (Drops)	The percentage of incoming packets that were dropped.*
% OUT Packet Loss (Drops)	The percentage of outgoing packets that were dropped.*
% IN Packet Loss (Errors)	The percentage of incoming packets that had errors.*
% OUT Packet Loss (Errors)	The percentage of outgoing packets that had errors.*
Packets IN	The number of incoming packets.*
Packets OUT	The number of outgoing packets.*
Packets IN Dropped	The number of incoming packets that were dropped.*
Packets OUT Dropped	The number of outgoing packets that were dropped.*

Packets IN Errors	The number of incoming packets that had errors.*
Packets OUT Errors	The number of outgoing packets that had errors.*
System Vendor	The name of the system vendor.
System Model	The name of the system model.
Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Last Booted	The date and time in which the host was last restarted.*
Timestamp	The date and time the data was last updated.

Single Host Summary

View the number of virtual machines running on a particular host, the most recent utilization data for the host, and the trend data for the host over a specified time range.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields.

Filter By:

The display might include these filtering options:

Server	The name of the server containing the host
Host	The host of the virtual machines for which you want to view data.
Last Update	The date and time that the data in the table was last updated.

Fields and Data:

Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Status	<p>The general health status of the host.*</p> <ul style="list-style-type: none"> Red indicates that the host is experiencing a problem. Yellow indicates that the host might have a problem. Grey indicates that the status of the host's health is unknown. Green indicates that host's status is OK.
Alert Severity	<p>The alert severity for the selected host:</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Connection State	Displays the current state of the connection for the host (connected/notConnected).*
Power State	Lists whether the host is powered on or powered off.*
Maintenance Mode	<p>Signifies whether or not the host is in maintenance mode.*</p> <p>true: host is in maintenance mode.</p> <p>false: host is not in maintenance mode.</p>


Standby Mode		The host's standby mode.*
Last Booted		The date and time in which the host was last restarted.*
VMs	Hosted...	The number of virtual machines on the host.*
	Powered...	The number of virtual machines on the host that are powered on.*
	Running...	The number of virtual machines currently up and running on the host.*
Stats for last 20 seconds	CPU Usage	The percentage of CPU used in the last 20 seconds.*
	Net Rx	The amount of network data received, in kilobytes per second, in the last 20 seconds.*
	Net Tx	The amount of network data transmitted, in kilobytes per second, in the last 20 seconds.*
	Pkts Dropped IN	The number of incoming packets that were dropped in the last 20 seconds.*
	Pkts Dropped OUT	The number of outgoing packets that were dropped in the last 20 seconds.*
	Memory Usage	The percentage of memory used in the last 20 seconds.*
	Memory Used	The amount of memory used, in megabytes, in the last 20 seconds.
	Memory Swapped	The amount of memory swapped, in megabytes, in the last 20 seconds.*
	Disk Read	The amount of data read from the disk, in kilobytes per second, in the last 20 seconds.*
	Disk Write	The amount of data written to the disk, in kilobytes per second, in the last 20 seconds.*

Trend Graphs

Traces the sum of process metrics for the host:

- **CPU Usage:** The percentage of CPU used.
- **Memory Usage:** The amount of memory used.
- **Net In KB/s:** The amount of network data received per second, in kilobytes per second.
- **Net Out KB/s:** The amount of network data transmitted per second, in kilobytes per second.
- **Disk Read KB/s:** The amount of data being read from the disk, in kilobytes per second.
- **Disk Write KB/s:** The amount of data being written to the disk, in kilobytes per second.

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



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields.

Filter By:

The display might include these filtering options:

Server	The name of the server containing the host
Host	The host of the virtual machines for which you want to view data.
Component Count	The number of components found on the host, which are displayed in the table.

Host Health Table

Server	The name of the server.
Host Name	The name of the host.
Component	The name of the component.
State	The general health status of the host.* <div> ● Red indicates that the host is experiencing a problem. ● Yellow indicates that the host might have a problem. ● Grey indicates that the status of the host's health is unknown. ● Green indicates that host's status is OK. </div>
Sensor Type	Lists the component's sensor type.*
Current Reading	Lists the current reading of the element indicated by the sensor.*
Units	Indicates the base units in which the sensor reading is specified.*
Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Timestamp	The date and time the data was last updated.

Host NICs

View data for all physical and virtual network adapters (NICs) for a particular host.

← VMware Host NICs - Table 22-Mar-2017 15:44 Data OK + ?

Server: All Servers Host: slesxi-1.sldemos-hq.loc

Physical NIC Count: 4

Server	Host Name	Device	Link Duplex?	Link Speed (Mb)	AutoNegotiate?	Resource Pool	Scheduler?
qavSphere1	slesxi-1.sldemos-hq.local	vmnic0	<input checked="" type="checkbox"/>	1000	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
qavSphere1	slesxi-1.sldemos-hq.local	vmnic1	<input checked="" type="checkbox"/>	1000	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
qavSphere1	slesxi-1.sldemos-hq.local	vmnic2	<input checked="" type="checkbox"/>	1000	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
qavSphere1	slesxi-1.sldemos-hq.local	vmnic3	<input type="checkbox"/>	0	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>

Virtual NIC Count: 3

Server	Host Name	Device	DHCP?	IP Address	Subnet Mask	Portgroup
qavSphere1	slesxi-1.sldemos-hq.local	vmk0	<input type="checkbox"/>	192.168.200.51	255.255.255.0	Management Network
qavSphere1	slesxi-1.sldemos-hq.local	vmk1	<input type="checkbox"/>	192.168.201.2	255.255.255.0	VMotion-01
qavSphere1	slesxi-1.sldemos-hq.local	vmk2	<input type="checkbox"/>	192.168.201.3	255.255.255.0	VMotion-02

Title Bar (possible features are):

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

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

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

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields.

Filter By:

The display might include these filtering options:

Server The name of the server containing the host

Host	The host of the virtual machines for which you want to view data.
Physical NIC Count	The number of NICs found on the host, which are displayed in the table.

Physical NICs Table

Server	The name of the server.
Host Name	The name of the host.
Device	The name of the device.
Link Duplex?	When checked, indicates that the link is capable of full-duplex. When unchecked, indicates that the link is only capable of half-duplex.*
Link Speed (MB)	The bit rate on the link, in megabytes.*
AutoNegotiate?	When checked, indicates that the physical network adapter supports autonegotiate.*
Resource Pool Scheduler?	When checked, indicates that the physical network adapter allows resource pool-based scheduling for network I/O control.*
VM Direct Path Gen2?	When checked, indicates that the NIC supports VMDirectPath Gen 2.*
Wake On LAN?	When checked, indicates that the NIC is wake-on-LAN capable.*
Driver	The name of the driver.*
MAC	The media access control (MAC) address of the physical network adapter.*
DHCP?	When checked, indicates that the network adapter uses a DHCP server.*
IP Address	The IP address of the physical network adapter.*
Subnet Mask	The subnet mask for the physical network adapter.*
Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Timestamp	The date and time the data was last updated.

Virtual NICs Table

Virtual NICs Count	The number of virtual NICs found on the host, which are displayed in the table.
Server	The name of the server.
Host Name	The name of the host.
Device	The name of the device.
DHCP?	When checked, indicates that the network adapter uses a DHCP server.*
IP Address	The IP address of the virtual network adapter.*

Subnet Mask	The subnet mask for the virtual network adapter.*
Port Group	The name of the port group in which the virtual network adapter resides.*
MAC	The media access control (MAC) address of the virtual network adapter.*
TSO Enabled?	When checked, indicates that TCP segment offloading (TSO) is enabled.*
Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Timestamp	The date and time the data was last updated.

Virtual Machines View

These displays present current and historical data for your virtual machines. Displays in this View are:


- [“All VMs Heatmap”](#): A color-coded heatmap view of utilization metrics.
- [“All VMs Table”](#): View data shown in the [“All VMs Heatmap”](#) display, as well as additional details, in a tabular format. Use this display to view all available data for each virtual machine by server and host.
- [“All VMs Disk Table”](#): View disk usage percentage, available disk space, and total capacity for one or all virtual machines on a specific server/host combination.
- [“Single VM Summary”](#): View current and historical utilization and performance metrics for a single virtual machine.

All VMs Heatmap

View the most critical CPU and memory usage, disk read and write utilization, and incoming and outgoing data metrics for your virtual machines. Use this display to quickly identify virtual machines with critical alerts.

Each rectangle in the heatmap represents a virtual machine. The rectangle color indicates the most critical alert state associated with the virtual machine, while the rectangle size represents the maximum memory used in the rectangle (a larger size is a larger value).

Choose a server and host from the drop-down menus to view their associated virtual machines. By default, this display shows **Alert Severity**, but you can choose a different metric to display from the **Color Metric** drop-down menu. Use the **Labels** check-box ☒ to include or exclude labels (virtual machine names for each rectangle) in the heatmap. You can hover your mouse over a rectangle to see additional metrics, and you can drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the [“Single VM Summary”](#) display.



Title Bar (possible features are):

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

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

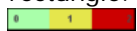












Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

Filter By:


The display might include these filtering options:

- Server:** Select the server for which you want to display data.
- Host** Select the host for which you want to display data.
- VM Count:** The total number of virtual machines in the heatmap display.


Fields and Data:

Labels	Select this check box to include labels in the heatmap.
Log Scale	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Show powered VMs	Select this check box to include only those VMs that are powered on.
Color Metric	Choose a metric to view in the display.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Overall Status	<p>The general health status of the virtual machine.*</p> <ul style="list-style-type: none">  Red indicates that the host is experiencing a problem.  Yellow indicates that the host might have a problem.  Grey indicates that the status of the host's health is unknown.  Green indicates that host's status is OK.
CPU Usage	The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Memory Usage	The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Disk Read KB/s	The amount of data being read from the disk per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Disk Write KB/s	The amount of data being written to the disk per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Net IN KB/s	The amount of network data received per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.


**Net OUT
KB/s**

The amount of network data transmitted per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**% Net IN
Pkts
Dropped**

The percentage of incoming packets that were dropped in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**% Net
OUT Pkts
Dropped**

The percentage of outgoing packets that were dropped in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.


All VMs Table

View data shown in the “[All VMs Heatmap](#)” display, and additional details, in a tabular format. Use this display to view all available data for each virtual machine by server and host.

Each row in the table lists the details for a virtual machine. Choose a server and a host from the drop-down menus to view all virtual machines running on that server/host combination. You can click a column header to sort column data in numerical or alphabetical order.

Drill-down and investigate by clicking a row to view details for the selected virtual machine in the “[Single VM Summary](#)” display.

VMware Virtual Machines - Table										
Server: vSphere2		Host: slesxi-1.sldemos-hq.local								
VM Count: 33		Virtual Machines								
Server	vSphere VM Name	Alert Severity	Alert Count	Overall Status	Heartbeat	Power State	Guest State	Number CPUs	CPU Usage(%)	Memory Usage(%)
vSphere2	2008-DC		0			poweredOff	notRunning	1	NA	
vSphere2	2008S-IT1		1			poweredOff	notRunning	1	NA	
vSphere2	2008S-WIN10		0			poweredOff	notRunning	2	NA	
vSphere2	2008S-WIN11		0			poweredOff	notRunning	1	NA	
vSphere2	2008S-WIN37		0			poweredOn	running	4	25.0	25.0
vSphere2	2008S-WIN38		0			poweredOff	notRunning	1	NA	
vSphere2	2008S-WIN44		1			poweredOn	running	1	14.7	14.7
vSphere2	64BIT-OL7-0		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-1		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-10		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-15		0			poweredOn	running	2	10.5	10.5
vSphere2	64BIT-OL7-15-BACKUP		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-16		0			poweredOn	running	4	3.5	3.5
vSphere2	64BIT-OL7-18		0			poweredOn	running	2	4.8	4.8
vSphere2	64BIT-OL7-2		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-3		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-5		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-6		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-8		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-9		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-UEK4-DOCK...		0			poweredOff	notRunning	4	NA	
vSphere2	64BIT-RH6-3-0		0			poweredOff	notRunning	1	NA	
vSphere2	64BIT-RH6-1		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-RH6-3-1		0			poweredOn	notRunning	1	8.1	8.1
vSphere2	64BIT-RH6-3-5		0			poweredOff	notRunning	1	NA	
vSphere2	64BIT-RH6-6-8		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-RHEL7-1		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-UBUNTU14-1		0			poweredOn	notRunning	4	2.0	2.0
vSphere2	64BIT-W2008S-R2-0		0			poweredOff	notRunning	2	NA	



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

Filter By:

The display might include these filtering options:


Server:	Select the server containing the virtual machines for which you want to view details.
Host	Select the host containing the virtual machines for which you want to view details.
VM Count:	The total number of virtual machines (rows) in the table.
Show powered VMs	Select to include only those VMs that are powered on.

Virtual Machines Table:

Column values describe the virtual machines running on the selected sever/host combination.

Server	The server on which the virtual machine resides.
vSphere VM Name	The name of the vSphere virtual machine.
Alert Severity	<p>The severity of the alert for the virtual machine.</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the virtual machine.
Overall Status	<p>The general health status of the virtual machine.*</p> <ul style="list-style-type: none"> Red indicates that the host is experiencing a problem. Yellow indicates that the host might have a problem. Grey indicates that the status of the host's health is unknown. Green indicates that host's status is OK.

Heartbeat	Displays whether or not the virtual machine has a heartbeat.*  Red indicates that heartbeating has stopped.  Grey indicates that heartbeat status is disabled.  Green indicates that heartbeat status is OK.
Power State	Displays whether or not the virtual machine is powered on.*
Guest State	The state of the guest operating system.*
Number CPUs	The number of CPUs used by the virtual machine.*
CPU Usage %	The percentage (%) of CPUs used.*
Memory Usage (%)	The percentage (%) of memory used by the virtual machine.*
Memory Used (MB)	The amount of used memory, in megabytes.*
Memory Total (MB)	The total amount of memory, in megabytes.*
Disk Reads (KB/sec)	The amount of data being read from the disk per second, in kilobytes.*
Disk Writes (KB/sec)	The amount of data being written to the disk per second, in kilobytes.*
Net IN (KB/sec)	The amount of network data received per second, in kilobytes.*
Net OUT (KB/sec)	The amount of network data transmitted per second, in kilobytes.*
% Packet Loss IN	The percentage of incoming packets that have been lost.*
% Packet Loss OUT	The percentage of outgoing packets that have been lost.*
Packets IN	The total number of incoming packets.*
Packets OUT	The total number of outgoing packets.*
Packet IN Dropped	The number of incoming packets that were dropped.*
Packets OUT Dropped	The number of outgoing packets that were dropped.*
Host	The name of the host.*
Guest Host Name	The name of the guest host.*
Guest IP Address	The IP address of the guest.*
Guest Operating System	The operating system used by the guest.*



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

Filter By:

The display might include these filtering options:

Server	Select the server containing the virtual machine(s) for which you want to view data.
Host	Select the host containing the virtual machine(s) for which you want to view data.
VM	Select a virtual machine for which you want to view data, or select All VMs to view data for all virtual machines on the server/host combination.
Count	Displays the current number of virtual machines listed in the table.

Fields and Data:

Server	The name of the server.*
vSphere VM Name	The name of the virtual machine.*
Disk Name	The name of the disk.*
% Disk Usage	Displays the current percentage of disk space that is being used.*
Capacity (GB)	Displays the total disk capacity, in gigabytes.*
Disk Used (GB)	Displays the total disk space currently being used.*
Disk Free (GB)	Displays the amount of available disk space, in gigabytes.*
Host	Displays the name of the host.*

Expired When checked, performance data for that cluster has not been received in the time specified in the **Duration** region on the RTView Configuration > (Project Name/**MISCMON-LOCAL**) > **Solution Package Configuration** > **VMWare** > **DATA STORAGE** tab.

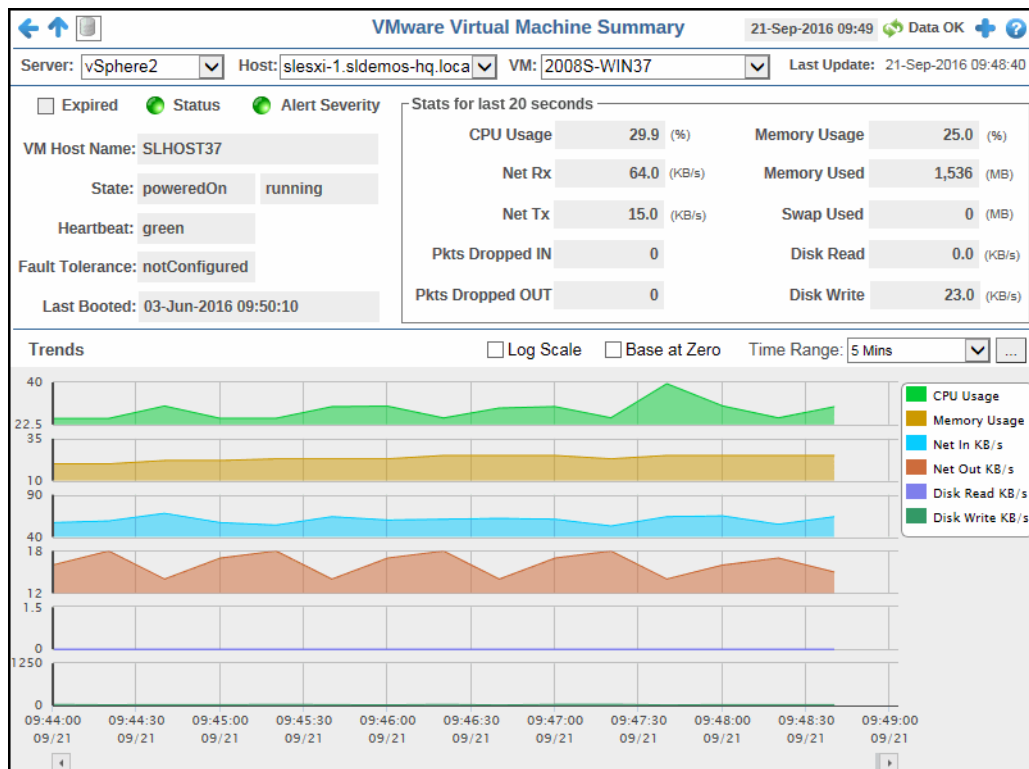
Timestamp The date and time the row data was last updated.


Single VM Summary

View current and historical utilization and performance metrics for a single virtual machine. You can use this display to investigate performance issues for a particular virtual machine.

This display includes trend graphs tracing CPU and memory usage, amount of network data transmitted and received, number of incoming and outgoing packets that have been lost, and disk usage.

Choose a server, host, and virtual machine from the drop-down menus to view details for a specific virtual machine. You can use the **Time-Range** in the **Trends** region to “zoom-in” or “zoom-out” on a specific time frame in the trend graph.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

Filter By:

The display might include these filtering options:

Server	Select the server containing the virtual machine for which you want to view data.
Host	Select the host containing the virtual machine for which you want to view data.
VM	Select the virtual machine for which you want to view data.
Last Update	The date and time that the data in the display was last updated.

Fields and Data:

Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Status	The general health status of the virtual machine.* <ul style="list-style-type: none"> Red indicates that the host is experiencing a problem. Yellow indicates that the host might have a problem. Grey indicates that the status of the host's health is unknown. Green indicates that host's status is OK.
Alert Severity	The current severity of alerts for the virtual machine. <ul style="list-style-type: none"> Red One or more alerts exceeded their ALARM LEVEL threshold. Yellow One or more alerts exceeded their WARNING LEVEL threshold. Green No alert thresholds have been exceeded.
VM Host Name	The name of the host.*
State	Displays whether or not the host/virtual machine is powered on.*
Heartbeat	Displays whether or not the virtual machine has a heartbeat.* <ul style="list-style-type: none"> Red indicates that heartbeating has stopped. Grey indicates that heartbeat status is disabled. Green indicates that heartbeat status is OK.

Fault Tolerance	Displays whether or not fault tolerance is configured (configured/notConfigured).*	
Last Booted	The date and time the virtual machine was last rebooted.*	
Status for last 20 seconds	CPU Usage	The percentage of CPU used in the last 20 seconds.*
	Net Rx	The amount of network data received, in kilobytes per second, in the last 20 seconds.*
	Net Tx	The amount of network data transmitted, in kilobytes per second, in the last 20 seconds.*
	Pkts Dropped IN	The number of incoming packets that were dropped in the last 20 seconds.*
	Pkts Dropped OUT	The number of outgoing packets that were dropped in the last 20 seconds.*
	Memory Usage	The percentage of memory used in the last 20 seconds.*
	Memory Used	The amount of memory used, in megabytes, in the last 20 seconds.*
	Swap Used	The amount of memory swapped, in megabytes, in the last 20 seconds.*
	Disk Read	The amount of data read from the disk, in kilobytes per second, in the last 20 seconds.*
	Disk Write	The amount of data written to the disk, in kilobytes per second, in the last 20 seconds.*

Trend Graphs

Traces the sum of process metrics for the virtual machine.


- **CPU Usage:** The percentage (%) CPU used.
- **Memory Usage:** The amount of memory used.
- **Net In KB/s:** The amount of network data received per second, in kilobytes.
- **Net Out KB/s:** The amount of network data transmitted per second, in kilobytes.
- **Disk Read KB/s:** The amount of data being read from the disk per second, in kilobytes.
- **Disk Write KB/s:** The amount of data being written to the disk per second, in kilobytes.


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


Base at Zero



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

Time Range

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



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

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

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

Datastores View

The displays in this view provide a list of datastores on one or all servers, a list of all hosts mounted to a particular datastore, a list of all virtual machines hosted by a particular datastore, or data for a particular datastore. This View contains the following displays:

- **"All Datastores Table"**: View all datastores, as well as data associated with the datastores, that exist on one server or on all servers.
- **"Hosts by Datastore Table"**: View all hosts that are mounted to a particular datastore.
- **"VMs by Datastore Table"**: View all virtual machines that are hosted by a particular datastore.
- **"Single Datastore Summary"**: View metrics and trend data for a single datastore, as well as those hosts and virtual machines that are using the datastore.

All Datastores Table

View all datastores, as well as data associated with the datastores, that exist on one server or on all servers.

← ↑ VMware Datastores - Table 22-Mar-2017 10:44 Data OK + ?

Server: All Servers ▾

Count: 4 Datastores

Server	Name	Overall Status	Maintenance Mode	Type	Accessible?	Multiple Hosts?	% Utili
qavSphere1	datastore1	●	normal	VMFS	✓	<input type="checkbox"/>	
qavSphere1	datastore1 (1)	●	normal	VMFS	✓	<input type="checkbox"/>	
qavSphere1	datastore-A1	●	normal	VMFS	✓	✓	
qavSphere1	datastore-A2	●	normal	VMFS	✓	✓	

< >

Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

Filter By:

The display might include these filtering options:

Server	Select the server for which you want to view data.
Count	The total number of datastores on the selected server(s), which are listed in the Datastores table.

Datastores Table:

Server	The name of the server.
Name	The name of the datastore.
Overall Status	<p>The general health status of the datastore.*</p> <ul style="list-style-type: none">  Red indicates that the datastore is experiencing a problem.  Yellow indicates that the datastore might have a problem.  Grey indicates that the status of the datastore's health is unknown.  Green indicates that datastore's status is OK.
Maintenance Mode	Lists current maintenance mode state of the datastore (normal, inMaintenance, enteringMaintenance).*
Type	Lists the type of file system volume, such as VMFS or NFS.*
Accessible?	The connectivity status of the datastore. When checked, indicates that the datastore is accessible.*
Multiple Hosts	When checked, indicates that more than one host has been configured with access to the datastore.*
% Utilization	Lists the current space utilization percentage for the datastore.*
Capacity (GB)	Displays the maximum capacity of the datastore, in gigabytes.*
Free Space (GB)	Displays the amount of available space in the datastore, in gigabytes.*
Space Uncommitted (GB)	Displays the amount of total additional storage space potentially used by all virtual machines on this datastore, in gigabytes.*
Expired	<p>When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.</p>
Timestamp	The date and time the row data was last updated.

Hosts by Datastore Table

View all hosts that are mounted to a particular datastore.

Server	Datastore Name	Host Name	Access Mode	Accessible?	Mounted?	
qavSphere1	datastore-A1	slesxi-1.sldemos-hq.local	readWrite	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	/vmfs/volumes/511bac...
qavSphere1	datastore-A1	slesxi-2.sldemos-hq.local	readWrite	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	/vmfs/volumes/511bac...

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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.


Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

Filter By:

The display might include these filtering options:

Server

Select the server containing the datastore for which you want to view data, or select **All Servers**.



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

Filter By:

The display might include these filtering options:

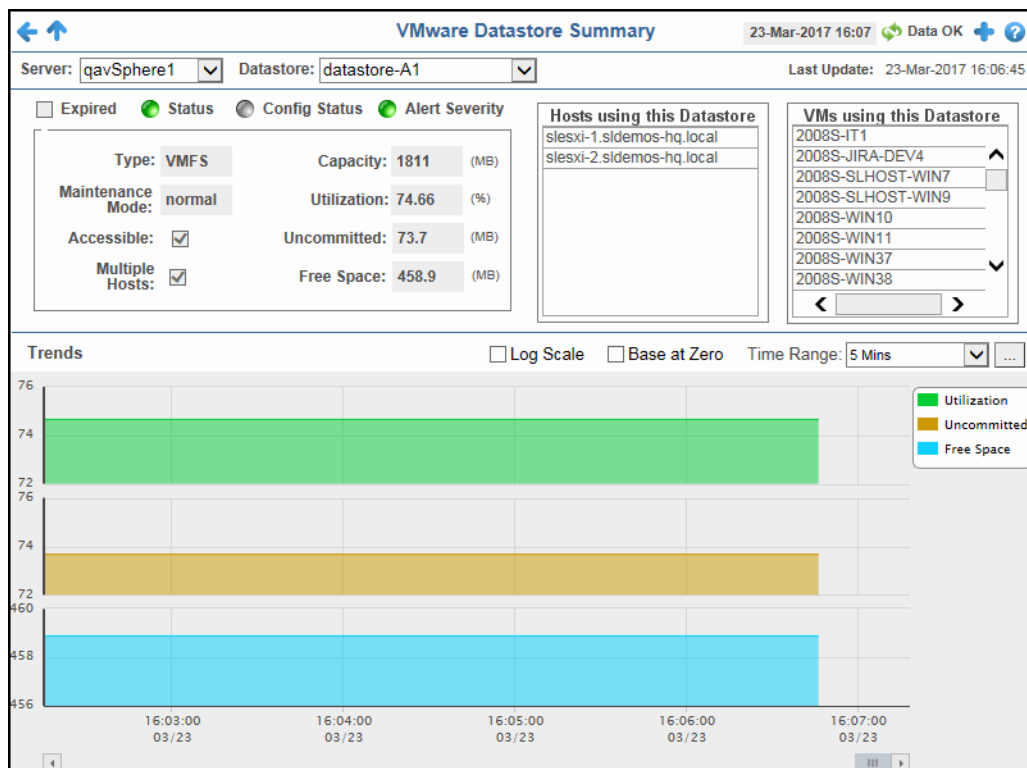
Server	Select the server containing the datastore for which you want to view data, or select All Servers .
Datastore	Select the datastore for which you want to view data.
Count	The total number of virtual machines connecting to a datastore, which are listed in the table.

VMs Hosted by Datastore Table:

Server	The name of the server.
Datastore Name	The name of the datastore.
VM Name	The name of the virtual machine hosted by the datastore.
Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Timestamp	The date and time the row data was last updated.

Single Datastore Summary

View metrics and trend data for a single datastore, as well as those hosts and virtual machines that are using the datastore.



Title Bar (possible features are):

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

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

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

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












Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vSphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

Filter By:

The display might include these filtering options:

Server	Select the server containing the datastore for which you want to view data.
Datastore	Select the datastore for which you want to view data.
Last Update	The date and time the data in the display was last updated.

Fields and Data:

Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Status	<p>The general health status of the datastore.*</p> <ul style="list-style-type: none">  Red indicates that the datastore is experiencing a problem.  Yellow indicates that the datastore might have a problem.  Grey indicates that the status of the datastore's health is unknown.  Green indicates that datastore's status is OK.
Config Status	<p>Indicates whether or not the system has detected a configuration issue involving the datastore.</p> <ul style="list-style-type: none">  Red indicates that a problem has been detected involving the datastore.  Yellow indicates a problem is about to occur or a transient condition has occurred.  Grey indicates that configuration status of the datastore is not being monitored.  Green indicates that no configuration issues have been detected.
Alert Severity	<p>The current severity of alerts for the datastore.</p> <ul style="list-style-type: none">  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
Type	Lists the type of file system volume, such as VMFS or NFS.*
Maintenance Mode	Lists current maintenance mode state of the datastore (normal, inMaintenance, enteringMaintenance).*
Accessible?	The connectivity status of the datastore. When checked, indicates that the datastore is accessible.*
Multiple Hosts	When checked, indicates that more than one host has been configured with access to the datastore.*
Capacity (MB)	Displays the maximum capacity of the datastore, in megabytes.*
Utilization (%)	Lists the current space utilization percentage for the datastore.*
Uncommitted (MB)	Displays the amount of total additional storage space potentially used by all virtual machines on this datastore, in megabytes.*
Free Space (MB)	Displays the amount of available space in the datastore, in megabytes.*

Hosts Using this Datastore -- Lists the hosts using the datastore.

VMs using this Datastore -- Lists the virtual machines using the datastore.

Trend Graphs

Traces the sum of process metrics for the virtual machine.

- **Utilization:** Traces the current space utilization percentage for the datastore.
- **Uncommitted:** Traces the amount of total additional storage space potentially used by all virtual machines on this datastore, in megabytes.
- **Free Space:** Traces the amount of available space in the datastore, in megabytes.


Log Scale

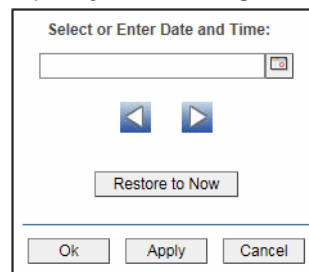
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log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


Base at Zero



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

Time Range

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



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

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

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

Networks View

The display in this View lists all networks, as well as data associated with the networks, that exist on one server or on all servers. The available display in this View is:

- **"All Networks Table":** View all networks, as well as data associated with the networks, that exist on one server or on all servers.

All Networks Table

View all networks, as well as data associated with the networks, that exist on one server or on all servers.

VMware Networks - Table 22-Mar-2017 11:37 Data OK

Server: All Servers

Count: 1

Server	Network Name	Overall Status	Accessible?	IP Pool Name	Expired	Timestamp
qavSphere1	VM Network		<input checked="" type="checkbox"/>		<input type="checkbox"/>	22-Mar-2017 11:36:4

Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

Filter By:

The display might include these filtering options:

- Server** Select the server for which you want to view data.
- Count** The total number of networks on the selected server(s), which are listed in the table.

VM Networks Table:

Server	The name of the server.
Network Name	The name of the network.
Overall Status	<p>The general health status of the network.*</p> <p> Red indicates that the network is experiencing a problem.</p> <p> Yellow indicates that the network might have a problem.</p> <p> Grey indicates that the status of the network's health is unknown.</p> <p> Green indicates that network's status is OK.</p>
Accessible?	The connectivity status of the virtual machine. When checked, indicates that the virtual machine is accessible.*
IP Pool Name	Lists the name of the IP pool that is assigned to the network.*
Expired	<p>When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.</p>
Timestamp	The date and time the row data was last updated.

Events/Alarms View

The displays in this View allow you to view event data and alarm data for one server or for all servers. Available displays in this View are:

- ["All Events Table"](#): View all events, as well as data associated with the events, that exist on one server or on all servers.
- ["All Alarms Table"](#): View all alarms, as well as data associated with the alarms, that exist on one server or on all servers.

All Events Table

View all events, as well as data associated with the events, that exist on one server or on all servers.

← ↑ VMware Events - Table 22-Mar-2017 11:44 Data OK + ?

Server: All Servers

Count: 245

Server	Created Time	Event ID	Chain ID	Event Class	Event Type	User Name	
qavSphere1	22-Mar-2017 08:47:22	184029	184029	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 08:42:54	184028	184028	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 08:41:24	184027	184027	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 08:24:55	184026	184026	UserLogoutSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 07:54:50	184025	184025	UserLoginSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 07:54:50	184024	184024	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 07:54:20	184023	184023	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 07:22:56	184022	184022	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 07:10:32	184021	184021	UserLogoutSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 07:02:20	184020	184020	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:52:56	184019	184019	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:40:29	184018	184018	UserLoginSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 06:40:28	184017	184017	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:38:13	184016	184016	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:35:38	184015	184015	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 00:35:09	184014	184014	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 00:07:02	184013	184013	UserLoginSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 18:06:58	184012	184012	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:44:07	184011	184011	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:34:10	184010	184010	UserLoginSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:21:10	184009	184009	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:10:00	184008	184008	UserLoginSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:03:30	184007	184007	AlarmStatusChangedEvent	info		Alar
qavSphere1	21-Mar-2017 17:03:09	184006	184006	AlarmStatusChangedEvent	info		Alar
qavSphere1	21-Mar-2017 17:01:49	184005	184005	AlarmStatusChangedEvent	info		Alar
qavSphere1	21-Mar-2017 15:02:39	184004	184004	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 14:53:56	184003	184003	UserLogoutSessionEvent	info	vmwarehost1	Use

Page 1 of 2 1 - 200 of 245 items

Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

Filter By:

The display might include these filtering options:

Server	Select the server for which you want to view data.
Count	The total number of events on the selected server(s), which are listed in the Events table.

Events Table:

Server	The name of the server.
Created Time	The date and time the event was created.*
Event ID	The ID of the event.*
Chain ID	The parent or group ID.*
Event Class	The type of event class.*
Event Type	The type of event.*
User Name	The user who caused the event.*
Message Text	A formatted text message describing the event.*
Host	The host object of the event.*
Virtual Machine	The event's virtual machine.*
Compute Resource	The event's compute resource.*
Datacenter	The event's datacenter.*
Datastore	The event's datastore.*
Distributed Virtual Switch	The event's DistributedVirtualSwitch.*
Network	The network associated with the event.*
Expired	When checked, performance data for that cluster has not been received in the time specified in the Duration region on the RTView Configuration > (Project Name/ MISCMON-LOCAL) > Solution Package Configuration > VMWare > DATA STORAGE tab.
Timestamp	The date and time the row data was last updated.

All Alarms Table

View all alarms, as well as data associated with the alarms, that exist on one server or on all servers.

← ↑ VMware Alarms - Table 22-Mar-2017 11:48 Data OK + ?

Server: All Servers ▾

Count: 820

Alarms					
Alarm Time	Server	Managed Object	Overall Status	Alarm Name	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host connection and power state	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Timed out starting Secondary VM	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	No compatible host for Secondary VM	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host processor status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host memory status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host hardware fan status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host hardware voltage	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host hardware temperature status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host hardware power status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host hardware system board status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host battery status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Virtual Machine Fault Tolerance vLockStep int	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Status of other host hardware objects	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host storage status	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host error	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Virtual machine error	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host connection failure	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Cannot connect to storage	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Migration error	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Exit standby error	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	License error	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Health status changed alarm	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Host cpu usage	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Virtual machine Fault Tolerance state change	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Network connectivity lost	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Network uplink redundancy lost	
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	🟡	Network uplink redundancy degraded	

◀ ▶ Page 1 of 5 1 - 200 of 820 items

Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- 🔔 Open the Alert Views - RTView Alerts Table display.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vSphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

Filter By:

The display might include these filtering options:

Server	Select the server for which you want to view data.
Count	The total number of datastores on the selected server(s), which are listed in the Datastores table.

Alarms Table:

AlarmTime	The date and time of the alarm.*
Server	The name of the server in which the alarm occurred.*
Managed Object	The name of the managed object.*
Overall Status	<p>The general health status of the alarm.*</p> <ul style="list-style-type: none">  Red indicates that the alarm is experiencing a problem.  Yellow indicates that the alarm might have a problem.  Grey indicates that the status of the alarm's health is unknown.  Green indicates that alarm's status is OK.
Alarm Name	The name of the alarm.*
Acknowledged	Lists whether or not the alarm has been acknowledged.*
Acknowledged By	Lists the user who acknowledged the alarm.*
Acknowledged Time	Lists the date and time when the alarm was acknowledged.*
Description	The description of the alarm.*
Timestamp	The date and time the row data was last updated.

CHAPTER 5 RTView DataServer for Kafka

The RTView DataServer for Kafka provides a way to create connections and modify default configuration settings for the Kafka solution package and sends collected data to RTView Central, which contains the displays associated with the RTView DataServer for Kafka that help you to monitor the health and performance across your Kafka components.

RTView Central contains the following Views and their associated displays that will be populated with data collected via the RTView DataServer for Kafka:

- ["Apache Kafka"](#)

The RTView *DataCollector* for Kafka is also available for use with the RTView DataServer for Kafka. RTView DataCollector for Kafka is used for collecting data and sending it to one or more RTView DataServers. The RTView DataCollector for Kafka is useful if you need to distribute data collection.

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

Apache Kafka

The following Apache Kafka Views (and their associated displays) can be found under **Components** tab > **Middleware** > **Kafka**.

The following views are available:

- **"Kafka Clusters View"**: The displays in this View allow you to view metrics for all Kafka clusters and view the performance metrics for all servers on a particular cluster.
- **"Kafka Topics View"**: This displays in this View allow you to view metrics for all topics for a particular broker in heatmap/table format, view current and trend data for a single topic, view the metrics for all topics on a particular cluster, and view metrics for a particular topic on a particular cluster.
- **"Kafka Brokers View"**: The displays in this View allow you to view the current and historical metrics for all brokers in heatmap/table formats, view various metrics for a particular broker, and view metrics and trend data for a particular broker.
- **"Kafka Zookeepers View"**: The displays in this View allow you to view the current and historical metrics for all zookeepers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single zookeeper.
- **"Kafka Producers View"**: The displays in this View allow you to view the current and historical metrics for all producers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single producer.
- **"Kafka Consumers View"**: The displays in this View allow you to view the current and historical metrics for all consumers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single consumers.

Kafka Clusters View

These displays allow you to view metrics for all Kafka clusters and view the performance metrics for all servers on a particular cluster. Displays in this View are:

- **"All Clusters Table"**: A tabular view of all clusters and their associated metrics.
- **"Cluster Performance"**: This display allows you to view performance metrics for all servers on a particular cluster.

All Clusters Table

The table in this display provides a view of all of your clusters and their associated metric data including maximum alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected cluster in the RT Cluster Performance display.

Cluster Name	MaxSeverity	AlertCount	# Brokers Monitored	# Brokers Running	# Active Controllers	Offli
SL-QA-Cluster1		0	3	2	1	
SL-QA-Cluster2		0	3	2	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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

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

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


Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Fields and Data:

Count	Lists the number of brokers found as a result of the cluster that was selected and displayed in the Kafka Brokers table.
--------------	--


Kafka Clusters Table:


cluster_name	The name of the cluster.*
MaxSeverity	<p>The current highest alert severity for any of clusters.</p> <p> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</p> <p> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</p> <p> Green indicates that no metrics have exceeded their alert thresholds.</p>
Alert Count	The total number of alerts for the host.
# Brokers Monitored	The current number of brokers being monitored for the cluster.*
# Brokers Running	The number of brokers currently running on the cluster.*
# Active Controllers	The number of active controllers on the cluster.*
Offline Partitions	The number of partitions without an active leader on the cluster.*
Under Replicated Partitions	The number of partition replicas that are out of sync (total number of replicas minus the total number of in-sync replicas) on all brokers on the cluster.*
% Max Deviation in Partition Count	The percentage of maximum deviation in partition count.*
% Max Deviation in Leader Count	The percentage of maximum deviation in leader count.*
# Zookeepers	The number of zookeepers on the cluster.*
# Zookeeper Connections	The number of connections on the zookeepers in the cluster.*
# Zookeeper Outstanding Reqs	The number of outstanding requests on the zookeepers in the cluster.
# Zookeeper Pkts Recvd	The number of packets received on the zookeepers in the cluster.*
# Zookeeper Pkts Sent	The number of packets sent by the zookeepers in the cluster.*
# Consumers	The number of consumers on the cluster.*
Bytes Consumed Rate	The rate of bytes being consumed by the consumers.
Records Consumed Rate	The rate of records being consumed by the consumers.

- # Producers** The number of producers on the cluster.
- Producer In
Byte Rate** The rate of incoming bytes for the producers.*
- Producer Out
Byte Rate** The rate of outgoing bytes for the producers.*
- Producer
Record Send
Rate** The rate of records being sent for the producers.*

Cluster Performance

This display provides a view of the current metrics for the brokers contained in a selected cluster.





Filter By:

The display might include these filtering options:

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

Fields and Data:

- | | |
|--|--|
| Broker Count | The number of brokers contained in the selected cluster. |
| Partitions Under Rep or Offline | Lists the number of partitions that are under-replicated or offline on each broker in the cluster. |
| % CPU | Lists the percentage of CPU used by the broker. |
| % Memory | Lists the percentage of memory used by the broker. |
| Msgs In/sec | Lists the rate of incoming messages (per second) for each broker in the cluster. |
| IO KBytes/sec | Lists the rate of incoming kilobytes (per second) for each broker in the cluster. |
| Total Time Ms | Lists the total time taken to service a request.* |
| Log Flush Latency 95P | Lists the 95th percentile value for the log flush latency for each broker on the cluster. |

Kafka Topics View

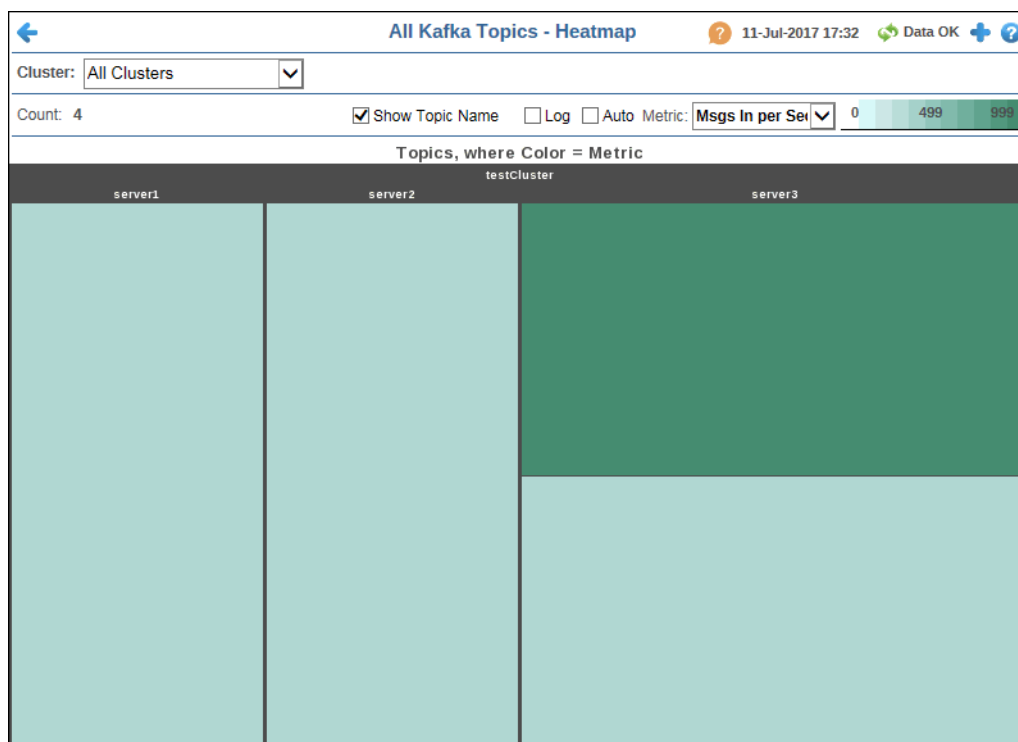
These displays allow you to view metrics for all Kafka topics on a particular topic in heatmap/table format, view the performance metrics for a single topic on a particular broker, view the metrics for all topics on a particular cluster, and view metrics for a particular topic on a particular cluster. Displays in this View are:

- [“All Topics Heatmap”](#): A heatmap view of all topics for a particular cluster.
- [“All Topics Table”](#): A tabular view of all topics for a particular cluster.
- [“Single Topic Summary”](#): This display allows you to view current metrics and trend data for a single topic.
- [“All Topics for Cluster”](#): This display allows you to view performance metrics for all topics on a particular cluster.
- [“Single Topic for Cluster”](#): This display allows you to view performance metrics for a particular topic on a particular cluster.

All Topics Heatmap

This heatmap provides an easy-to-view interface that allows you to quickly identify the current status of each of your topics for each available metric. You can view the topics in the heatmap based on the following metrics: the rate of incoming messages, the rate of incoming bytes, the rate of outgoing bytes, the rate of rejected bytes, the rate of total fetch requests, the rate of failed fetch requests, the rate of total produce requests, and the rate of failed produce requests. By default, this display shows the heatmap based on the **Msgs In per Sec** metric.

You can use the **Show Topic Name** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a topic. Clicking one of the rectangles in the heatmap opens the [“Single Topic Summary”](#) display, which allows you to see additional details for the topic metrics for the selected topic.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data:

- Cluster** Select the cluster for which you want to view data.
- Count** Lists the number of topics displayed in the heatmap.
- Show Topic Name** Select this check box to display the names of the topics at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.

Metric

Choose a metric to view in the display.

Msgs In per Sec

The rate of incoming messages (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.

0 499 999

Bytes In per Sec

The rate of incoming bytes (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming bytes per second. The middle value in the gradient bar indicates the middle value of the range.

0 80,895 161,790

Bytes Out per Sec

The rate of outgoing bytes (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of outgoing bytes per second. The middle value in the gradient bar indicates the middle value of the range.

0 80,895 161,790

Bytes Rejected per Sec

The rate of bytes being rejected (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of bytes rejected per second. The middle value in the gradient bar indicates the middle value of the range.

0 50 100

Total Fetch Requests per Sec

The rate of fetch requests (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of fetch requests per second. The middle value in the gradient bar indicates the middle value of the range.

0 465 930

Failed Fetch Requests per Sec

The rate of failed fetch requests (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of failed fetch requests per second. The middle value in the gradient bar indicates the middle value of the range.

0 50 100

Total Produce Requests per Sec

The rate of total producer requests (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of produce requests per second. The middle value in the gradient bar indicates the middle value of the range.

0 233.1 466.2

Failed Produce Requests per Sec

The rate of failed producer requests (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of failed produce requests per second. The middle value changes accordingly to indicate the color of the middle value of the range.

0 50 100

All Topics Table

The table in this display provides a view of all of your topics for a particular cluster and their associated metric data. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected topic in the ["Single Topic Summary"](#) display.

← ↑ All Kafka Topics - Table 29-Aug-2017 11:09 Data OK + ?

Cluster: SL-QA-Cluster1


Count: 2 MeanRate


Cluster Name	Topic	BytesInPerSec	BytesOutPerSec	BytesRejectedPerSec	FailedFetchRequestsPe
SL-QA-Cluster1	my-replicated-topic	161,756.84	161,757.25	0.00	
SL-QA-Cluster1	my-replicated-topic1	161,756.80	165,651.96	0.00	

< >

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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By

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

Kafka Topics Table:

Count The total number of topics listed in the table.

**<Rate Drop
Down List>**

Select the option for which you want to view data:

MeanRate -- Select this option to view the average rate for each metric for the topics in the display.

One Minute -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the topics in the display.

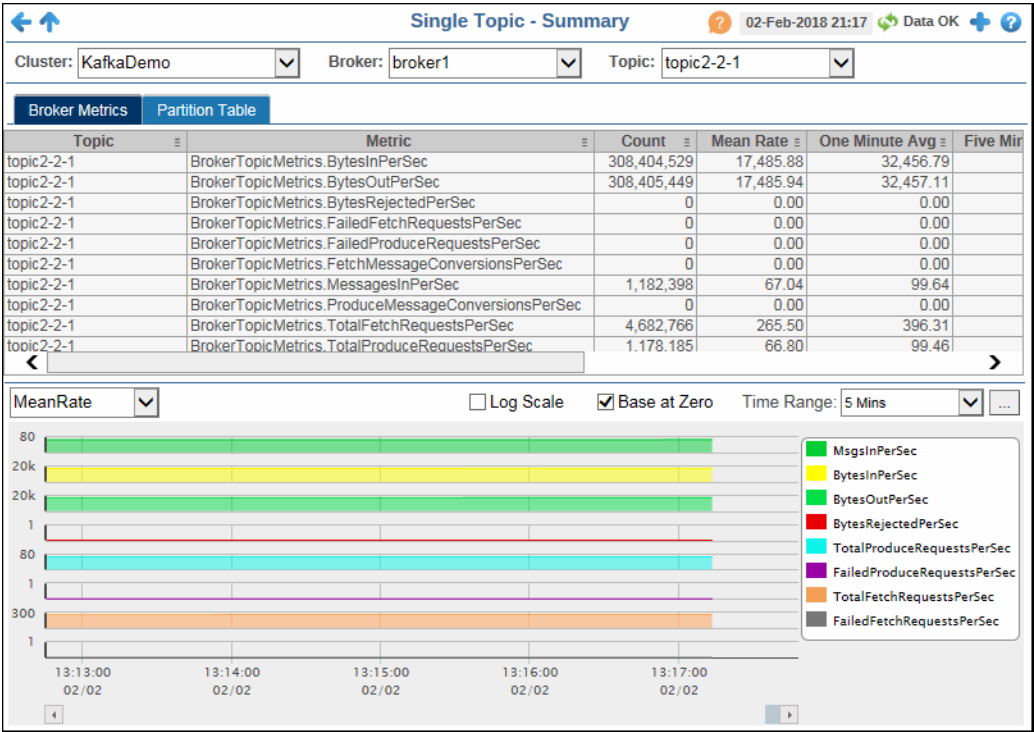
Five Minute -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the topics in the display.

Fifteen Minute -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the topics in the display.

Cluster Name	Lists the name of the cluster.
Topic	Lists the name of the topic.
Bytes In Per Sec	The rate of incoming bytes.
Bytes Out Per Sec	The rate of outgoing bytes.
Bytes Rejected Per Sec	The rate of rejected bytes.
Failed Fetch Requests Per Sec	The rate of failed fetch requests.
Failed Produce Requests Per Sec	The rate of failed produce requests.
Messages In Per Sec	The rate of incoming messages.
Total Fetch Requests Per Sec	The rate of total fetch requests.
Total Produce Requests Per Sec	The rate of total produce requests.
Replicas	The total number of replicas associated with the topic.
Replicas Out Of Sync	The number of replicas that are out of sync for the topic.

Single Topic Summary

This display provides a view of the current metrics and trend data for a single topic. Selecting the **Cluster/Broker** combination populates the **Topic** drop down list, from which you can select the topic for which you want to view metric and trend data.



Title Bar (possible features are):

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

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

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

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

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

Broker Select the broker for which you want to see data.

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

Broker Metrics Tab

Broker Metrics		Partition Table			
Topic	Metric	Count	Mean Rate	One Minute Avg	Five Min
topic2-2-1	BrokerTopicMetrics.BytesInPerSec	312,980,410	17,605.18	32,577.67	
topic2-2-1	BrokerTopicMetrics.BytesOutPerSec	312,981,330	17,605.24	32,577.70	
topic2-2-1	BrokerTopicMetrics.BytesRejectedPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.FailedFetchRequestsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.FailedProduceRequestsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.FetchMessageConversionsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.MessagesInPerSec	1,196,438	67.30	99.96	
topic2-2-1	BrokerTopicMetrics.ProduceMessageConversionsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.TotalFetchRequestsPerSec	4,738,723	266.55	398.50	
topic2-2-1	BrokerTopicMetrics.TotalProduceRequestsPerSec	1,192,206	67.06	99.81	

Topic The name of the topic.

Metric The name of the metric.

Count The total number of the particular metric.

Mean Rate The mean rate of the metric.

One Minute Avg The rate of incoming messages (per second) averaged over a one minute period for the metric, based on the **Rate Units**.

Five Minute Avg The rate of incoming messages (per second) averaged over a five minute period for the metric, based on the **Rate Units**.

Fifteen Minute Avg The rate of incoming messages (per second) averaged over a fifteen minute period for the metric, based on the **Rate Units**.

Event Type The event type for the topic metric.

Rate Units The unit of measure used to calculate the **Mean Rate**, the **One Minute Avg**, the **Five Minute Avg**, and the **Fifteen Minute Avg**.

Expired When checked, performance data in the row has not been received within the time specified (in seconds) in the **Expire Time** field in the RTView Configuration Application > (**KAFKAMON-LOCAL/Project Name**) > **Solution Package Configuration** > **Apache Kafka** > **DATA STORAGE** > **Duration** > **Expire Time** property. The RTView Configuration Application > (**KAFKAMON-LOCAL/Project Name**) > **Solution Package Configuration** > **Apache Kafka** > **DATA Storage** > **Duration** > **Delete Time** property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

For example, if **Expire Time** was set to 120 and **Delete Time** was set to 3600, then the **Expired** check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.

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

Rate Trends

Select the option for which you want to view data:

MeanRate -- Select this option to view the average rate for each metric in the trend graph.

One Minute-- Select this option to view the rate averaged over a one minute period for each metric in the trend graph.

Five Minute -- Select this option to view the rate averaged over a five minute period for each metric in the trend graph.

Fifteen Minute -- Select this option to view the rate averaged over a fifteen minute period for each metric in the trend graph.

Traces the following:

MsgsInPerSec -- traces the selected rate of incoming messages.

BytesInPerSec -- traces the selected rate of incoming bytes.

BytesOutPerSec -- traces the selected rate of outgoing bytes.

BytesRejectedPerSec -- traces the selected rate of rejected bytes.

TotalProduceRequestsPerSec -- traces the selected rate of total produce requests.

FailedProduceRequestsPerSec -- traces the selected rate of failed produce requests.

TotalFetchRequestsPerSec -- traces the selected rate of total fetch requests.

FailedFetchRequestsPerSec -- traces the selected rate of failed fetch requests.


Log Scale

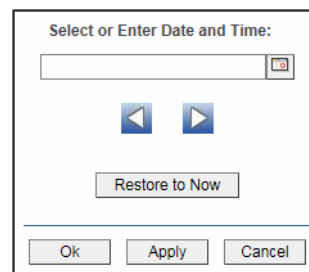
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log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


Base at Zero



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

Time Range

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

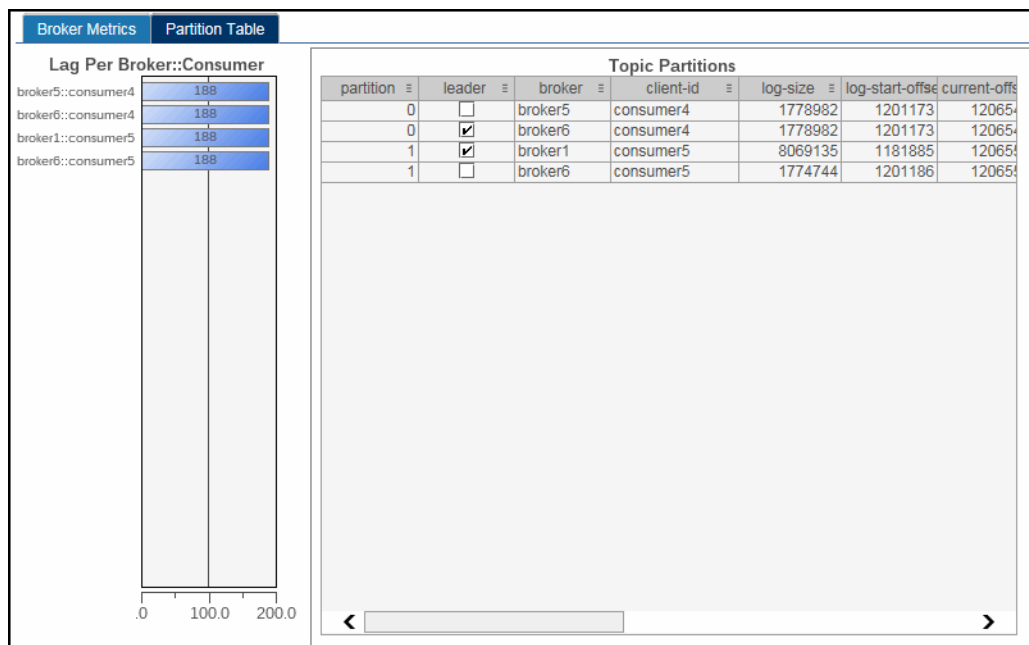


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Partition Table Tab



Lag Per Broker::Consumer

Lag per partition, where the partitions are identified by the broker hosting the partition and the consumer reading the partition.

Topic Partitions Table

partition	The name of the partition.
leader	When checked, signifies that the broker is a leader on the partition.
broker	The name of the broker.
client-id	The ID of the consumer reading the topic.*
log-size	The current number of messages in the log.*
log-start-offset	The offset of the first message written to a log.*
current-offset	The offset of the message currently being consumed.*
log-end-offset	The offset of the last message written to a log.*
lag	The difference between the current consumer position in the partition and the end of the log.*
lag-delta	The difference in the amount of lag from the previous polling period to the current polling period.*
lag-rate	The rate of change in the amount of lag.*
current-delta	The difference between the current consumer position in the partition from the previous polling period to the current polling period.*
current-rate	The rate of change of the current consumer position.*

- log-end-delta


The difference between the offset of the last message in the partition from the previous polling period to the current polling period.*
- log-end-rate


The rate of change of the last message offset.*
- time_stamp

The date and time the row data was last updated.

All Topics for Cluster

This display provides a view of the activity metrics on all topics for a particular cluster. You can view the metrics based on the mean rate, a 1 minute average rate, a 5 minute average rate, or a 15 minute average rate.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

Cluster

Select the cluster for which you want to see data.

<Rate Drop Down List>

Select the option for which you want to view data:

Mean Rate -- Select this option to view the average rate for each metric for the topics in the display.

1 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the topics in the display.

5 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the topics in the display.

15 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the topics in the display.

Topic Count

The number of topics found in the cluster.

Sort By

Select the metric by which you want to sort the data in the display. When using this option with the **Sort Descending** toggle, the topics will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select **Msgs In/sec** from this drop down and select the **Sort Descending** toggle, the topics listed in the display will be sorted so that the topic with the most **Msgs In/sec** will be listed at the top followed by the topic with the next most **Msgs In/sec**, and so on.

Sort Descending

When toggled on, the topics listed in the display are sorted in descending order based on the selected metric in the **Sort By** drop down. When toggled off, the topics are listed in ascending order.

Cluster Activity for Each Topic:

Consumer Lag

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

KBytes In/sec

The number of incoming kilobytes per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of incoming kilobytes per second for 1 minute.

KBytes Out/ sec

The number of outgoing kilobytes per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of outgoing kilobytes per second for 1 minute.

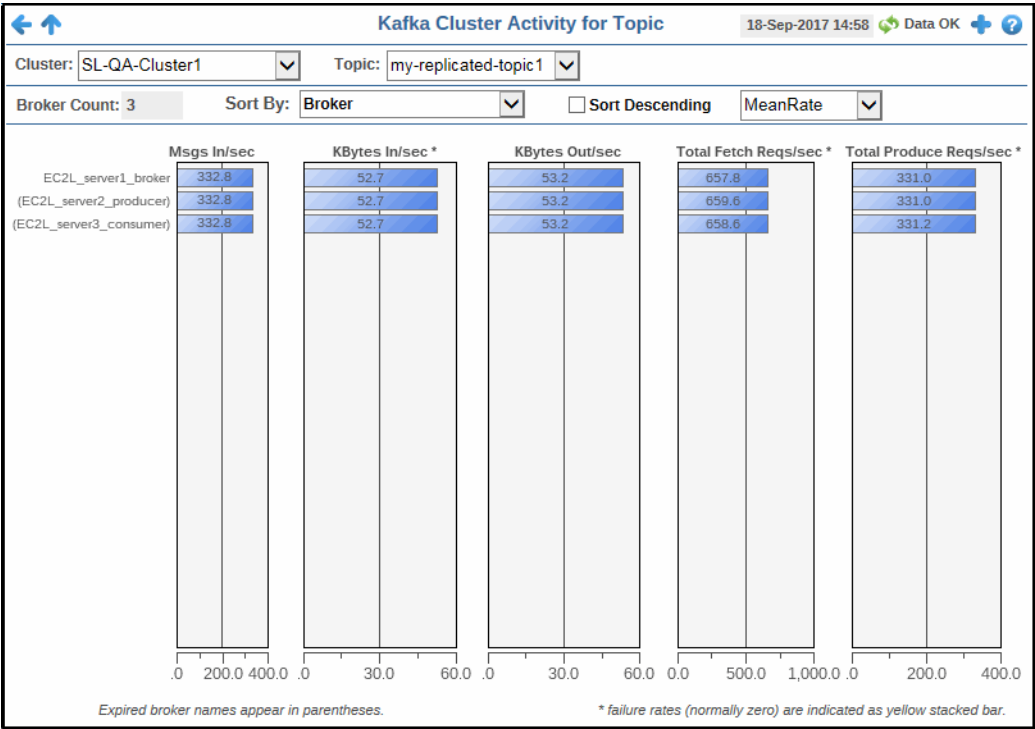
- Total Fetch Reqs/sec**


The total number of fetch requests per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of fetch requests per second for 1 minute.
- Total Produce Reqs/sec**

The total number of produce requests per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of producer requests per second for 1 minute.

Single Topic for Cluster

This display provides a view of the activity metrics on all brokers for a particular topic. You can view the metrics based on the mean rate, a 1 minute average rate, a 5 minute average rate, or a 15 minute average rate.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

- Cluster** Select the cluster for which you want to see data.
- Topic** Select the topic for which you want to see data.
- <Rate Drop Down List>** Select the option for which you want to view data:
 - Mean Rate** -- Select this option to view the average rate for each metric for the brokers in the display.
 - 1 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the brokers in the display.
 - 5 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the brokers in the display.
 - 15 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the brokers in the display.

Broker Count The number of brokers found in the cluster with the associated topic.

Sort By Select the metric by which you want to sort the data in the display. When using this option with the **Sort Descending** toggle, the brokers will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select **Msgs In/sec** from this drop down and select the **Sort Descending** toggle, the brokers listed in the display will be sorted so that the broker with the most **Msgs In/sec** will be listed at the top followed by the broker with the next most **Msgs In/sec**, and so on.

Sort Descending When toggled on, the brokers listed in the display are sorted in descending order based on the selected metric in the **Sort By** drop down. When toggled off, the brokers are listed in ascending order.

Cluster Activity for Each Broker:

- Msgs In/sec** The number of incoming messages per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of incoming messages per second for 1 minute.
- KBytes In/sec** The number of incoming kilobytes per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of incoming kilobytes per second for 1 minute.

KBytes Out/sec	The number of outgoing kilobytes per second. For example, if you select 1 Minute Avg from the drop down list, the average rate of outgoing kilobytes per second for 1 minute.
Total Fetch Reqs/sec	The total number of fetch requests per second. For example, if you select 1 Minute Avg from the drop down list, the average rate of fetch requests per second for 1 minute.
Total Produce Reqs/sec	The total number of produce requests per second. For example, if you select 1 Minute Avg from the drop down list, the average rate of producer requests per second for 1 minute.

Kafka Brokers View


These displays provide detailed data for all brokers in heatmap and tabular form, provide details for all metrics for a particular broker in tabular form, and provide JVM runtime and broker status details for a particular broker. Displays in this View are:

- **"All Brokers Heatmap"**: A heatmap view of all brokers in a heatmap format and their associated metrics.
- **"All Brokers Table"**: A tabular view of your brokers and their associated metrics.
- **"All Broker Metrics Table"**: A tabular and trend graph view of meter metrics, histogram metrics, and timer metrics for a particular broker.
- **"Single Broker Summary"**: Contains JVM runtime data, broker status, topic, and topic trend details for a particular broker.

All Brokers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your brokers for each available metric. You can view the brokers in the heatmap based on the following metrics: the current alert severity, the current alert count, the under replicated partitions count, the offline partitions count, the rate of incoming messages, the rate of incoming bytes, the rate of outgoing bytes, and the log flush latency value. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Show Broker Name** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a broker. Clicking one of the rectangles in the heatmap opens the **"Single Broker Summary"** display, which allows you to see additional details for the selected broker.



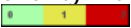
Title Bar (possible features are):




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

Fields and Data:


- Cluster** Select the cluster for which you want to view data.
- Show Broker Name** Select this check box to display the names of the brokers at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

Alert Severity

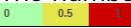
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

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

Alert Count

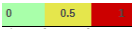
The total number of critical and warning unacknowledged alerts in the brokers. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient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

Under Replicated Partitions

The number of under-replicated partitions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerUnderReplicatedPartns**. The middle value in the gradient bar indicates the middle value of the range.

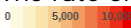
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Offline Partitions

The number of offline partitions. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerOfflinePartitionCnt**. The middle value in the gradient bar indicates the middle value of the range.


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

Msgs In Per Sec

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

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

Bytes In Per Sec

The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerBytesInPerSec**. The middle value in the gradient bar indicates the middle value of the range.

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

Bytes Out Per Sec

The rate of outgoing bytes (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerBytesOutPerSec**. The middle value in the gradient bar indicates the middle value of the range.

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

Log Flush Latency 95 Pctile


The log flush latency for the top five percent of values. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerLogFlushLatency95P**. The middle value in the gradient bar indicates the middle value of the range.

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

All Brokers Table

The table in this display provides a view of all of your brokers and their associated metric data including cluster name, broker name, broker ID, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected broker in the [“Single Broker Summary”](#) display.

All Kafka Brokers - Table							
Cluster: All Clusters							
Count: 3							
Kafka Brokers							
Cluster Name	Broker Name	Broker ID	Alert Level	Alert Count	Broker State	Active Controller	CPU %
KafkaTest	broker1	0	0	0	RunningAsBroker	<input checked="" type="checkbox"/>	12.0%
KafkaTest	broker2	1	0	0	RunningAsBroker	<input type="checkbox"/>	1.5%
KafkaTest	broker3	2	0	0	RunningAsBroker	<input type="checkbox"/>	5.9%



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By

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

Count Lists the number of brokers found as a result of the cluster that was selected and displayed in the Kafka Brokers table.

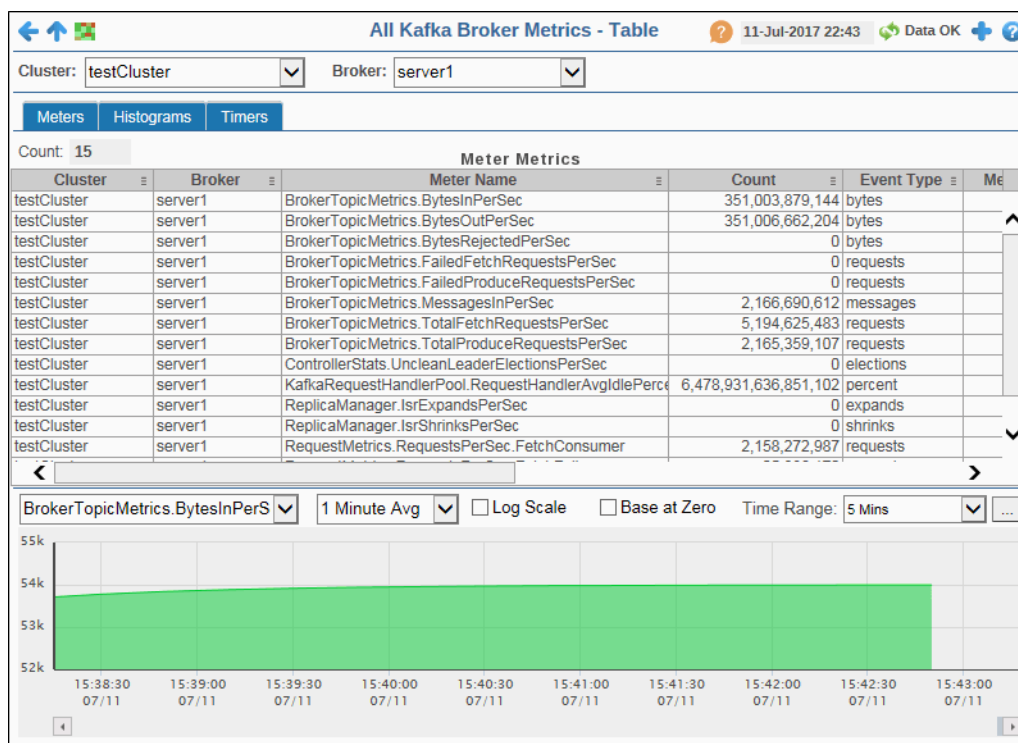
Kafka Brokers Table:

Cluster Name	The name of the cluster.
Broker Name	The name of the broker.
Broker ID	The broker ID for the server.
Alert Level	The current alert severity. <div> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. </div> <div> ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. </div> <div> ● Green indicates that no metrics have exceeded their alert thresholds. </div>
Alert Count	The total number of alerts for the host.
Broker State	The current state of the kafka broker.*
Active Controller	Denotes whether the broker is an active controller.*
CPU %	The percentage of CPU being used by the broker.*
Mem Used %	The percentage of JVM memory being used by the broker.*
Leader Count	The number of leaders on the broker.*
# Partitions	The number of partitions on the broker.*
# Offline Partitions	The number of partitions without an active leader on the broker.*
Under Replicated Partitions	The number of partition replicas that are out of sync (total number of replicas minus the total number of in-sync replicas) on the broker.*

Preferred Replica Imbalance Count	The number of topics whose replicas are not balanced on the broker.*
ZK Disconnect Rate	The mean rate of zookeeper disconnects per broker, in seconds.
# Purgatory Fetch	The number of fetch requests currently in purgatory (and waiting to be satisfied).*
# Purgatory Heartbeat	The number of requests in purgatory due to failed heartbeat tests.*
# Purgatory Produce	The number of produce requests currently in purgatory (and waiting to be satisfied).*
# Purgatory Rebalance	The number of changes that need to be propagated to the replicas so that the partitions are no longer in purgatory.*
# Purgatory Topic	The number of requests (based on topics) currently in purgatory.*
Network Processor Avg % Idle	The average fraction of time the network processors are idle.*
Version	The current version of Kafka.*
JMX Connection String	The JMX connection string used.*
Connected?	Denotes whether or not the broker is connected.*
Expired	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the Expire Time field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA STORAGE > Duration > Expire Time property. The RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA Storage > Duration > Delete Time property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if Expire Time was set to 120 and Delete Time was set to 3600, then the Expired check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
Timestamp	The date and time the row data was last updated.

All Broker Metrics Table

This display contains broker metrics broken down into three tabs: **Meters**, **Histograms**, and **Timers**. The table in the **Meters** tab provides a view of all of your metered metrics by broker and their associated metric data including count, event type, and rate data. The table in the **Histograms** tab provides a view of the histogram metrics for the selected broker. The table in the **Timers** tab provides a view of all the timers for the selected broker and their associated metrics. Each of the tabs also contains a trend graph, which provides a trend chart for each of the metrics listed in the associated table.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:**Cluster**

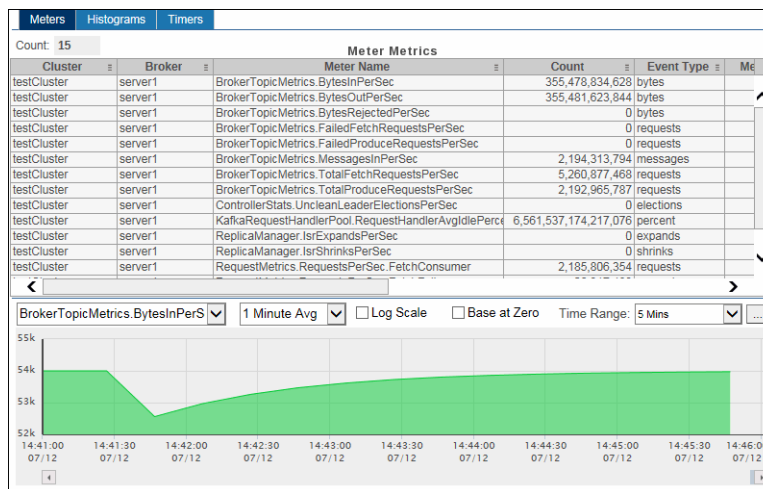
Select the cluster containing the broker for which you want to see data.

Broker

Select the broker for which you want to see data.

Count

The number of meters/histograms/timers found using the filter.

Meters Tab:**Cluster**

The name of the cluster.

Broker

The name of the broker.

Meter Name

The name of the metered metric.

Count

The total count for the metered metric.

Event Type

The type (unit) of metered metric.*

Mean Rate

The average rate for the meter, based on the **Rate Unit**.*

1 Minute Avg

The rate of incoming messages (per second) averaged over a one minute period for the meter, based on the **Rate Unit**.*

5 Minute Avg

The rate of incoming messages (per second) averaged over a five minute period for the meter, based on the **Rate Unit**.*

15 Minute Avg

The rate of incoming messages (per second) averaged over a fifteen minute period for the meter, based on the **Rate Unit**.*


Rate Unit

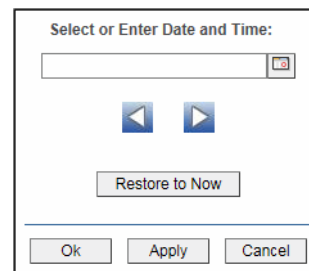
The unit of measure used to calculate **1 Minute Avg**, **5 Minute Avg**, **15 Minute Avg**, and **Mean Rate**.*


Expired

When checked, performance data in the row has not been received within the time specified (in seconds) in the **Expire Time** field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > **Solution Package Configuration** > **Apache Kafka** > **DATA STORAGE** > **Duration** > **Expire Time** property. The RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > **Solution Package Configuration** > **Apache Kafka** > **DATA STORAGE** > **Duration** > **Delete Time** property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

For example, if **Expire Time** was set to 120 and **Delete Time** was set to 3600, then the **Expired** check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.

Timestamp	The date and time the row data was last updated.
Metric Performance Trends	The trend chart provides a moving chart over the selected time range for each of the Meter Names listed in the Meters Metrics table.
<Meter Names Drop Down List>	Select the meter name for which you want to view data.
<Rate Drop Down List>	<p>Select the option for which you want to view data:</p> <p>Mean Rate -- Select this option to view the average rate for each metric for the metrics in the display.</p> <p>1 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display.</p> <p>5 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the metrics in the display.</p> <p>15 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display.</p>
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Select to use zero (0) as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar  .

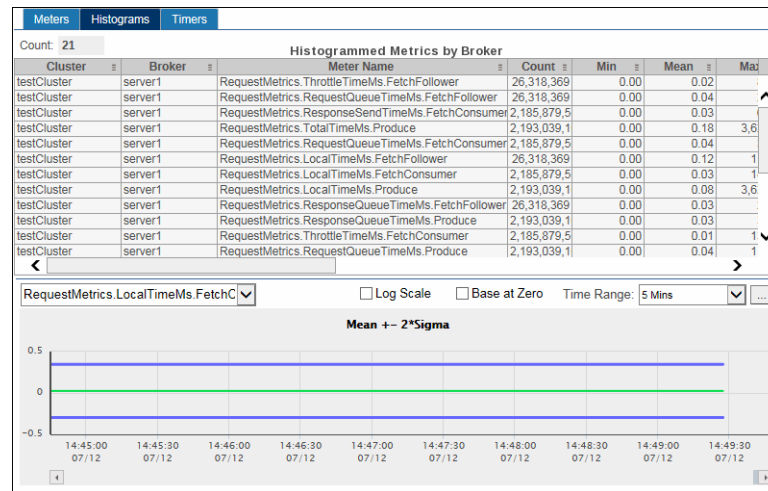


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Histograms Tab



Cluster	The name of the cluster.
Broker	The name of the broker.
Meter Name	The name of the metered metric.
Count	The total count for the metered metric.*
Min	The minimum number of occurrences for the meter during the current polling period.*
Mean	The average number of occurrences for the meter during the current polling period.*
Max	The maximum number of occurrences for the meter during the current polling period.*
Std Dev	The standard deviation for the number of occurrences for the meter during the current polling period.*
50th Percentile	The 50th percentile value for the number of occurrences for the meter during the current polling period.*
75th Percentile	The 75th percentile value for the number of occurrences for the meter during the current polling period.*
95th Percentile	The 95th percentile value for the number of occurrences for the meter during the current polling period.*
98th Percentile	The 98th percentile value for the number of occurrences for the meter during the current polling period.*
99th Percentile	The 99th percentile value for the number of occurrences for the meter during the current polling period.*
999th Percentile	The 999th percentile value for the number of occurrences for the meter during the current polling period.*

Expired

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

```
#####
# CACHE / HISTORIAN SETTINGS
#
# Cache history settings
#
sl.rtvadm.sub=$kafkaRowExpirationTime:120
sl.rtvadm.sub=$kafkaRowExpirationTimeForDelete:0
```

In the example above, the **Expired** check box would be checked after 120 seconds, and the row would never be deleted. If **\$kafkaRowExpirationTimeForDelete** was set to 3600, then the row would be removed from the table after 3600 seconds.

Timestamp

The date and time the row data was last updated.

Metric Performance Trends

The trend chart provides a moving chart over the selected time range for each of the **Metric Names** listed in the **Histogrammed Metrics by Broker** table.

<Metric Names Drop Down List>

Select the meter name for which you want to view data.

<Rate Drop Down List>

Select the option for which you want to view data:

Mean Rate -- Select this option to view the average rate for each metric for the metrics in the display.

1 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display.

5 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the metrics in the display.


15 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the metrics in the display.

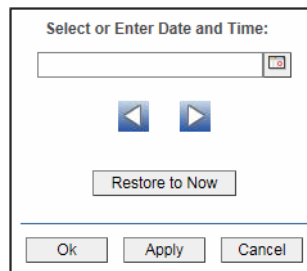
Log Scale


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

Base at Zero

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

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




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

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

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

Timers Tab:



Cluster	The name of the cluster.
Broker	The name of the broker.
Timer Name	The name of the timer.
Count	The total count for the timer.
Event Type	The event type for the timer.
Mean	The average number of events for the timer during the current polling period.*
Mean Rate	The average rate (based on Rate Unit) of events for the timer during the current polling period.*
Min	The minimum number of events for the timer during the current polling period.*
Max	The maximum number of events for the timer during the current polling period.*

Std Dev	The standard deviation for the number of events for the timer during the current polling period.*
1 Minute Avg	The rate of incoming messages (per second) averaged over a one minute period for the timer, based on the Rate Unit .*
5 Minute Avg	The rate of incoming messages (per second) averaged over a five minute period for the timer, based on the Rate Unit .*
15 Minute Avg	The rate of incoming messages (per second) averaged over a fifteen minute period for the timer, based on the Rate Unit .*
50th Percentile	The 50th percentile value for the number of events for the timer during the current polling period.*
75th Percentile	The 75th percentile value for the number of events for the timer during the current polling period.*
95th Percentile	The 95th percentile value for the number of events for the timer during the current polling period.*
98th Percentile	The 98th percentile value for the number of events for the timer during the current polling period.*
999th Percentile	The 999th percentile value for the number of events for the timer during the current polling period.*
99th Percentile	The 99th percentile value for the number of events for the timer during the current polling period.*
Latency Unit	The unit of measure used to calculate the latency.*
Rate Unit	The unit of measure used to calculate 1 Minute Rate , 5 Minute Rate , 15 Minute Rate , and Mean Rate .*
Expired	<p>When checked, performance data has not been received within the time specified (in seconds) in the \$kafkaRowExpirationTime field in the conf\rtvapm_kafkamonitor.properties file. The \$kafkaRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response. To view/edit the current values, modify the following lines in the .properties file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS # # Cache history settings # sl.rtvapm.sub=\$kafkaRowExpirationTime:120 sl.rtvapm.sub=\$kafkaRowExpirationTimeForDelete:0</pre> <p>In the example above, the Expired check box would be checked after 120 seconds, and the row would never be deleted. If \$kafkaRowExpirationTimeForDelete was set to 3600, then the row would be removed from the table after 3600 seconds.</p>
Timestamp	The date and time the row data was last updated.
Timers Performance Trends	<p>The trend chart provides a moving chart over the selected time range for each of the Timer Names listed in the Timer Metrics by Broker table.</p> <p><Timer Names Drop Down List> Select the timer name for which you want to view data.</p>

**<Rate Drop
Down List>**

Select the option for which you want to view data:

Mean Rate -- Select this option to view the average rate for each metric for the metrics in the display.

1 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display.

5 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the metrics in the display.

15 Minute Avg -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the metrics in the display.


Log Scale

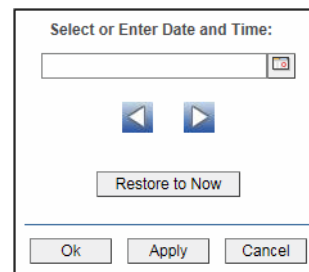
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log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


Base at Zero

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

Time Range

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



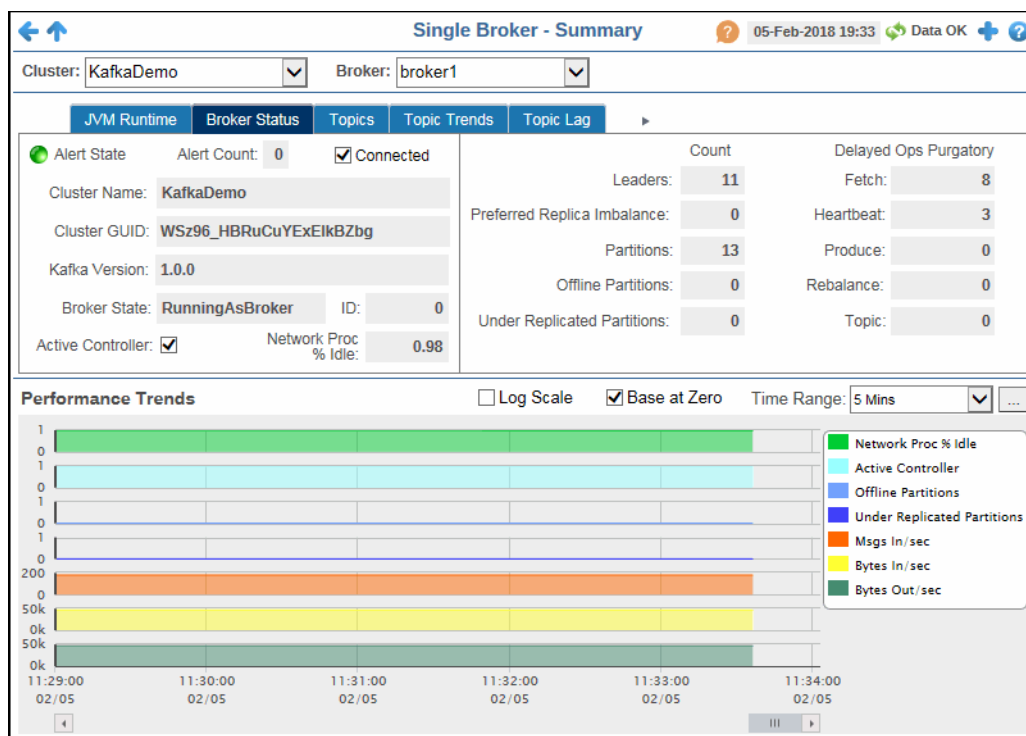
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Single Broker Summary

This display provides a view of the current and historical metrics for a single broker, including JVM runtime data, broker status, topic data, and topic trend data.



Title Bar (possible features are):

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

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

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

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

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

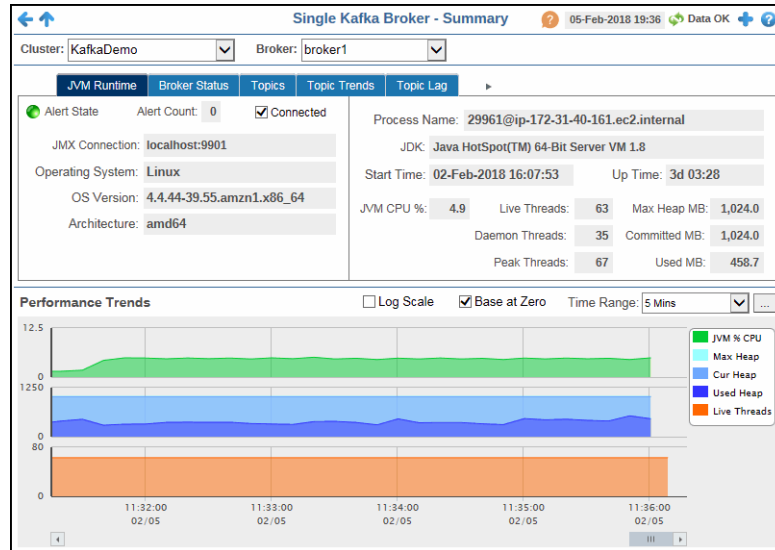
Cluster Select the cluster for which you want to show data in the display.

Broker

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

Topic

Only displays when the **Topic Trends** tab is selected. Select the topic for which you want to show data in the display.

JVM Runtime Tab**Alert State**

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

When checked, denotes that the broker is connected.

JMX Connection

The name of the JMX connection.*

Operating System

The operating system installed on the broker.*

OS Version

The version number of the operating system.*

Architecture

The type of processor being used.*

Process Name

The name of the process.*

JDK

The JDK version number.*

Start Time

The date and time when the broker was started.*

Up Time

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

JVM CPU %


The percentage of CPU used by the JVM.*

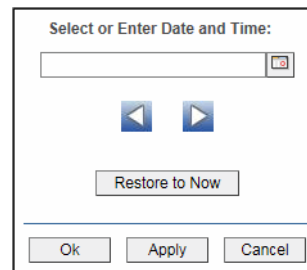
Live Threads


The number of live threads on the broker.*



Max Heap MB

The maximum amount of available heap, in megabytes.*

Daemon Threads	The number of daemon threads running on the broker.*
Committed MB	The total number of megabytes committed on the broker.*
Peak Threads	The highest number of threads running at one time during the current polling period.*
Used MB	The number of used megabytes on the broker.*
Performance Trends	<p>Traces the following:</p> <p>JVM % CPU -- traces the percentage of CPU used by the JVM.</p> <p>Max Heap -- traces the maximum amount of available heap.</p> <p>Cur Heap -- traces the current amount of heap being used.</p> <p>Used Heap -- traces the highest amount of heap used.</p> <p>Live Threads -- traces the number of live threads.</p> <p>Log Scale Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p> <p>Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.</p> <p>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 .</p>

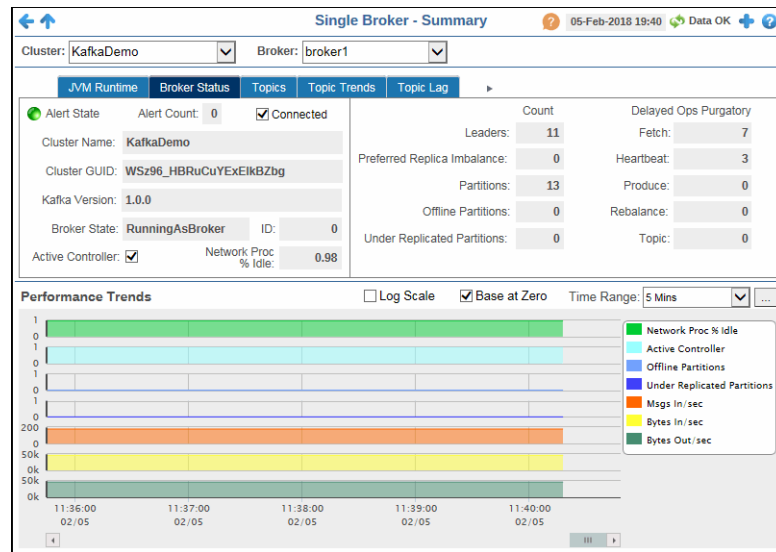


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Broker Status Tab



Alert State

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

When checked, denotes that the broker is connected.

Cluster Name

The name of the cluster in which the broker is contained.

Cluster GUID

Lists the cluster's globally unique identifier.

Note: This field will not be populated for brokers running on Kafka Version 0.9.*, and the **KafkaClusterSplitBrain** alert will not work properly for those brokers.

Kafka Version

The current version of Kafka installed on the broker.

Broker State

The current state of the Kafka broker.

ID

The broker ID for the server.

Active Controller

Denotes whether or not the broker is an active controller.

Network Proc % Idle

The average fraction of time the network processors are idle.*

Count

Leaders

The number of leaders on the broker.*

Preferred Replica Imbalance

The number of topics whose replicas are not balanced on the broker.*

Partitions


The number of partitions on the broker.*

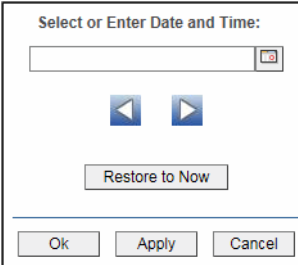
Offline Partitions

The number of partitions on the broker that are currently offline.*


Under Replicated Partitions	The number of partition replicas that are not in sync on the broker.*
Delayed Ops Purgatory	
Fetch	The number of fetch requests currently in purgatory (and waiting to be satisfied).*
Heartbeat	The number of requests in purgatory due to failed heartbeat tests.*
Produce	The number of produce requests currently in purgatory (and waiting to be satisfied).*
Rebalance	The frequency with which the partition rebalance check is triggered by the controller.*
Topic	The number of requests (based on topics) currently in purgatory.*
Performance Trends	<p>Traces the following:</p> <p>Network Proc % Idle -- traces the average fraction of time the network processors are idle.</p> <p>Active Controller -- traces whether or not the broker is/was an active controller.</p> <p>Offline Partitions -- traces the number of offline partitions.</p> <p>Under Replicated Partitions -- traces the number of partition replicas out of sync on the broker.</p> <p>Msgs In/sec -- traces the number of incoming messages per second.</p> <p>Bytes In/sec -- traces the number of incoming bytes per second.</p> <p>Bytes Out/sec -- traces the number of outgoing bytes per second.</p> <p>Log Scale Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p> <p>Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.</p>



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 .



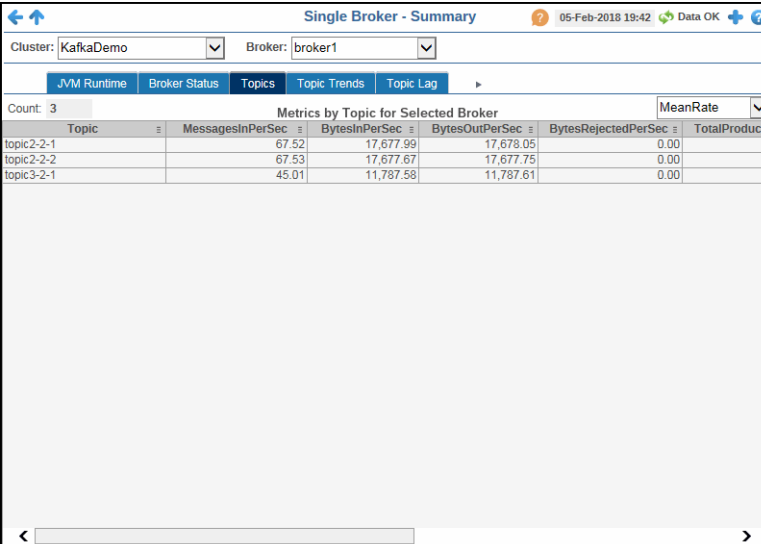
The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right) and a "Restore to Now" button. At the bottom are "Ok", "Apply", and "Cancel" buttons.

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

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

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

Topics Tab



The screenshot shows the "Single Broker - Summary" window. It includes a header with "Cluster: KafkaDemo" and "Broker: broker1". Below the header is a tabbed interface with "Topics" selected. The "Topics" tab displays a table with 3 topics. The table has columns for "Topic", "MessagesInPerSec", "BytesInPerSec", "BytesOutPerSec", "BytesRejectedPerSec", and "TotalProduce". The "MeanRate" dropdown is set to "MeanRate".

Topic	MessagesInPerSec	BytesInPerSec	BytesOutPerSec	BytesRejectedPerSec	TotalProduce
topic2-2-1	67.52	17,677.99	17,678.05	0.00	
topic2-2-2	67.53	17,677.67	17,677.75	0.00	
topic3-2-1	45.01	11,787.58	11,787.61	0.00	

Count

The total number of topics listed in the table.

<Rate Drop Down List>

Select the option for which you want to view data.

Mean Rate

Select this option to view the average rate for each metric for the topics in the display.

One Minute

Select this option to view the 1 minute rate for each metric for the topics in the display.

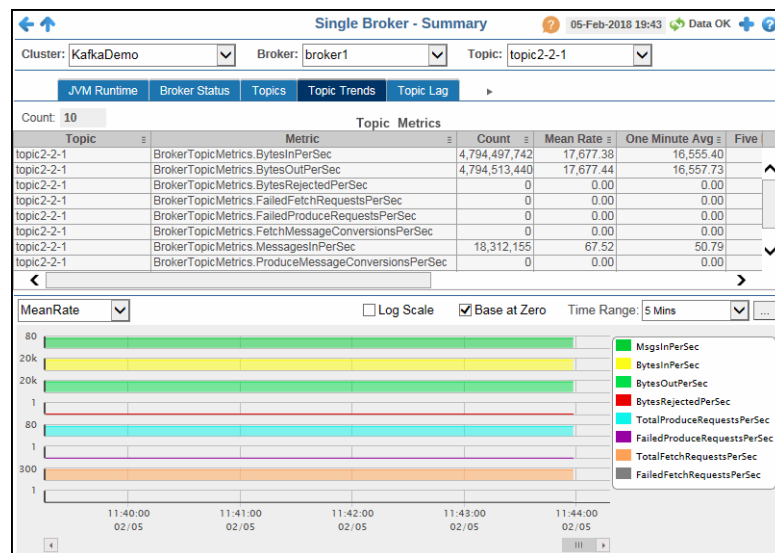
Five Minute

Select this option to view the 5 minute rate for each metric for the topics in the display.

Fifteen Minute Select this option to view the 15 minute rate for each metric for the topics in the display.

Topic	Lists the name of the topic.
Messages In Per Sec	The rate of incoming messages
Bytes In Per Sec	The rate of incoming bytes
Bytes Out Per Sec	The rate of outgoing bytes.
Bytes Rejected Per Sec	The rate of rejected bytes.
Total Produce Requests Per Sec	The rate of total produce requests.
Failed Produce Requests Per Sec	The rate of failed produce requests.
Total Fetch Requests Per Sec	The rate of total fetch requests.
Failed Fetch Requests Per Sec	The rate of failed fetch requests.


Topic Trends Tab

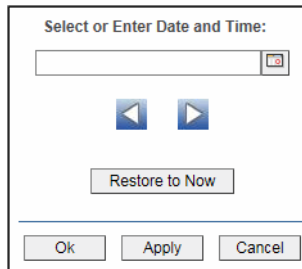


Count	The total number of topic metrics listed in the table.
Topic	The name of the topic
Metric	The name of the metric.
Count	The total number of the particular metric.
Mean Rate	The mean rate of the metric.


One Minute Avg	The (one minute) rate for the metric, based on the Rate Units.
Five Minute Avg	The (five minute) rate for the metric, based on the Rate Units.
Fifteen Minute Avg	The (fifteen minute) rate for the metric, based on the Rate Units.
Event Type	The event type for the topic metric.
Rate Units	The unit of measure used to calculate the Mean Rate , the One Minute Avg , the Five Minute Avg , and the Fifteen Minute Avg .
Expired	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the Expire Time field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA STORAGE > Duration > Expire Time property. The RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA Storage > Duration > Delete Time property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if Expire Time was set to 120 and Delete Time was set to 3600, then the Expired check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
time_stamp	The date and time the row data was last updated.
Rate Trends	<p>Select the option for which you want to view data:</p> <p>MeanRate -- Select this option to view the average rate for each metric in the trend graph.</p> <p>One Minute-- Select this option to view the rate averaged over a one minute period for each metric in the trend graph.</p> <p>Five Minute -- Select this option to view the rate averaged over a five minute period for each metric in the trend graph.</p> <p>Fifteen Minute -- Select this option to view the rate averaged over a fifteen minute period for each metric in the trend graph.</p> <p>Traces the following:</p> <p>MsgsInPerSec -- traces the selected rate of incoming messages.</p> <p>BytesInPerSec -- traces the selected rate of incoming bytes.</p> <p>BytesOutPerSec -- traces the selected rate of outgoing bytes.</p> <p>BytesRejectedPerSec -- traces the selected rate of rejected bytes.</p> <p>TotalProduceRequestsPerSec -- traces the selected rate of total produce requests.</p> <p>FailedProduceRequestsPerSec -- traces the selected rate of failed produce requests.</p> <p>TotalFetchRequestsPerSec -- traces the selected rate of total fetch requests.</p> <p>FailedFetchRequestsPerSec -- traces the selected rate of failed fetch requests.</p> <p>Log Scale Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p> <p>Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.</p>



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 .

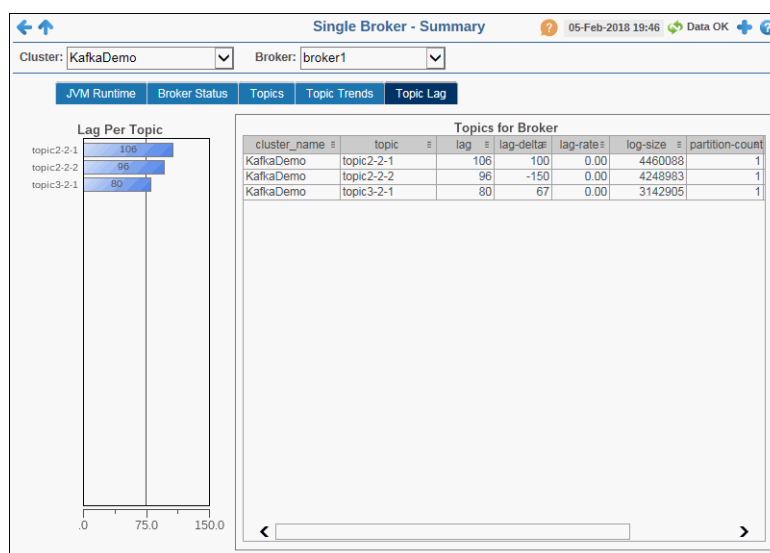


The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right). At the bottom of the dialog are three buttons: "Restore to Now", "Ok", "Apply", and "Cancel".

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   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.

Topic Lag Tab**Lag Per Topic Bar Graph**

Displays the lag per topic in a bar graph format.

Topics for Broker Table

cluster_name The name of the cluster in which the topic resides.

topic The name of the topic.

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

lag-delta The difference in the amount of lag from the previous polling period to the current polling period.*

lag-rate	The rate of change in the amount of lag.*
log-size	The current number of messages in the log.*
partition-count	The number of partitions containing the topic.
time_stamp	The date and time the row data was last updated.

Kafka Zookeepers View


These displays provide detailed data for all zookeepers or for a particular zookeeper. The available displays in this View are:

- [“All Zookeepers Heatmap”](#): Heatmap view of all zookeepers and their associated metrics in a particular cluster.
- [“All Zookeepers Table”](#): Tabular view of all zookeepers and their associated metrics in a particular cluster.
- [“Zookeepers Summary”](#): Contains current and historical metrics, as well as trend data, for a single zookeeper.
-

All Zookeepers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your zookeepers for each available metric. You can view the zookeepers in the heatmap based on the following metrics: the current alert severity, the current alert count, the number of clients connections, the number of queued requests, the number of incoming packets per second, and the number of outgoing packets per second. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a zookeeper. Clicking one of the rectangles in the heatmap opens the [“Zookeepers Summary”](#) display, which allows you to see additional details for the selected zookeeper.



Title Bar (possible features are):

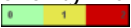
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.




- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

Filter By:


- Cluster** Select the cluster for which you want to view data.
- Names** Select this check box to display the names of the zookeepers at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

Alert Severity

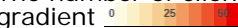
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

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

Alert Count

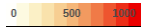
The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient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

Alive Connections

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

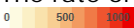
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Outstanding Reqs

The number of queued requests. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaZookeeperOutstandingReqs**. The middle value in the gradient bar indicates the middle value of the range.

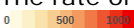
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Packets In Per Sec

The rate of incoming packets (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaZookeeperRatePktsRcvd**. The middle value in the gradient bar indicates the middle value of the range.

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

Packets Out Per Sec

The rate of outgoing packets (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaZookeeperRatePktsSent**. The middle value in the gradient bar indicates the middle value of the range.

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

All Zookeepers Table

The table in this display provides a view of all of the zookeepers for a specific cluster and their associated metric data including connection, cluster name, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the [“Zookeepers Summary”](#) display.

Cluster Name	Zookeeper Name	Alert Level	Alert Count	Role	Max Request Latency (ms)	Avg Request Latency (ms)
testCluster	zoo1		0	Follower	13	
testCluster	zoo2		0	Follower	18	
testCluster	zoo3		0	Leader	70	

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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

Cluster	Select the cluster for which you want to see data.
Count	The number of zookeepers that were found in the selected cluster.

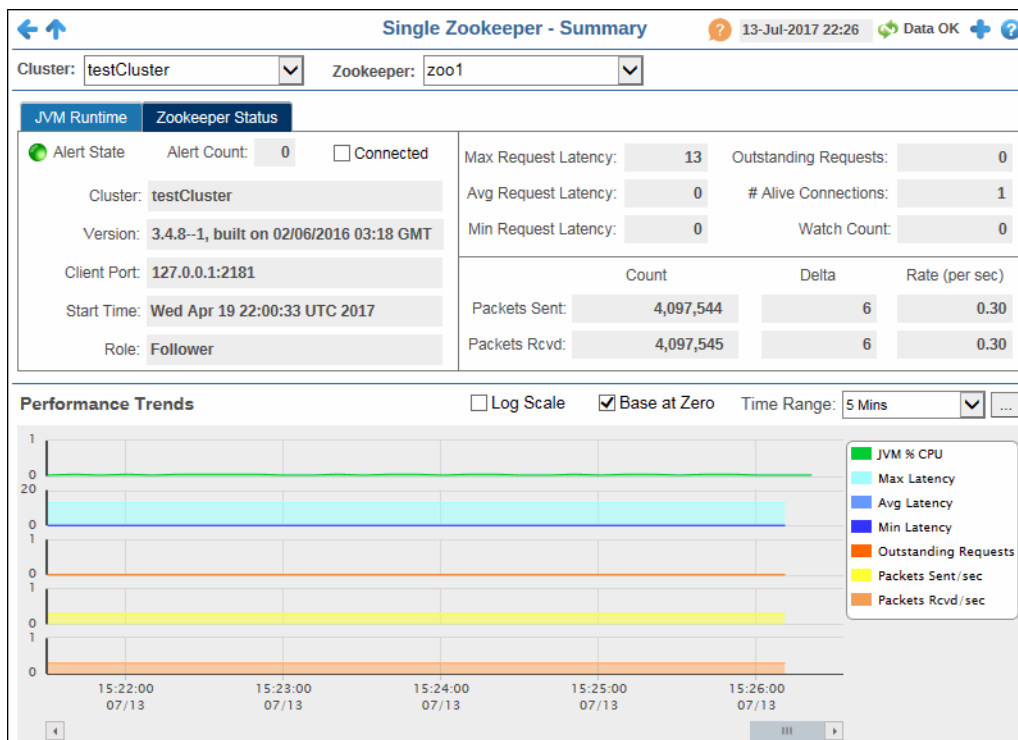
All Zookeepers Table:

Cluster Name	The name of the cluster.
Zookeeper Name	The name of the zookeeper.
Alert Level	<p>The current alert severity.</p> <p> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</p> <p> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</p> <p> Green indicates that no metrics have exceeded their alert thresholds.</p>
Alert Count	The total number of alerts for the host.
Role	The role of the zookeeper (Leader or Follower).*
Max Request Latency (ms)	The longest amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Avg Request Latency	The average amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Min Request Latency	The least amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Num Alive Connections	The number of clients connected to the zookeeper.*
Outstanding Requests	The number of queued requests.*
Node Count	The total number of nodes.*
Watch Count	The number of watchers set up over the zookeeper nodes.*
Packets Recvd	The number of packets received.*
Packets Sent	The number of packets sent.*
Delta Packets Recvd	The increase in the amount of packets received by the zookeeper (from the previous polling period to the current polling period).*
Delta Packets Sent	The increase in the amount of packets sent from the zookeeper (from the previous polling period to the current polling period).*
Rate Packets Recvd	The rate at which packets are being received by the zookeeper.*
Rate Packets Sent	The rate at which packets are being sent by the zookeeper.*
Max Client Cnxns Per Host	The maximum number of connections allowed from each host.*
Max Session Timeout	The maximum allowed session timeout allowed for registered consumers.*

Min Session Timeout	The minimum allowed session timeout allowed for registered consumers.*
Version	The current version of Kafka being used.*
Client Port	Lists the client's port.*
JMX Connection String	Lists the connection string.*
Connected	Denotes whether or not the zookeeper is connected.
Expired	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the Expire Time field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA STORAGE > Duration > Expire Time property. The RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA Storage > Duration > Delete Time property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if Expire Time was set to 120 and Delete Time was set to 3600, then the Expired check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
Start Time	The date and time the zookeeper was started.*
Timestamp	The date and time the row data was last updated.

Zookeepers Summary

This display provides a view of the current and historical metrics for a single zookeeper. You can view JVM runtime statistics and trend data as well as zookeeper status and trend data for the selected zookeeper.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

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

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

Open the Alert Views - RTView Alerts Table display.

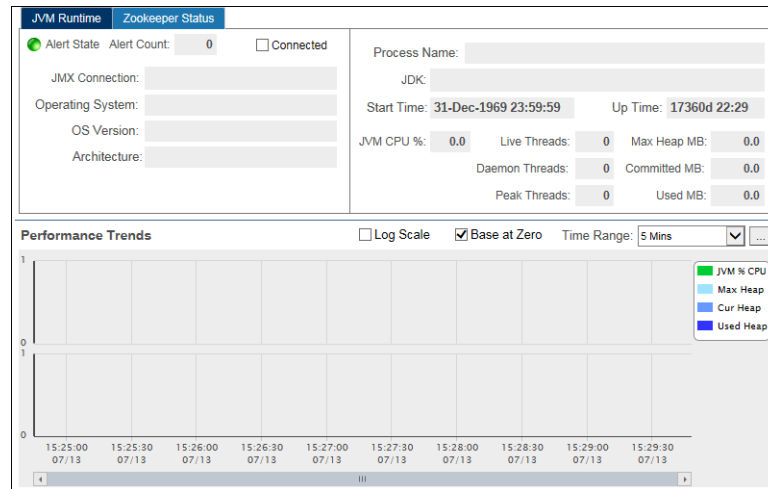
Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

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

Zookeeper

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

JVM Runtime Tab:**Alert State**

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

Denotes whether or not the jmx connection is connected.

JMX Connection

The name of the JMX connection.*

Operating System

The operating system installed on the zookeeper.*

OS Version

The version number of the operating systems.*

Architecture

The type of processor being used.*

Process Name

The name of the process.*

JDK

The JDK version number.*

Start Time

The date and time when the zookeeper was started.*

Up Time

The amount of time the zookeeper has been up and running.*

JVM CPU %

The percentage of CPU being used by the JVM.*

Live Threads

The number of live threads.*

Max Heap MB

The maximum amount of available heap, in megabytes.*

Daemon Threads

The number of daemon threads running.*

Committed MB

The total number of megabytes committed.*

Peak Threads The highest number of threads running at one time during the current polling period.*

Used MB The number of used megabytes.*

Performance Trends Traces the following:

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

Max Heap -- traces the maximum amount of available heap.

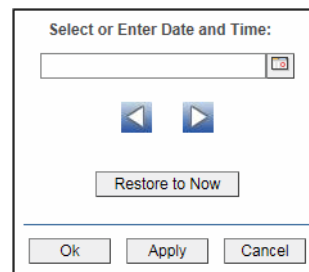
Cur Heap-- traces the current amount of heap being used.


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



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

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

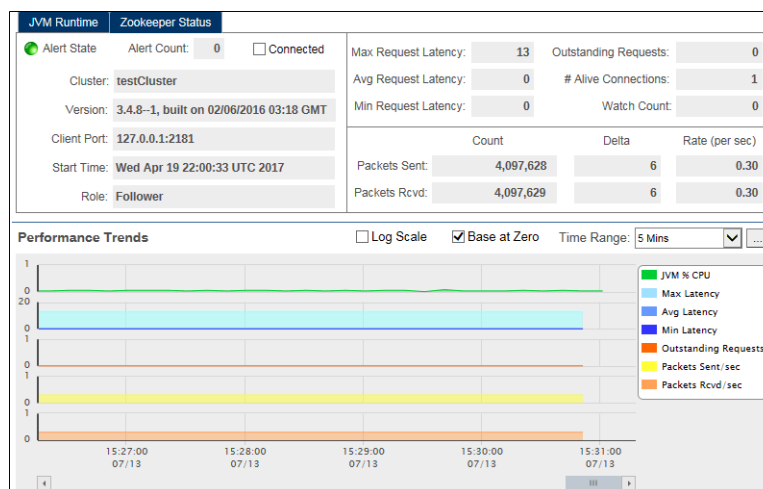
Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



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

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

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

Zookeeper Status Tab:**Alert State**

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

When checked, denotes that the zookeeper is connected.

Cluster

The name of the cluster in which the zookeeper is contained.

Version

The current version of Apache Kafka installed.*

Client Port

The client's IP address and port.*

Start Time

The date and time when the zookeeper was started.*

Role

The zookeeper's role (**Leader/Follower**).*

Max Request Latency

The longest amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*

Avg Request Latency

The average amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*

Min Request Latency

The least amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*

Outstanding Requests

The number of queued requests.*

Alive Connections

The number of clients connected to the zookeeper.*

Watch Count


The number of watchers set up over the zookeeper nodes.*

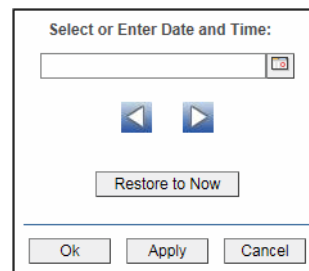
Packets Sent


Count -- The number of packets sent.*



Delta -- The increase in the amount of packets sent from the zookeeper (from the previous polling period to the current polling period).*

Rate (per sec) -- The rate at which packets are being sent (per second) by the zookeeper.*

- Packets Rcvd Count** -- The number of packets received.*
- Delta** -- The increase in the amount of packets received by the zookeeper (from the previous polling period to the current polling period).*
- Rate (per sec)** -- The rate at which packets are being received (per second) by the zookeeper.*
- Performance Trends** Traces the following:
- JVM % CPU** -- traces the percentage of CPU used by the JVM.
 - Max Latency** -- traces the longest amount of time taken to respond to a client request.
 - Avg Latency** -- traces the average amount of time taken to respond to a client request.
 - Min Latency** -- traces the least amount of time taken to respond to a client request.
 - Outstanding Requests** -- traces the number of queued requests.
 - Packets Sent/sec** -- traces the rate at which packets are being sent (per second) by the zookeeper.
 - Packets Rcvd/sec** -- traces the rate at which packets are being received (per second) by the zookeeper.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



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

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

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

Kafka Producers View

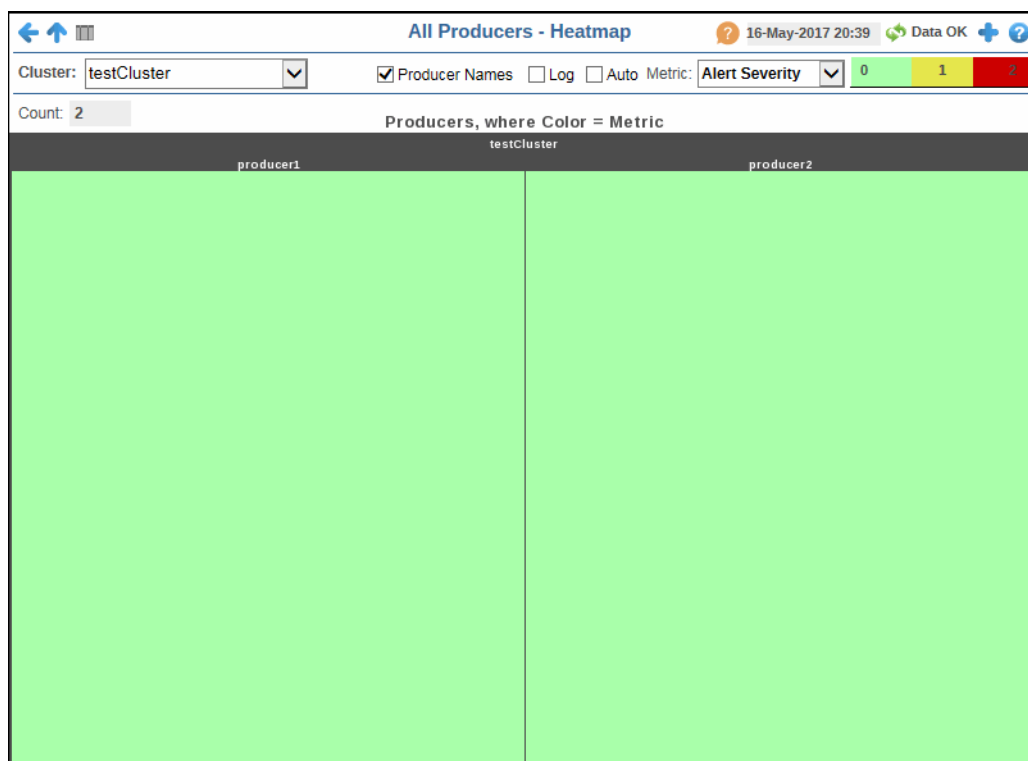
These displays provide detailed data for all producers or for a particular producer. The available displays in this View are:

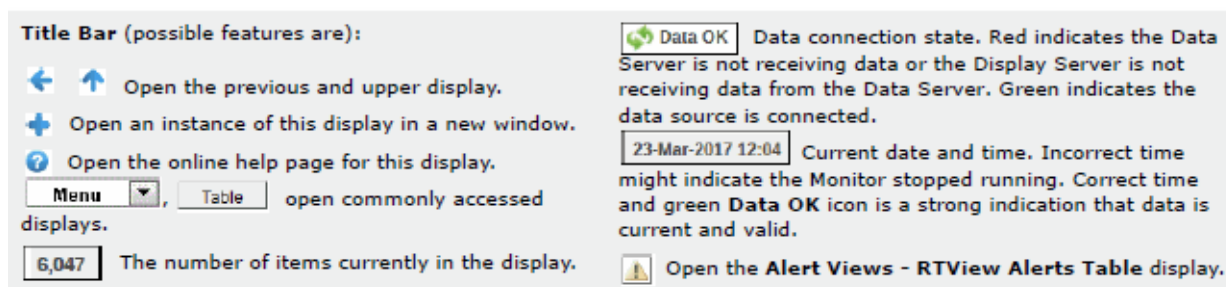
- **"All Producers Heatmap"**: Heatmap view of all producers and their associated metrics in a particular cluster.
- **"All Producers Table"**: Tabular view of all producers and their associated metrics in a particular cluster.
- **"Producer Summary"**: Contains current and historical metrics, as well as trend data, for a single producer.

All Producers Heatmap

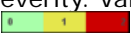


This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your producers for each available metric. You can view the producers in the heatmap based on the following metrics: the current alert severity, the current alert count, the incoming/outgoing byte rate, the IO wait time, the request latency, and the request/response rates. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Producer Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a producer. Clicking one of the rectangles in the heatmap opens the **"Producer Summary"** display, which allows you to see additional details for the selected producer.

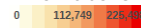




Fields and Data:

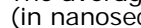
Cluster	Select the cluster for which you want to view data.
Producer Names	Select this check box to display the names of the producers at the top of each rectangle in the heatmap.
Log	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. Note: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Incoming Byte Rate	<p>The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of KafkaProducerIncomingByteRate. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

**Outgoing
Byte Rate**

The rate of outgoing bytes (per second). The color gradient bar  shows the range of the value/color mapping. The current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerOutgoingByteRate**. The middle value in the gradient bar indicates the middle value of the range.


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

**IO Wait Time
NSec Avg**

The average length of time the IO thread spent waiting for a socket (in nanoseconds). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerIoWaitTimeMS**. The middle value in the gradient bar indicates the middle value of the range.

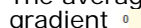
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Request
Latency**

The amount of time between when a producer is called and when the producer receives a response from the broker. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerRequestLatency**. The middle value in the gradient bar indicates the middle value of the range.

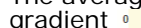
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Request Rate

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

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

**Response
Rate**

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

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

All Producers Table

The table in this display provides a view of all of your producers and their associated metric data including connection, alert level, alert count, cluster name, client ID, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the “Producer Summary” display.

All Producers - Table26-Nov-2018 14:18Data OK

Cluster:KafkaMirror

Count: 1

All Producers Table

Cluster Name	Producer Name	Alert Level	Alert Count	Client ID	CPU %	Mem Used %	Batch Size Avg
KafkaMirror	mirror-maker			producer-1	0.19	16.96	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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

- Cluster


Select the cluster for which you want to view data.
- Count

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

All Producers Table:

- Cluster Name

The name of the cluster.*

Producer Name	The name of the producer.
Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the host.
Client ID	The ID of the producer.*
CPU %	The percentage of CPU being used by the producer.*
Mem Used %	The percentage of JVM memory being used by the producer.*
Batch Size Average	The average batch size sent by the producer.*
Batch Size Max	The maximum number of messages that can be added to a batch before being sent to the event handler.*
Buffer Available Bytes	The number of available bytes in the buffer.*
Buffer Exhausted Rate	The average per-second number of record sends that are dropped due to buffer exhaustion.*
Buffer Total Bytes	The total number of bytes allowed in the buffer.*
Buffer Pool Wait Ratio	The fraction of time an appender waits for space allocation.*
Compression Rate Avg	The average compression rate of record batches.*
Connection Close Rate	The rate of connections being closed.*
Connection Count	The number of active connections.*
Connection Creation Rate	The rate of connections being created.*
Incoming Byte Rate	The average number of incoming bytes per second.*
IO Ratio	The rate of input/output operations.*
IO Time NS Avg	The average length of time the I/O thread spent waiting for a socket (in nanoseconds).*
IO Wait Ratio	The percent of time the producer was performing I/O operations while the CPU was idle.*
IO Wait Time Millisec Avg	The average length of time the I/O thread spent waiting for a socket (in milliseconds).*
Metadata Age	The age (in seconds) of the current producer metadata being used.*
Network IO Rate	The rate of input/output network operations.*

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

Expired

When checked, performance data in the row has not been received within the time specified (in seconds) in the **Expire Time** field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > **Solution Package Configuration** > **Apache Kafka** > **DATA STORAGE** > **Duration** > **Expire Time** property. The RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > **Solution Package Configuration** > **Apache Kafka** > **DATA Storage** > **Duration** > **Delete Time** property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

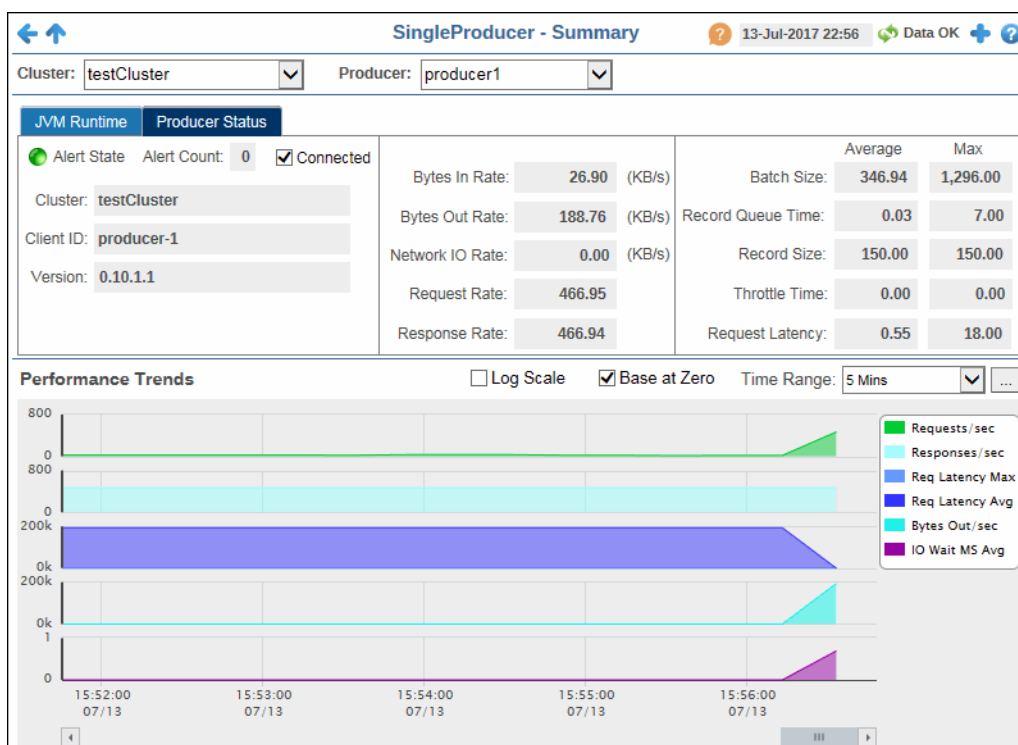
For example, if **Expire Time** was set to 120 and **Delete Time** was set to 3600, then the **Expired** check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.

Timestamp

The date and time the row data was last updated.

Producer Summary

This display provides a view of the current and historical metrics for a single producer. You can view JVM runtime statistics and trend data as well as producer status and trend data for the selected producer.



Title Bar (possible features are):

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

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

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

🔔 Open the Alert Views - RTView Alerts Table display.


Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

Producer

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

JVM Runtime Tab:



Alert State

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

Denotes whether or not the jmx connection is connected.

JMX Connection


The name of the JMX connection.*

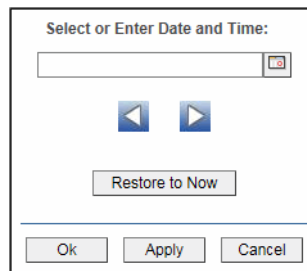
Operating System


The operating system installed on the producer.*



OS Version	The version number of the operating systems.*
Architecture	The type of processor being used.*
Process Name	The name of the process.*
JDK	The JDK version number.*
Start Time	The date and time when the producer was started.*
Up Time	The amount of time the producer has been up and running.*
JVM CPU %	The percentage of CPU being used by the JVM.*
Live Threads	The number of live threads.*
Max Heap MB	The maximum amount of available heap, in megabytes.*
Daemon Threads	The number of daemon threads running.*
Committed MB	The total number of megabytes committed.*
Peak Threads	The highest number of threads running at one time during the current polling period.*
Used MB	The number of used megabytes.*
Performance Trends	<p>Traces the following:</p> <ul style="list-style-type: none"> JVM % CPU -- traces the CPU being used by the JVM. Max Heap -- traces the maximum amount of available heap. Cur Heap-- traces the current amount of heap being used. Used Heap-- traces the highest amount of heap used. <p>Log Scale Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p> <p>Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.</p>

Time Range

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

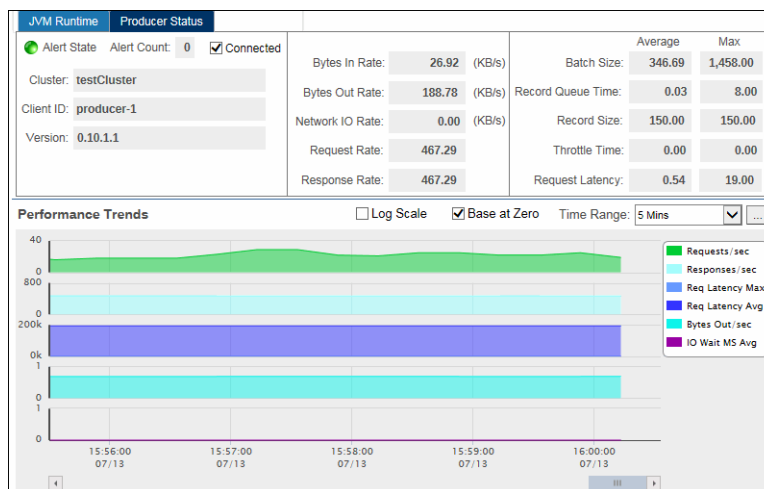


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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




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

Producer Stats Tab:



Alert State

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

When checked, denotes that the producer is connected.

Cluster

The name of the cluster in which the producer is contained.

Client ID

The ID of the client.

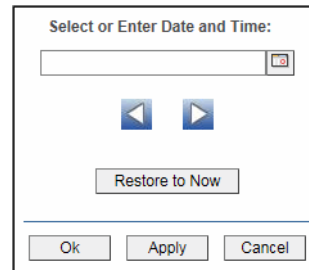
Version

The current version of Apache Kafka installed.*


Bytes In Rate	The rate of incoming bytes (kilobytes per second).*
Bytes Out Rate	The rate of outgoing bytes (kilobytes per second).*
Network IO Rate	The rate of input/output network operations.*
Request Rate	The average number of requests sent per second.*
Response Rate	The average number of responses received per second.*
Batch Size	Average -- The average batch size sent by the producer.* Max -- The maximum number of messages that can be added to a batch before being sent to the event handler.*
Record Queue Time	Average -- The average time (in milliseconds) record batches spent in the record accumulator.* Max -- The maximum time (in milliseconds) record batches spent in the record accumulator.*
Record Size	Average -- The average record size.* Max -- The maximum record size.*
Throttle Time	Average -- The average throttle time (in milliseconds).* Max -- The maximum time (in milliseconds) a request was throttled by a broker.*
Request Latency	Average -- The average request latency (in milliseconds).* Max -- The maximum request latency (in milliseconds).*
Performance Trends	Traces the following: Requests/sec -- traces the number of requests per second. Responses/sec -- traces the number of responses per second. Req Latency Max -- traces the maximum request latency (in milliseconds). Req Latency Avg -- traces the average request latency (in milliseconds). Bytes Out/sec -- traces the rate of outgoing bytes (kilobytes per second). IO Wait MS Avg -- traces the average length of time the I/O thread spent waiting for a socket (in milliseconds). Log Scale Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.



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

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



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath these arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

Kafka Consumers View

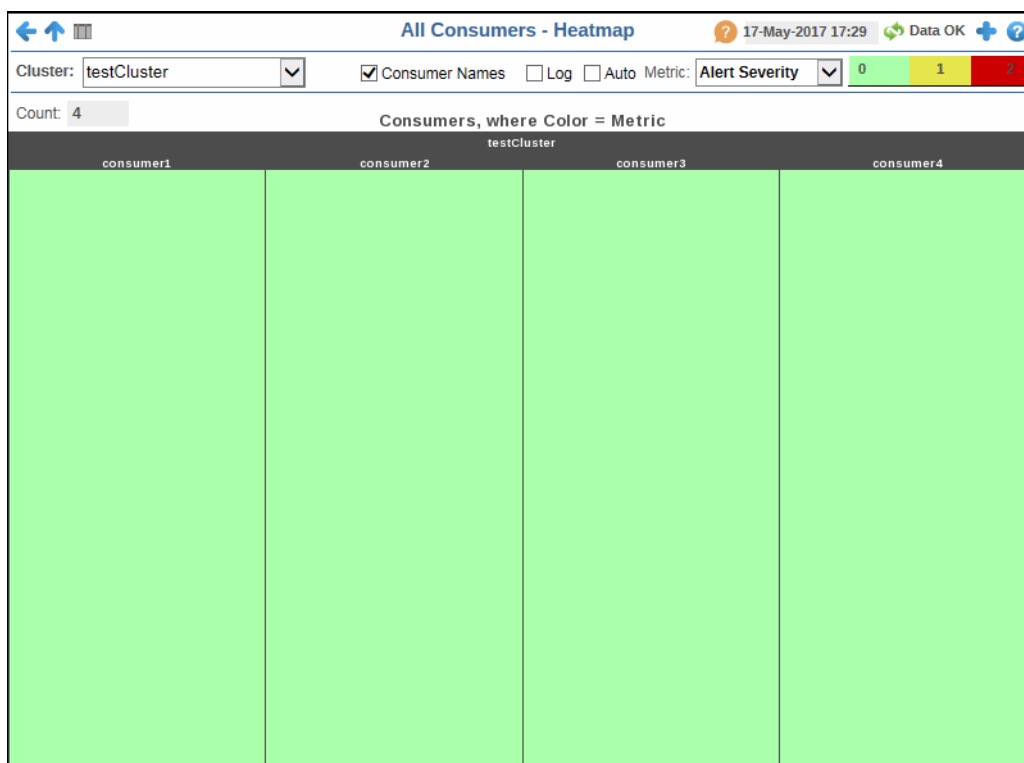
These displays provide detailed data for all consumers or for a particular consumer. The available displays in this View are:

- **"All Consumers Heatmap"**: Heatmap view of all consumers and their associated metrics in a particular cluster.
- **"All Consumers Table"**: Tabular view of all consumers and their associated metrics in a particular cluster.
- **"Consumers Summary"**: Contains current and historical metrics, as well as trend data, for a single consumer.

All Consumers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your consumers for each available metric. You can view the consumers in the heatmap based on the following metrics: the current alert severity, the current alert count, the bytes consumed rate, the fetch latency average, the fetch rate, the maximum consumer lag, and the records consumed rate. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Consumer Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a consumer. Clicking one of the rectangles in the heatmap opens the **"Consumers Summary"** display, which allows you to see additional details for the selected consumer.



Title Bar (possible features are):

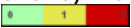
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.




- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data:


- Cluster** Select the cluster for which you want to view data.
- Consumer Names** Select this check box to display the names of the consumers at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

Alert Severity

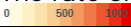
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

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

Alert Count

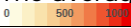
The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient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

Bytes Consumed Rate

The rate of bytes being consumed (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.


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

Fetch Latency Avg

The average time taken for fetch request. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.


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

Fetch Rate

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

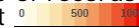
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Consumer
Max Lag**

The maximum lag in the number of records for any partition. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.

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

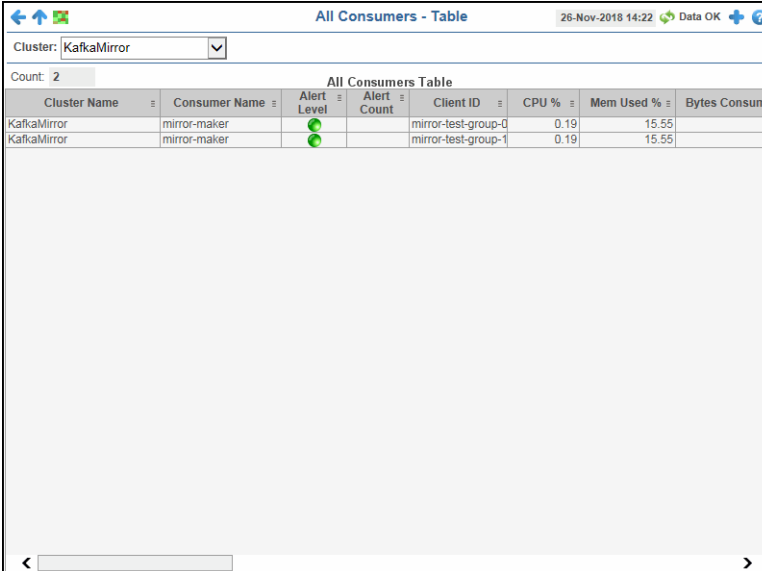
**Records
Consumed
Rate**



The rate of records being consumed (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.


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

All Consumers Table

The table in this display provides a view of all of your consumers and their associated metric data including connection, alert level, alert count, cluster name, client ID, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected consumer in the “[Consumers Summary](#)” display



Cluster Name	Consumer Name	Alert Level	Alert Count	Client ID	CPU %	Mem Used %	Bytes Consumed
KafkaMirror	mirror-maker			mirror-test-group-0	0.19	15.55	
KafkaMirror	mirror-maker			mirror-test-group-1	0.19	15.55	



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

Cluster	Select the cluster for which you want to view data.
Count	The number of consumers found on the selected cluster, which are listed in the All Consumers Table .

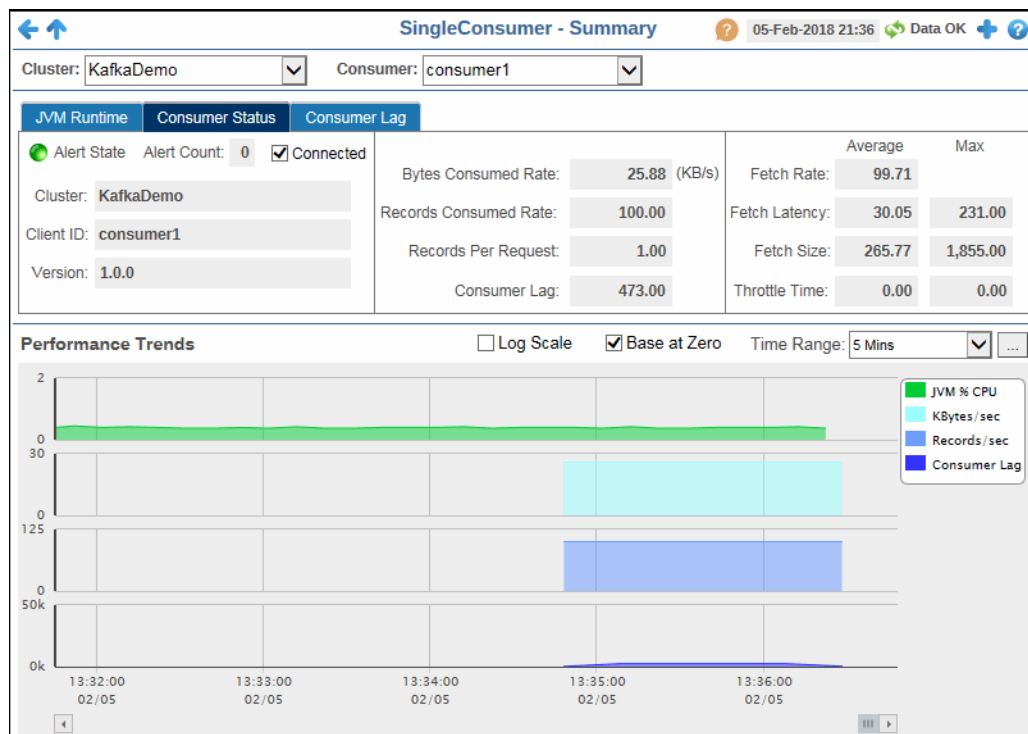
All Consumers Table:

Cluster Name	The name of the cluster.
Consumer Name	The name of the consumer.
Alert Level	The current alert severity. ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. ● Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the host.
Client ID	The ID of the consumer.*
CPU %	The percentage of CPU being used by the consumer.*
Mem Used %	The percentage of JVM memory being used by the consumer.*
Bytes Consumed Rate	The average number of bytes consumed per second.*
Fetch Latency Avg	The average time taken for a fetch request.*
Fetch Latency Max	The maximum time taken for a fetch request.*
Fetch Rate	The number of fetch requests per second.*
Fetch Size Avg	The average number of bytes fetched per request.*
Fetch Size Max	The maximum number of bytes fetched per request.*

Fetch Throttle Time Avg	The average throttle time in milliseconds.*
Fetch Throttle Time Max	The maximum throttle time in milliseconds.*
Records Consumed Rate	The average number of records consumed per second.*
Records Lag Max	The maximum lag in the number of records for any partition.*
Records per Request Avg	The average number of records in each request.*
JMX Connection String	The JMX connection string.*
Version	The current version of Apache Kafka installed.*
Connected	Denotes whether or not the consumer is connected.*
Expired	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the Expire Time field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA STORAGE > Duration > Expire Time property. The RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA Storage > Duration > Delete Time property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if Expire Time was set to 120 and Delete Time was set to 3600, then the Expired check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
Timestamp	The date and time the row data was last updated.

Consumers Summary

This display provides a view of the current and historical metrics for a single consumer. You can view JVM runtime statistics and trend data as well as consumer statistics and trend data for the selected consumer.



Title Bar (possible features are):

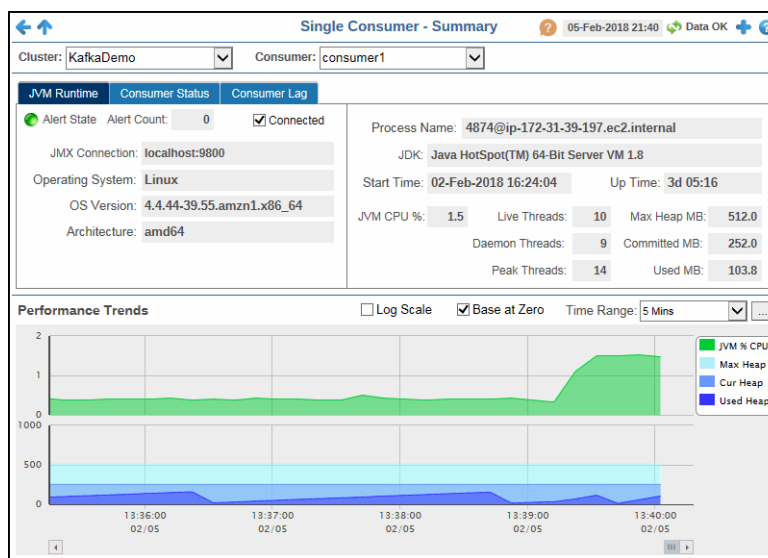
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

Filter By:

- Cluster** Select the cluster for which you want to show data in the display.
- Consumer** Select the consumer for which you want to show data in the display.

JVM Runtime Tab:**Alert State**

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

Denotes whether or not the JMX connection is connected.

JMX Connection

The name of the JMX connection.*

Operating System

The operating system installed on the producer.*

OS Version

The version number of the operating systems.*

Architecture

The type of processor being used.*

Process Name

The name of the process.*

JDK

The JDK version number.*

Start Time

The date and time when the producer was started.*

Up Time

The amount of time the producer has been up and running.*

JVM CPU %

The percentage of CPU used by the JVM.*

Live Threads

The number of live threads.*

Max Heap MB

The maximum amount of available heap, in megabytes.*

Daemon Threads

The number of daemon threads running.*

Committed MB

The total number of megabytes committed.*

Peak Threads The highest number of threads running at one time during the current polling period.*

Used MB The number of used megabytes.*

Performance Trends Traces the following:

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

Max Heap -- traces the maximum amount of available heap.

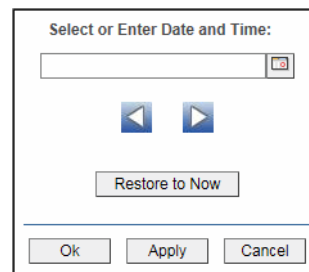
Cur Heap-- traces the current amount of heap being used.


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



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

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

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

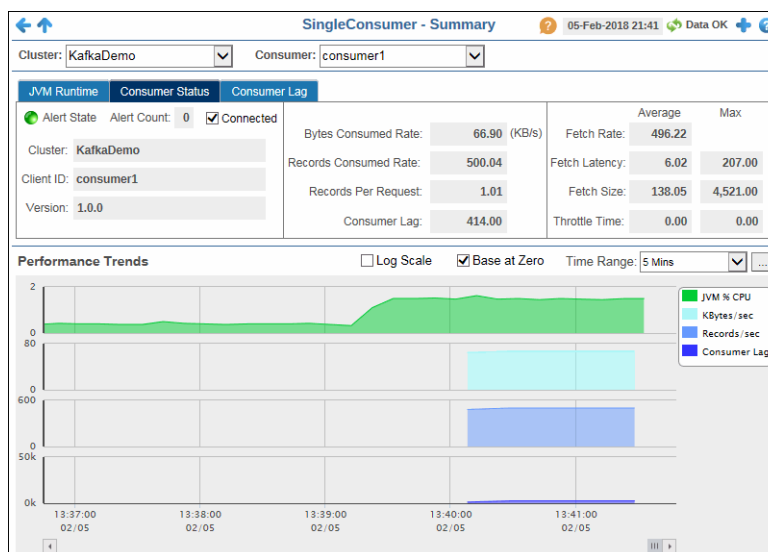


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Consumer Status Tab:



Alert State

The current alert severity.

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

Alert Count

The total number of current alerts.

Connected

When checked, denotes that the consumer is connected.

Cluster

The name of the cluster in which the consumer is contained.

Client ID

The ID of the client.

Version

The current version of Apache Kafka installed.*

Bytes Consumed Rate

The average number of bytes consumed per second.*

Records Consumed Rate

The average number of records consumed per second.*

Records Per Request

The average number of records in each request.*

Consumer Lag

The maximum lag in number of records for any partition.*

Fetch Rate

Average -- The average number of fetch requests per second.*

Max -- The highest number of fetch requests per second.*

Fetch Latency

Average -- The average time taken for a fetch request.*

Max -- The maximum amount of time taken for a fetch request.*

Fetch Size

Average -- The average number of bytes fetched per request.*

Max -- The highest number of bytes fetched per request.*

Throttle Time **Average** -- The average throttle time, in milliseconds.*
Max -- The maximum throttle time, in milliseconds.*

Performance Trends Traces the following:

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

KBytes/sec -- traces the number of kilobytes consumed per second.

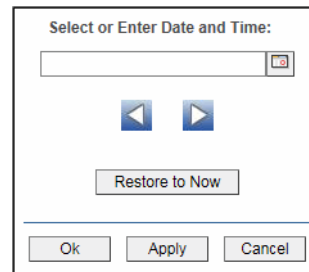
Records/sec -- traces the number of records being fetched per second.


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



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

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

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

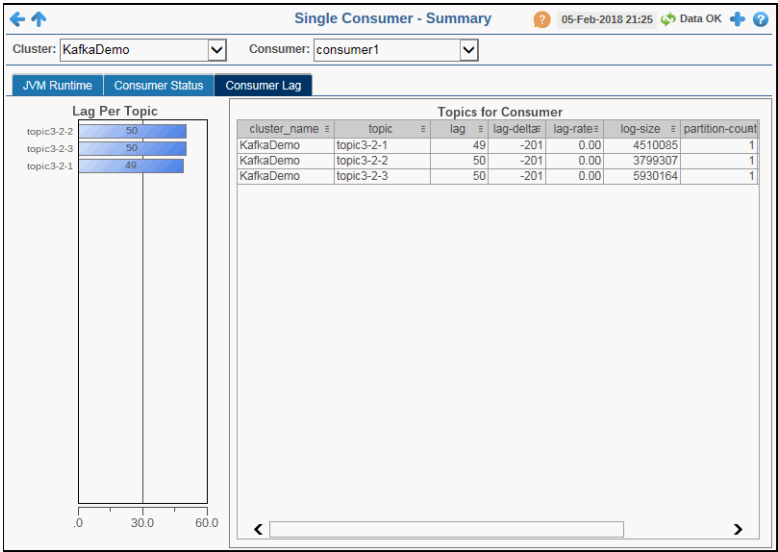


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Consumer Lag
Tab



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

Topics for Consumer Table

cluster_name	The name of the cluster in which the topic resides.
topic	The name of the topic.
lag	The difference between the current consumer position in the partition and the end of the log.*
lag-delta	The difference in the amount of lag from the previous polling period to the current polling period.*
lag-rate	The rate of change in the amount of lag.*
log-size	The current number of messages in the log.*
partition-count	The number of partitions containing the topic.
time_stamp	The date and time the row data was last updated.

■

CHAPTER 6 RTView DataServer for Oracle

The RTView DataServer for Oracle provides a way to create connections and modify default configuration settings for the various Oracle solution packages and sends collected data to RTView Central, which contains the displays associated with the RTView DataServer for Oracle that help you to monitor the health and performance across your Oracle components.

RTView Central contains the following Views and their associated displays that will be populated with data collected via the RTView DataServer for Oracle:

- "Oracle Coherence"
- "Oracle Database"
- "Connector for Oracle Enterprise Manager"
- "Oracle WebLogic"

The RTView *DataCollector* for Oracle is also available for use with the RTView DataServer for Oracle. RTView DataCollector for Oracle is used for collecting solution package data and sending it to one or more RTView DataServers. The RTView DataCollector for Oracle is useful if you need to distribute data collection.

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

Oracle Coherence

The Solution Package for Oracle Coherence provides information about the health and configuration of your Oracle Coherence cluster elements, including caches, nodes, services and clients.



The Solution Package for Oracle Coherence collects metrics from all your Coherence elements simultaneously, and does so at frequent intervals (typically every 10 seconds). At each interval, the OC 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.

The Solution Package for Oracle Coherence can be configured for a single Coherence cluster or multiple Coherence clusters. The Solution Package for Oracle Coherence is also often used in pre-production environments for conducting load testing and performance tuning.

The following Oracle Coherence Views can be found under **Components** tab > **Middleware** > **Oracle Coherence**:

- **"Cluster Selector"**: See all your Coherence clusters and Data Servers and choose which cluster to display data for.
- **"Cluster Views"**: Use these displays to assess Coherence cluster-level performance and utilization.
- **"Proxy Services"**: Use these displays to assess proxy service performance metrics.
- **"Cache Services"**: Use these displays to assess performance and utilization of all caches in the cluster.
- **"Federated Clusters"**: See all your Federated clusters and Data Servers and choose which cluster to display data for.

- **"All Caches"**: Use these displays to investigate performance, utilization and activity metrics of a single cache.
- **"Single Cache"**: Use these displays to assess node-level performance and utilization in the cluster.
- **"All Nodes"**: Use these displays to investigate performance and utilization metrics of a single node.
- **"Single Node"**: Use these displays to investigate performance and utilization metrics of a single node.
- **"Time Range Analysis"**: 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.

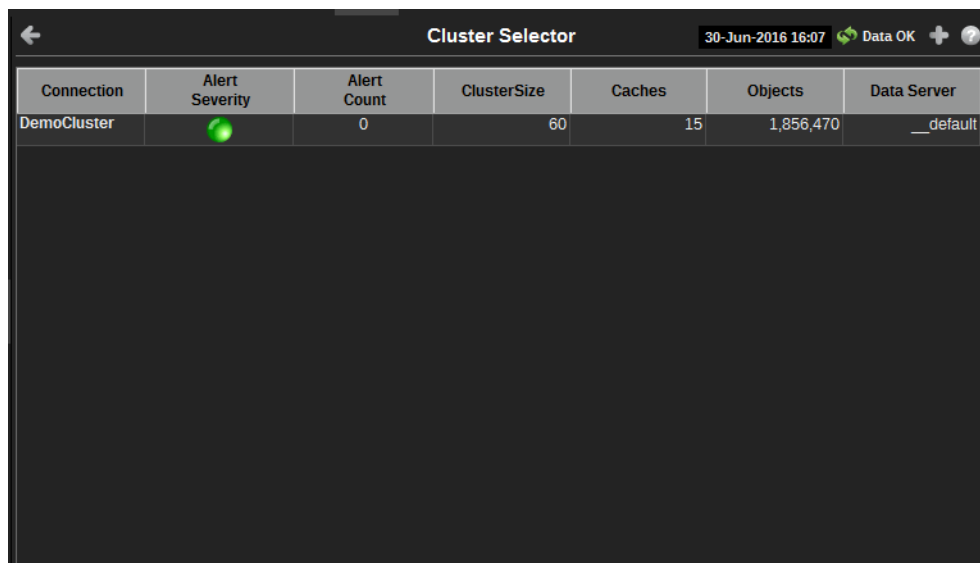
Cluster Selector

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.



The screenshot shows a web-based interface titled "Cluster Selector". At the top right, it displays the date and time "30-Jun-2016 16:07" and a status "Data OK" with a green checkmark icon. Below the title is a table with the following columns: Connection, Alert Severity, Alert Count, ClusterSize, Caches, Objects, and Data Server. The table contains one row for "DemoCluster". The "Alert Severity" cell contains a green circle icon. Below the table, there is a large, empty dark gray rectangular area.

Connection	Alert Severity	Alert Count	ClusterSize	Caches	Objects	Data Server
DemoCluster		0	60	15	1,856,470	__default



Connection	The name of the user defined connection that is used to connect to the monitored Coherence cluster.
Alert Severity	<p>The maximum level of alerts on the cluster.</p> <ul style="list-style-type: none"> 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 529](#): 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 533](#): View cache and node utilization hot spots and currently active alerts.
- ["Memory/Network Health" on page 536](#): Assess cluster memory utilization and packet transmission success/failure trends, and see weakest nodes.
- ["Stability Metrics" on page 538](#): Troubleshoot nodes joining and leaving the cluster, view HA status for cache services.
- ["All Services History" on page 540](#): Assess capacity utilization, over time, by all services in a cluster.
- ["All Caches History" on page 543](#): Assess capacity utilization and distribution for all caches in a cluster, and quickly identify potential bottlenecks.
- ["All Nodes History" on page 547](#): 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):

- 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.

Fields and Data:

Coherence Cluster Configuration

- 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.
Cluster Memory Usage Totals	
Senior Node	Node ID of the senior node of the cluster.
Client Nodes	Monitor client 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	<p>The maximum level of alerts for all nodes in the cluster. Click to drill down to the Alert Detail Table.</p> <p> Red indicates that one or more exceeded their ALARM LEVEL threshold.</p> <p> Yellow indicates that one or more exceeded their WARNING LEVEL threshold.</p> <p> Green indicates that none have exceeded their alert thresholds.</p> <p>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.</p> <p>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.</p> <p>Stability Represents the current most critical state of alerts related to cluster stability for all nodes in the cluster. For example, the DepartedNodePercentage alert.</p> <p>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.</p> <p>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.</p> <p>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 CapacityLimitAllCaches alert.</p> <p>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.</p>

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 Name	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-SAFE	If 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).
ENDANGERED	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

Most Gets	Track services performing the greatest number of gets in the cluster. The total is 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.
Largest Cache	Track caches that consume the greatest amount of capacity. Click to drill-down to the All Caches - "Current Size Chart" display.

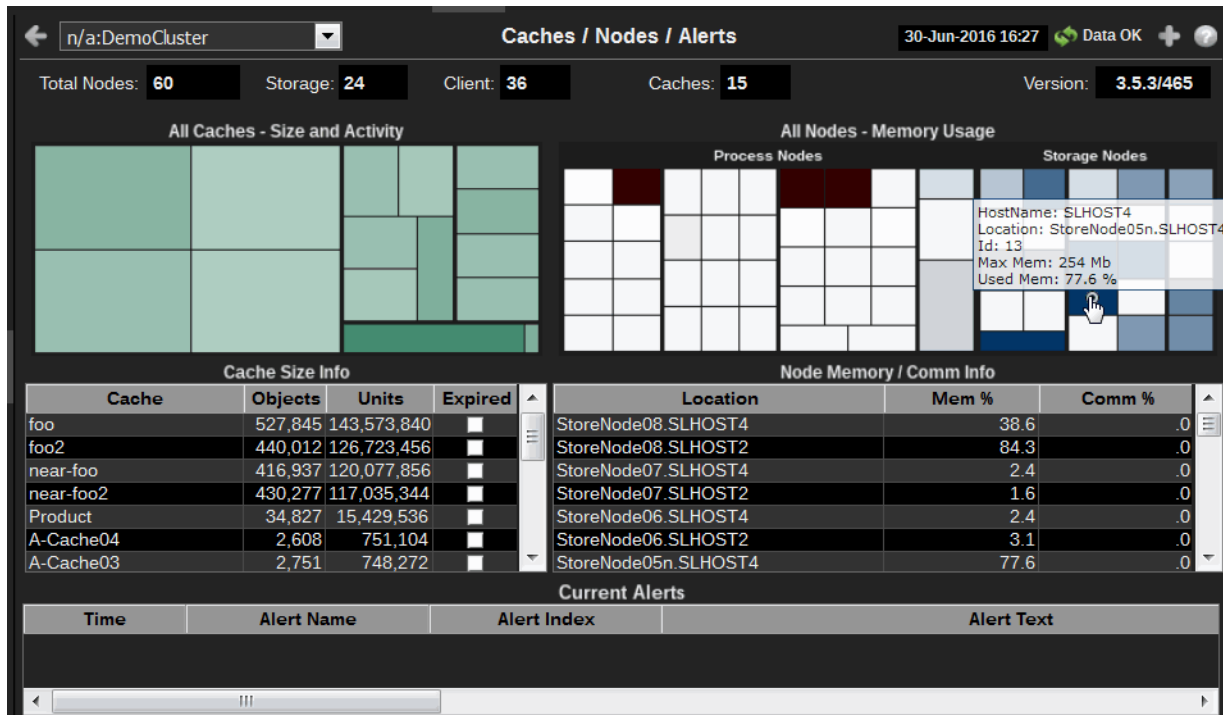
Cluster Stability

Node Uptimes	Monitor cluster stability and how often nodes are restarted (for example, every month, every day, every hour, and so forth). If the number of nodes running for 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 Utilization%	Monitor memory utilization for all nodes in the cluster.
Average	The average memory utilization for all nodes in the cluster.

	<p>Worst Node The most amount of memory consumed by a single node in the cluster. A slow node that provides data to other nodes can cause 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.</p>
Communication Success%	<p>Monitor cluster packet loss--an 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.</p> <p>The 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.</p> <p>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.</p> <p>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.</p>
	<p>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.</p>
	<p>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.</p>

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.



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 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.

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 & Memory 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. Click 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 is the maximum gets seen.

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.

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 tool-tip. 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 Memory/Comm Info

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.

● 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.
Acknowledged	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 Reason	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 Time	The time the alert was cleared.
Alert Index Value	The Oracle Coherence source (node or cache) from which the alert originated.
Cluster Connection	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.



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.
- 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.

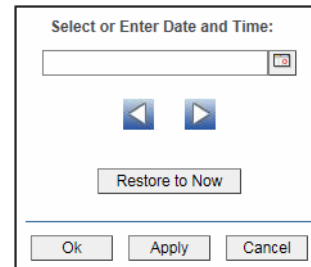
23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

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 voted by Coherence as the weakest in the cluster. The Weakest Node often points to a server/node that is causing performance issues. The node value most often appears in the "weakest node" attribute of all the JMX "node" objects. The format of this string is **<Node IP Address>:< Node Port >/<NodeID>**.



	Weak	The percent of the Coherence nodes that "elected" the node as the weakest.
Worst Network		The node that has the longest network queue in the cluster.
	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.
Worst Memory		The node that has the lowest available memory of any node in the cluster.
	Mem Used	The percent of memory consumed on the Worst Memory node.
Average over all Process / Storage Nodes	Trend Graphs	The trend graphs show aggregated performance metrics for storage and process nodes.

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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field at the top with a calendar icon to its right. Below the input field are two blue navigation arrows, a left-pointing arrow and a right-pointing arrow. Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

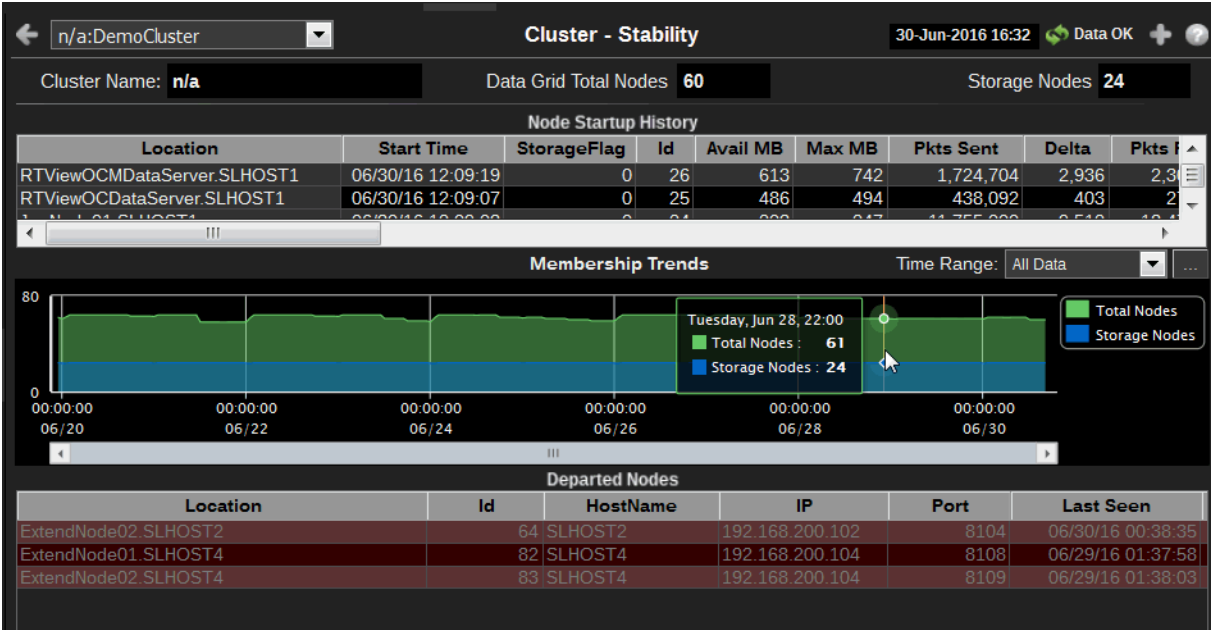
Click **Restore to Now** to reset the time range end point to the current time.

Process Nodes	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.



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.

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.


Storage Nodes

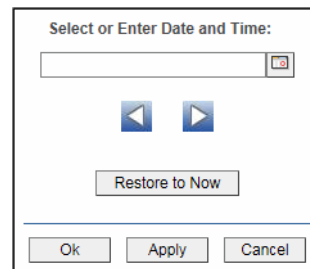
The total number of nodes in the cluster which have storage enabled for any cache. This value is equal to the total nodes when replicated caches are being used. The number is less when only distributed cache types are utilized.


Node Startup History

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).

Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site .
Start Time	The date and time that the node joined the cluster.
StorageFlag	Indicates whether storage is enabled (0 or 1).
Id	The short member id that uniquely identifies this member.
Avail MB	The amount of available memory for this node, in megabytes.
Max MB	The maximum amount of memory for this node, in megabytes.
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.
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-to-point communication.

Membership Trends	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.
	Product-Edition	The product edition this node is running. Possible values are: Standard Edition (SE), Enterprise Edition (EE), Grid Edition (GE).
	Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar  .



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

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

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


Departed Nodes	Track departed nodes by IP address, port number and time last seen.	
Location	A unique identifier for each node. It is defined as: member_name.machine.rack.site .	
HostName	The name of the host on which the node resides.	
IP	The node IP address.	
Port	The unicast port the node used while in the cluster. This is the port of the node's DatagramSocket for point-to-point communication.	
Last Seen	The 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.



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.

Filter By:


The display might include these filtering options:

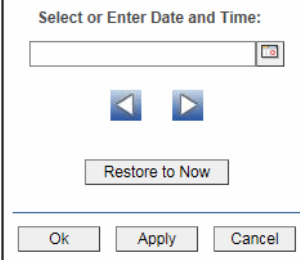
Service Metric:	Choose a service metric for which to display data in the heatmap. Use the mouse-over tool-tip to view metrics. Identify a service with high utilization. Perform node analysis by clicking One to view the "Single Service History" display.
CPU%	Percent of CPU utilization in the specified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

TaskBacklog The size of the backlog queue that holds tasks scheduled to be executed by one of the service threads. Use this metric for 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.


RequestPending-Count The number of pending requests issued by the service.



RequestAverage-Duration The average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon to its right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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 service details visible upon mouseover.

History Heatmap of Selected Metric by Service Use the heatmap to view utilization trends for all services, 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 service details in the ["Single Service Summary"](#) display.

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 on Process 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 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.


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.

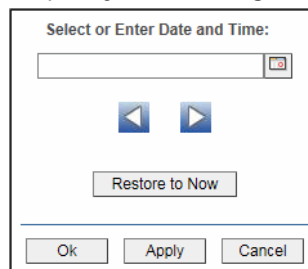
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.


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.



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, how often requests are made for data that does not exist in the cache. If a cache has a high rate of misses, consider performing a lower level analysis by viewing the cache in the "Single Cache 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.
StoreFailures (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.
StoreReadMillis (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.
StoreWritesMillis (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 range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



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

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


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

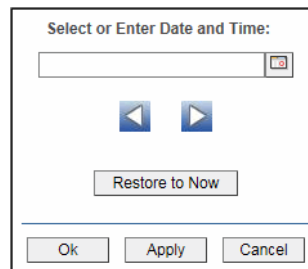
AppName: Choose an AppName to show data for in the display.


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



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

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

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

**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 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.

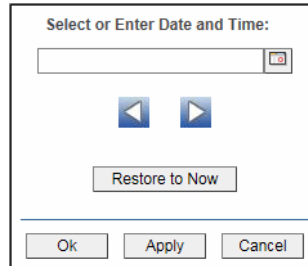
The heatmap shows cache utilization trends for the selected service and metric, for all caches running in the cluster. Each row represents a cache. Cells in a row are sized uniformly and represent one process node. Each column represents a time period (typically in 10 second intervals). The heatmap is 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.

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 .

A dialog box titled "Select or Enter Date and Time:". It contains a text input field with a calendar icon to its right. Below the input field are two blue navigation arrows, one pointing left and one pointing right. Below the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "Cancel".

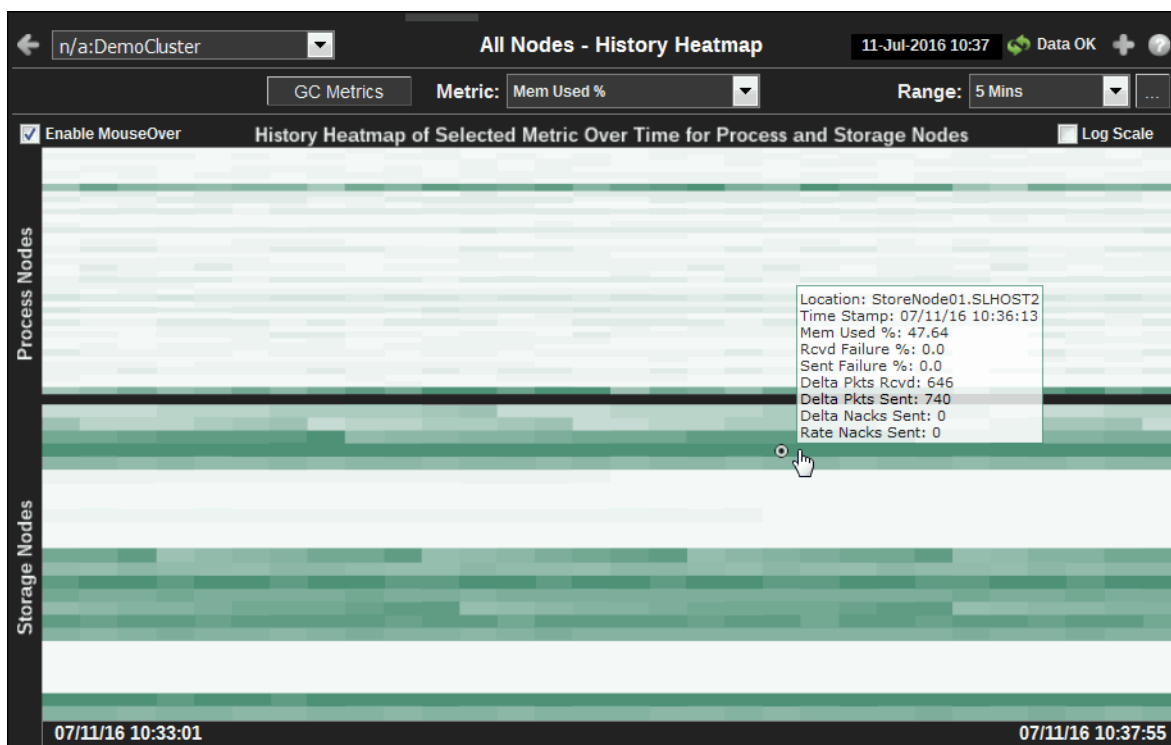
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   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):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.


Filter By:


- Cluster:** Select a cluster for which to display data in the heatmap.
- GC Metrics** Click to open the [“All Nodes History”](#) display which shows GC Duty Cycle for all the nodes in a cluster.
- Metric:** 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.



Delta Packets Sent	The number of packets sent by this node since the last data sample.
Delta Packets Rcvd	The number of packets received by this node since the last data sample.
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   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 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.

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 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 Services

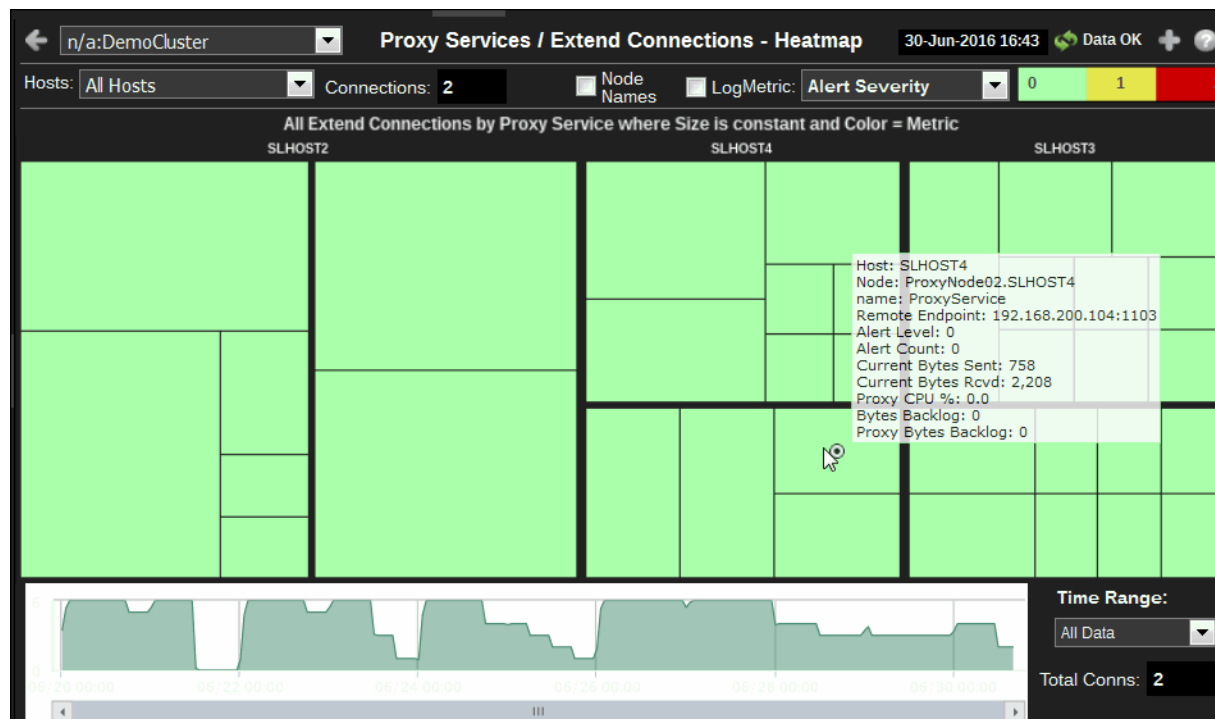
Proxy Services displays present detailed proxy server performance metrics for the cluster. Use the Proxy Services displays to quickly identify overloaded proxy services and locate the extend client connection causing the issue.

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 551](#): Table shows proxy services data with trend graphs/tables of extend connection detail for a specified location.
- [“Proxy / Extend Detail” on page 557](#): 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 561](#): Heatmap shows performance utilization, over time, for all proxy service nodes in the selected cluster.
- [“Extend Connections History” on page 563](#): 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.




Title Bar (possible features are):

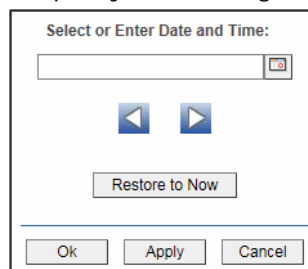
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.

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.



- 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 Severity** The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient bar, where 2 is the highest Alert Severity:
- 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 Sent	Total number of bytes sent by the selected proxy in the time range specified.
Current Bytes Received	Total number of bytes received by the selected proxy in the time range specified.
Proxy CPU%	The average percent CPU utilization for the selected proxy.
Bytes Backlog	The number of pending bytes in the Extend outgoing queue.
Proxy Bytes Backlog	The number of pending bytes in the Proxy outgoing queue.
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  .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right) and a "Restore to Now" button. At the bottom of the dialog are "Ok", "Apply", and "Cancel" buttons.

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

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

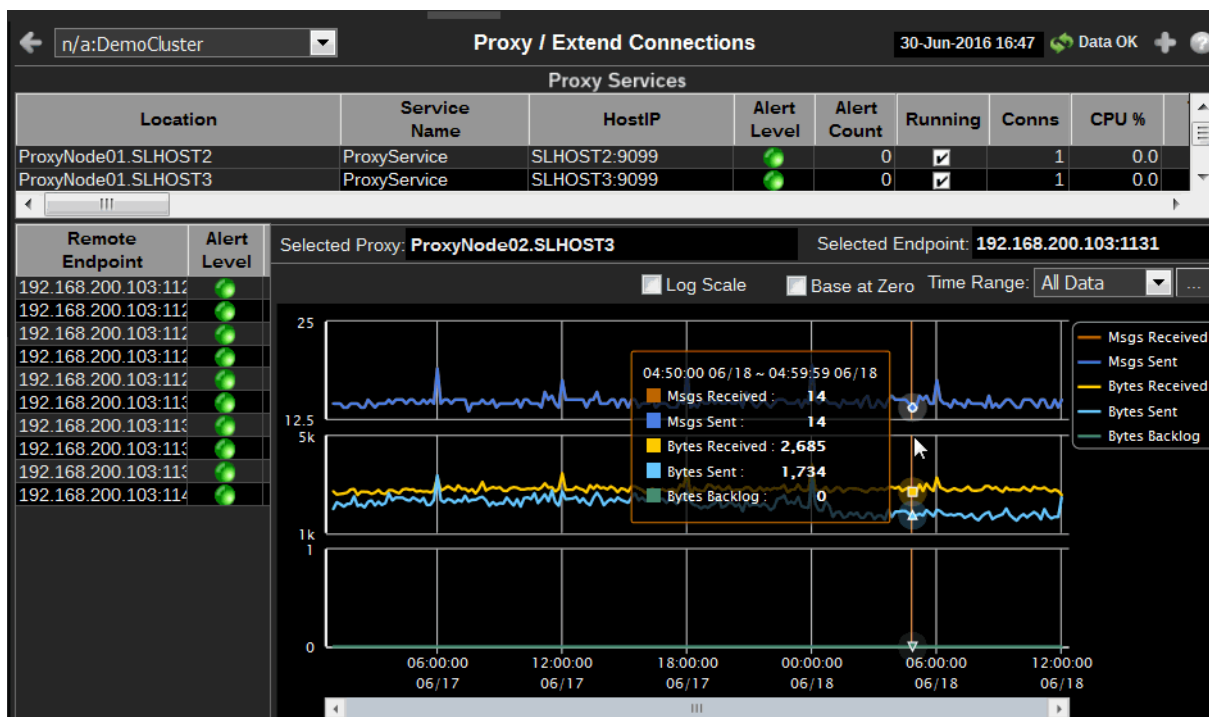
Click **Restore to Now** to reset the time range end point to the current time.

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.



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.

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.
RequestAverageDuration	The average duration (in milliseconds) of an individual synchronous request issued by the proxy service since the last time the statistics were reset.
RequestMaxDuration	Maximum duration (in milliseconds) of an individual proxy service request since the last time the statistics were reset.
RequestTotalCount	The number of requests issued and received by the proxy service.
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 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.
TaskHungDuration	The longest currently executing hung task duration in milliseconds.
TaskHungTaskId	The id of the of the longest currently executing hung task.

TaskHungThresholdMillis	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.
TaskTimeoutCount	The total number of timed-out proxy service tasks since the last time the statistics were reset.
RequestPendingCount	The number of pending proxy service requests.
RequestPendingDuration	The average duration (in milliseconds) that an individual proxy service request waits before being executed.
RequestTimeoutCount	The total number of timed-out proxy service requests since the last time the statistics were reset.
RequestTimeoutMillis	The duration (in milliseconds) for a proxy service request to reach the specified timeout threshold.
TaskTimeoutMillis	The default timeout value (in milliseconds) for tasks that can be timed-out but do not explicitly specify the task execution timeout value.
IncomingBufferPoolSize	The number of buffers in the incoming pool.
ThreadAbandonedCount	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.
ThreadIdleCount	The number of currently idle threads in the proxy service thread pool.
AverageActiveThreadCount	The average number of proxy service active threads, not currently idle, since the last time the statistics were reset.
ThreadAverageActiveCount	The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.
AverageTaskDuration	The average duration (in milliseconds) to perform a proxy service task since the last time the statistics were reset.
MaximumBacklog	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.
ThroughputInbound	The amount of data (in kilobytes) that is transferred from clients to the proxy service.
ThroughputOutbound	The amount of data (in kilobytes) that is transferred from the proxy service to clients.
IncomingBufferPoolCapacity	The size (in kilobytes) of the proxy service incoming buffer pool.
OutgoingBufferPoolCapacity	The size (in kilobytes) of the proxy service outgoing buffer pool.
OutgoingBufferPoolSize	The number of buffers in the proxy service outgoing pool.
nodeId	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.
	SeniorMemberId	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.
	Max	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 populated 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.

Time Range

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



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

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

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

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

Proxy / Extend Detail

30-Jun-2016 16:49

Data OK

Proxy Services

Location	Service Name	HostIP	Alert Level	Alert Count	Running	Conns	CPU %	Total
ProxyNode01.SLHOST2	ProxyService	SLHOST2:9099		0	<input checked="" type="checkbox"/>	1	0.0	2
ProxyNode01.SLHOST3	ProxyService	SLHOST3:9099		0	<input checked="" type="checkbox"/>	1	0.0	2
ProxyNode01.SLHOST4	ProxyService	SLHOST4:9099		0	<input checked="" type="checkbox"/>	0	0.0	4
ProxyNode02.SLHOST2	ProxyService	SLHOST2:9099		0	<input checked="" type="checkbox"/>	0	0.0	
ProxyNode02.SLHOST3	ProxyService	SLHOST3:9099		0	<input checked="" type="checkbox"/>	0	0.0	
ProxyNode02.SLHOST4	ProxyService	SLHOST4:9099		0	<input checked="" type="checkbox"/>	0	0.0	

III

Extend Client Connections for ProxyNode02.SLHOST3:*

Location	Remote Endpoint	Alert Level	Alert Count	Time Stamp	Connect Time (mins)	Bytes Backlog	Message Backlog
ProxyNode02.SLHOST3	192.168.200.103:1131		0	06/17/16 01:03:29	2,098.6	0	
ProxyNode02.SLHOST3	192.168.200.103:1136		0	06/17/16 01:03:34	2,098.5	0	
ProxyNode02.SLHOST3	192.168.200.103:1128		0	06/19/16 01:03:33	659.1	0	
ProxyNode02.SLHOST3	192.168.200.103:1121		0	06/19/16 01:03:27	1,437.2	0	
ProxyNode02.SLHOST3	192.168.200.103:1124		0	06/22/16 01:03:32	2,818.6	0	
ProxyNode02.SLHOST3	192.168.200.103:1132		0	06/22/16 01:03:37	2,098.7	0	
ProxyNode02.SLHOST3	192.168.200.103:1136		0	06/24/16 01:03:40	1,018.4	0	
ProxyNode02.SLHOST3	192.168.200.103:1140		0	06/24/16 01:03:45	1,018.3	0	
ProxyNode02.SLHOST3	192.168.200.103:1121		0	06/26/16 01:03:29	1,437.2	0	
ProxyNode02.SLHOST3	192.168.200.103:1127		0	06/26/16 01:03:34	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.

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.

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.
RequestAverageDuration	The average duration (in milliseconds) of an individual synchronous request issued by the proxy service since the last time the statistics were reset.
RequestMaxDuration	Maximum duration (in milliseconds) of an individual proxy service request since the last time the statistics were reset.
RequestPendingCount	The number of pending proxy service requests.
RequestPendingDuration	The average duration (in milliseconds) that an individual proxy service request waits before being executed.
RequestTimeoutCount	The total number of timed-out proxy service requests since the last time the statistics were reset.
RequestTimeoutMillis	The duration (in milliseconds) for a proxy service request to reach the specified timeout threshold.
RequestTotalCount	The number of requests issued and received by the proxy service.
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 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.
TaskHungDuration	The longest currently executing hung task duration in milliseconds.
TaskHungTaskId	The id of the of the longest currently executing hung task.

TaskHungThresholdMillis	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.
TaskTimeoutCount	The total number of timed-out proxy service tasks since the last time the statistics were reset.
TaskTimeoutMillis	The default timeout value (in milliseconds) for tasks that can be timed-out but do not explicitly specify the task execution timeout value.
IncomingBufferPoolSize	The number of buffers in the incoming pool.
ThreadAbandonedCount	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.
ThreadIdleCount	The number of currently idle threads in the proxy service thread pool.
AverageActiveThreadCount	The average number of proxy service active threads, not currently idle, since the last time the statistics were reset.
ThreadAverageActiveCount	The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.
AverageTaskDuration	The average duration (in milliseconds) to perform a proxy service task since the last time the statistics were reset.
MaximumBacklog	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.
ThroughputInbound	The amount of data (in kilobytes) that is transferred from clients to the proxy service.
ThroughputOutbound	The amount of data (in kilobytes) that is transferred from the proxy service to clients.
IncomingBufferPoolCapacity	The size (in kilobytes) of the proxy service incoming buffer pool.
OutgoingBufferPoolCapacity	The size (in kilobytes) of the proxy service outgoing buffer pool.
OutgoingBufferPoolSize	The number of buffers in the proxy service outgoing pool.
nodeId	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.
SeniorMemberId	The proxy service senior member id. If the proxy service is not running, it is -1.

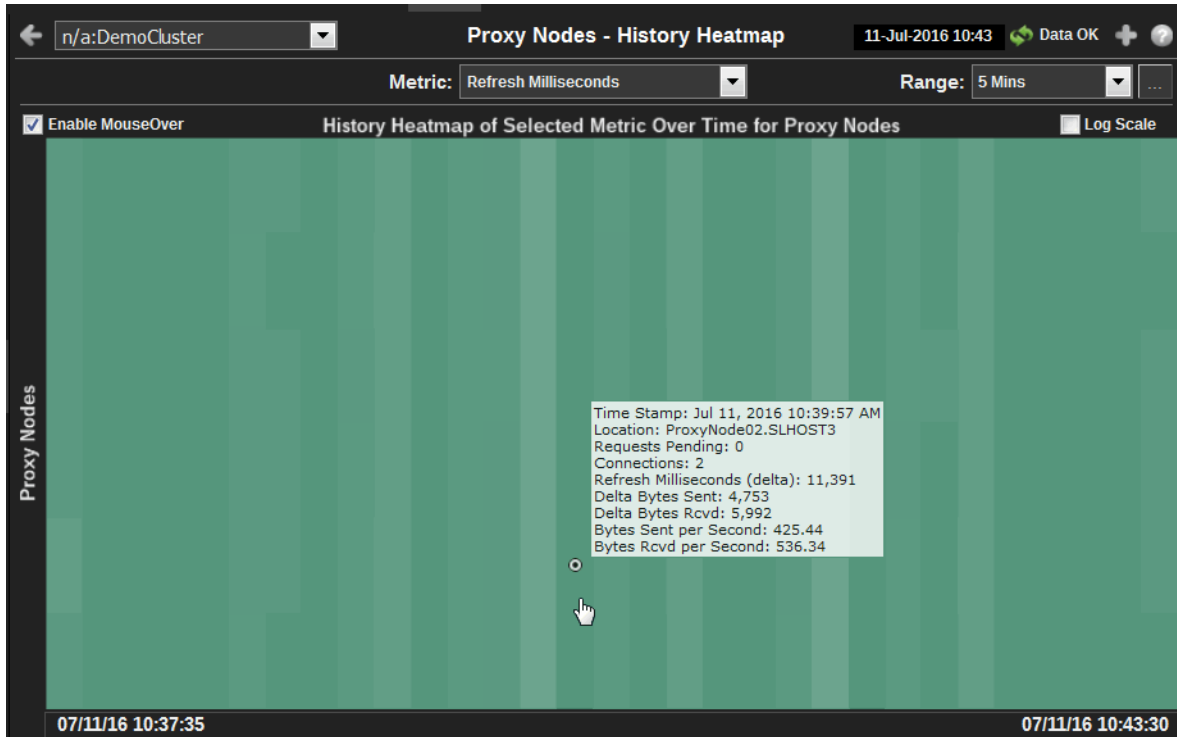
**Extend
Client
Connections**

Select a row from the 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.
OutgoingByteBacklog	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.
OutgoingMessageBacklog	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.
TotalBytesReceived	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.
TotalMessagesReceived	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.
TotalMessagesSent	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.
nodeId	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.



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.


Cluster Select a cluster from the drop-down menu.

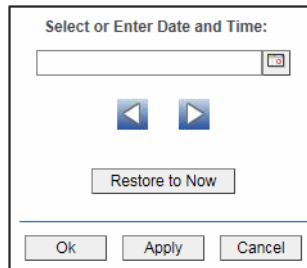
Metric Select a metric from the drop-down menu.


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



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

Range

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



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

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

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

Enable MouseOver

Select this option to make details visible upon 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.

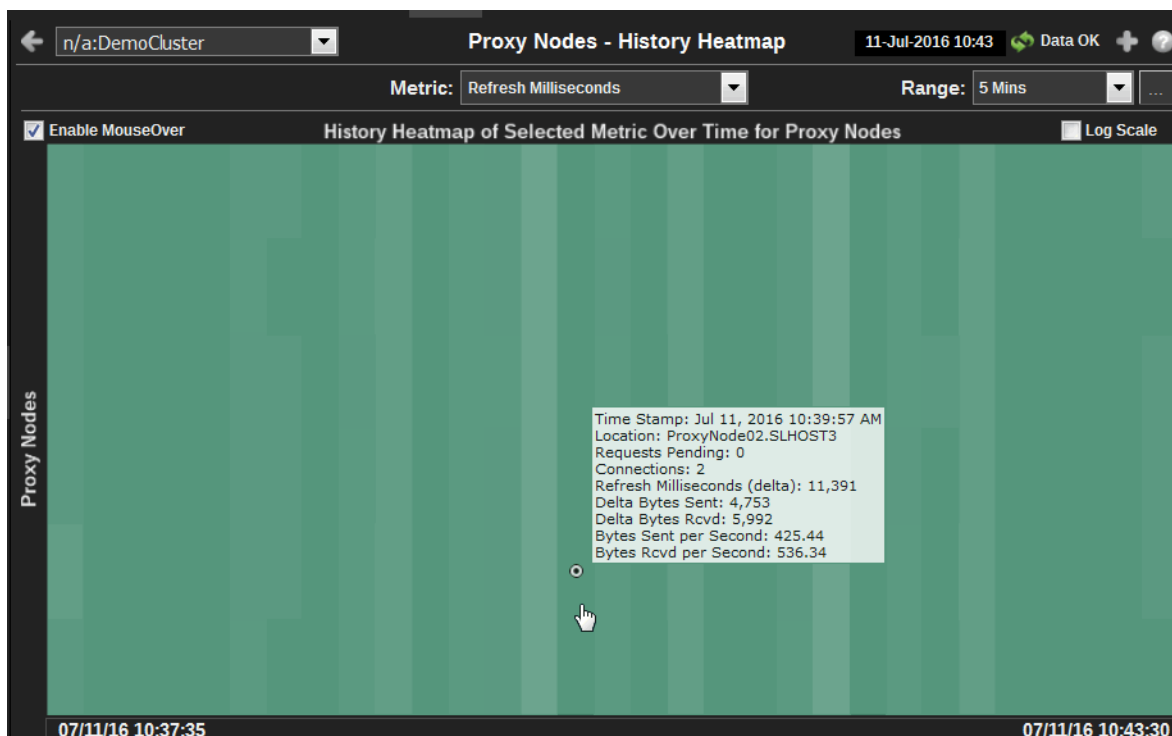
Proxy Nodes Heatmap

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 mouseover 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.



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.


Cluster Select a cluster from the drop-down menu.

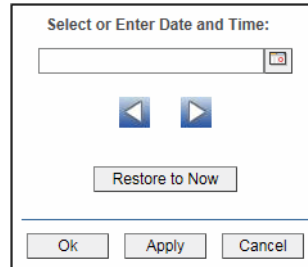
Metric Select a metric from the drop-down menu.

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


Bytes per Second	Total bytes received, per second, by the node.
Msgs Sent per Second	Total messages sent, per second, by the node.
Msgs Rcvd per Second	Total messages received, per second, by the node.



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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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 details visible upon 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 mouseover tool-tip to see the node hostname and data for the selected metric.

Cache Services

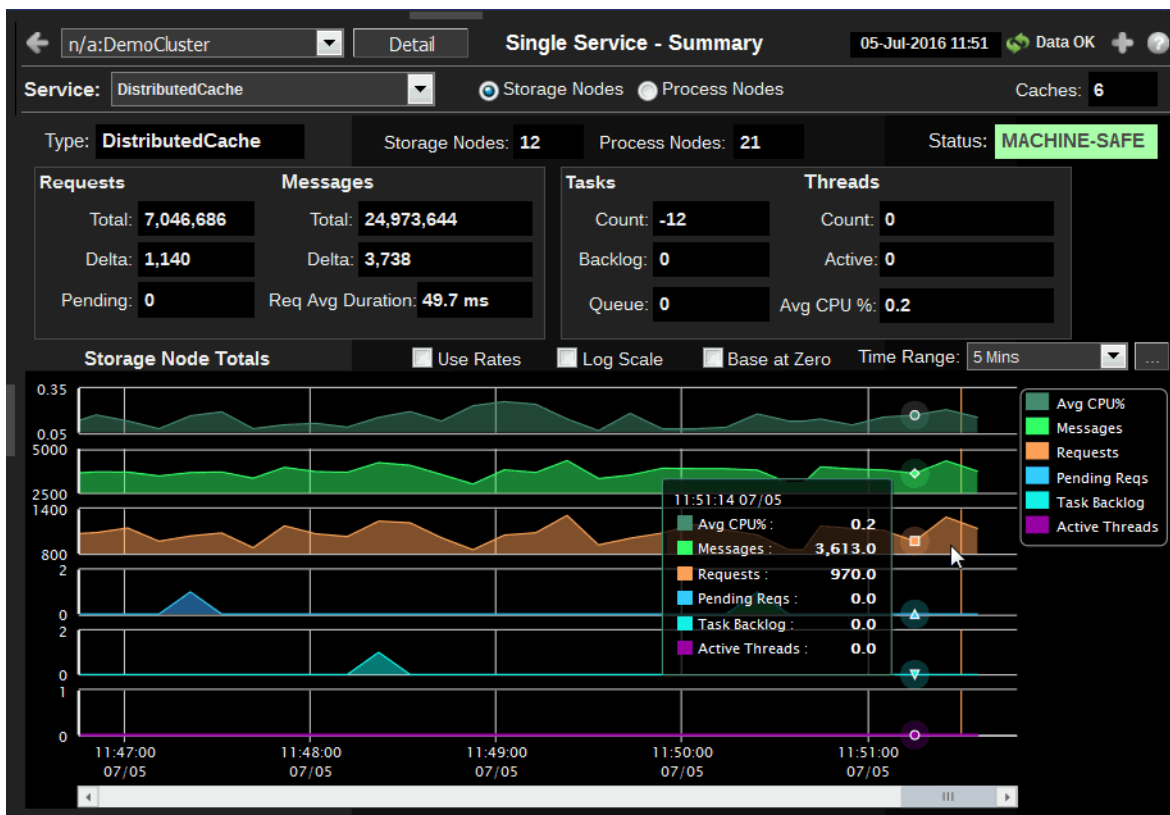
Cache Services displays present detailed service performance metrics for the cluster. Use the Cache Services displays to quickly identify overloaded services and locate the client connection causing the issue.


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 565](#): Trend graphs show performance metrics for a single service aggregated across all nodes.
- [“Service Metrics Overview” on page 569](#): 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 570](#): 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 572](#): 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 574](#): 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.





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.
Type	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	<p>The high availability status of the service:</p> <ul style="list-style-type: none"> ● 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	<p>Requests executed by the service.</p> <p>Total The number of requests executed.</p> <p>Rate / Delta Use the Use Rates checkbox to toggle between two value types: Rate and Delta (as labeled in the display upon selection).</p> <p>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.</p> <p>Pending The number of pending requests.</p>

Messages

Messages executed by the service.

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.

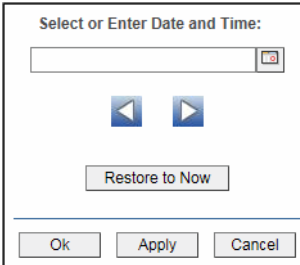
Rate is the value per second. The **Rate** value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the **Rate** value does not vary if the sample period changes (whereas the **Delta** value does vary). The **Rate** value enables you to 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 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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "Cancel".

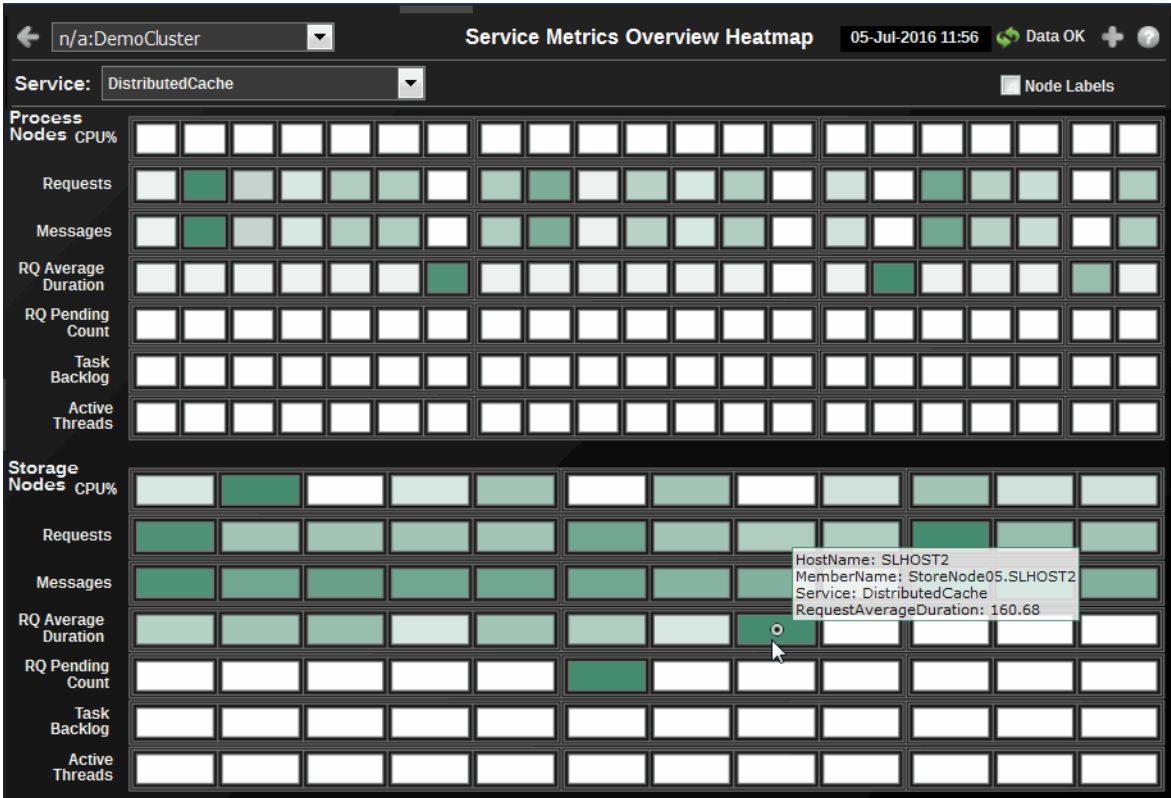
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   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.



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.

- 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.

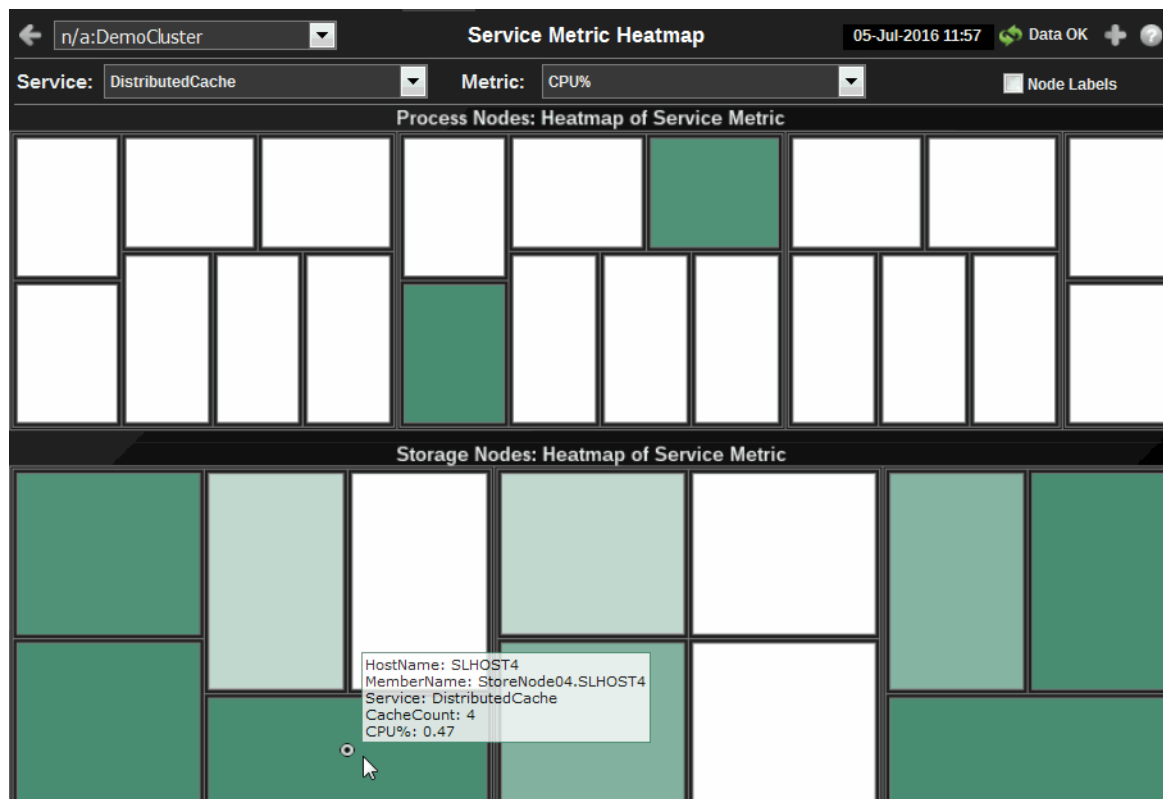
Request Pending Count Number of pending requests issued by the service.


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.

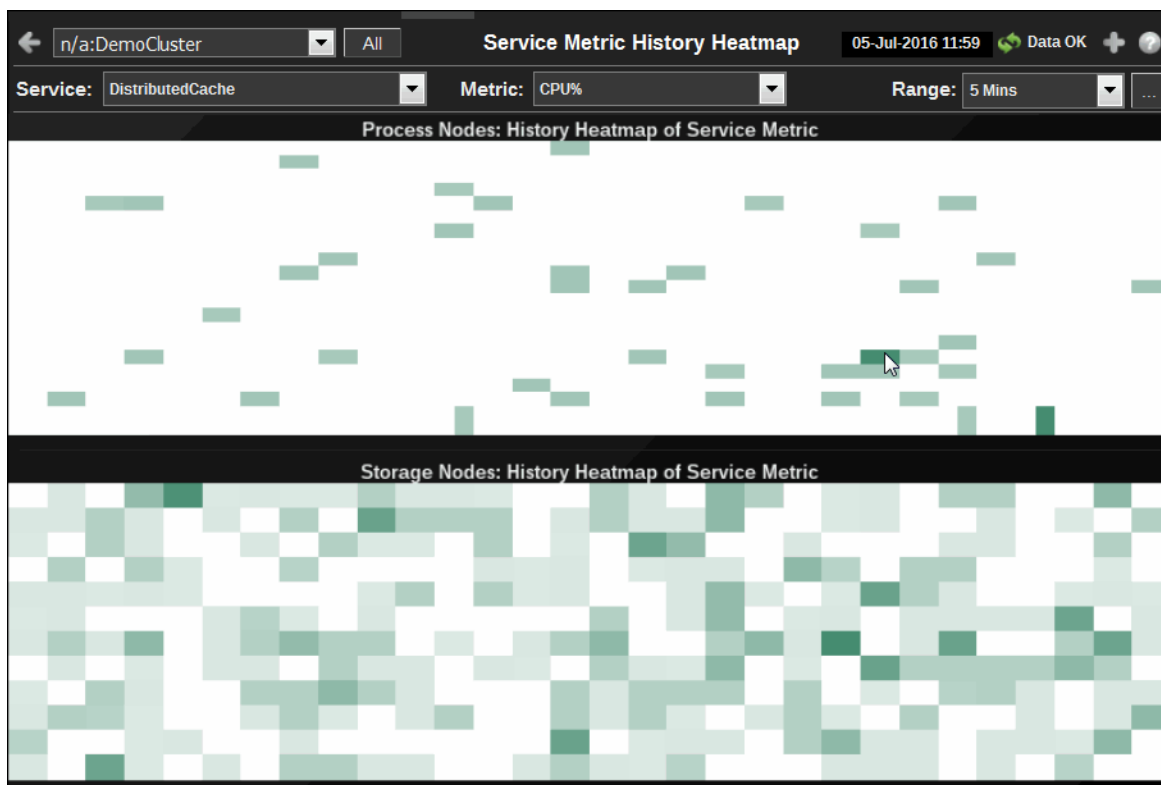




Cluster	Select a cluster to display.
Service	Select a service to display.
Node Labels	Select to display node labels.
Metric	<p>CPU% Percent of CPU utilization on the given node.</p> <p>Requests Number of requests issued by the service in the measured period.</p> <p>Request Average Duration Average duration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

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):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Cluster Select a cluster to display.

All 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.

Time Range

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



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

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

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

**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.

Location	Service	Running	StatusHA	Storage	CPU %	Messages
JmxNode01.SLHOST1	DistributedCache	✓	MACHINE-SAFE	■	0.0	195,635
ProcessNode01.SLHOST2	DistributedCache	✓	MACHINE-SAFE	■	0.0	376,689
ProcessNode01.SLHOST3	DistributedCache	✓	MACHINE-SAFE	■	0.0	1,149,263
ProcessNode01.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.0	1,131,618
ProcessNode04.SLHOST2	DistributedCache	✓	MACHINE-SAFE	■	0.0	188,285
ProcessNode04.SLHOST3	DistributedCache	✓	MACHINE-SAFE	■	0.0	574,088
ProcessNode04.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.0	564,263
ProcessNode05.SLHOST3	DistributedCache	✓	MACHINE-SAFE	■	0.0	1,150,706
ProcessNode05.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.0	1,134,618
ProcessNode05n.SLHOST2	DistributedCache	✓	MACHINE-SAFE	■	0.0	754,247
ProcessNode05n.SLHOST3	DistributedCache	✓	MACHINE-SAFE	■	0.0	2,334,486
ProcessNode05n.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.0	2,248,394
ProcessNode08.SLHOST3	DistributedCache	✓	MACHINE-SAFE	■	0.0	573,289
ProcessNode08.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.0	565,373
ProxyNode01.SLHOST2	DistributedCache	✓	MACHINE-SAFE	■	0.0	375,960
ProxyNode01.SLHOST3	DistributedCache	✓	MACHINE-SAFE	■	0.0	2,431
ProxyNode01.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.0	1,788
ProxyNode02.SLHOST2	DistributedCache	✓	MACHINE-SAFE	■	0.0	307
ProxyNode02.SLHOST3	DistributedCache	✓	MACHINE-SAFE	■	0.0	584,044
ProxyNode02.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.0	1,127,855
RTViewOCDataServer.SLHOST1	DistributedCache	✓	MACHINE-SAFE	■	0.0	92
StoreNode01.SLHOST1	DistributedCache	✓	MACHINE-SAFE	■	0.3	632,135
StoreNode01.SLHOST2	DistributedCache	✓	MACHINE-SAFE	■	0.0	1,223,758
StoreNode01.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.3	3,676,489
StoreNode04.SLHOST1	DistributedCache	✓	MACHINE-SAFE	■	0.1	640,945
StoreNode04.SLHOST4	DistributedCache	✓	MACHINE-SAFE	■	0.5	3,647,694

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.

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.

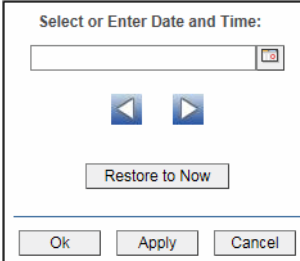
Cache Service Detail by Node:


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	<p>The high availability status of the service:</p> <ul style="list-style-type: none"> ● 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.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar  .



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

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

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

**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 577](#): Shows current information for all participating nodes for a selected cluster.
- [“Federated Destination Summary” on page 580](#): Shows current information and trended historical rate information.
- [“Federated Origin Detail” on page 582](#): Shows current information for all participating nodes for a selected cluster.
- [“Federated Origin Summary” on page 585](#): 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.

Location	BytesSentSecs	ConnectRetryTimeoutMillis	Connection	CurrentBandwidth	Delta
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):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

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.
EstimatedReplicateAllRemainingTime	The estimated remaining time, in milliseconds, to complete the replicateAll request.
Expired	When checked, this connection is expired due to inactivity.
GeoIp	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.
MsgApplyTimePercentileMillis	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the destination.
MsgNetworkRoundTripTimePercentileMillis	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.
RecordBacklogDelayTimePercentileMillis	The 90-percentile value , in milliseconds, of the time the journal records are in the cache waiting to be replicated.
ReplicateAllPercentComplete	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.



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.
- 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.

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**.
- Id:** 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 IP	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.

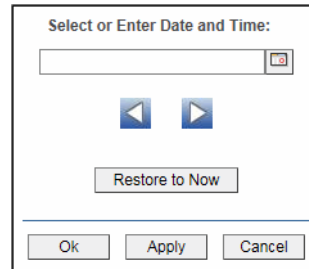
RateTotalMsgSent: Traces the total number of messages sent per second.

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.
- 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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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 Federated Origin Detail 24-Mar-2017 16:54 Data OK

Host: All Hosts

Federated Origin Detail by Node

Location	BytesReceivedSecs	Connection	DeltaTIME	STAMP	DeltaTotalBytesReceived	Delta
1.SLNB50	5,910	Cluster2Jmx	10,540		52,791	
2.SLNB50	9,305	Cluster2Jmx	10,540		99,502	

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.
- 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.

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.
MsgApplyTimePercentileMillis	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.
RecordBacklogDelayTimePercentileMillis	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

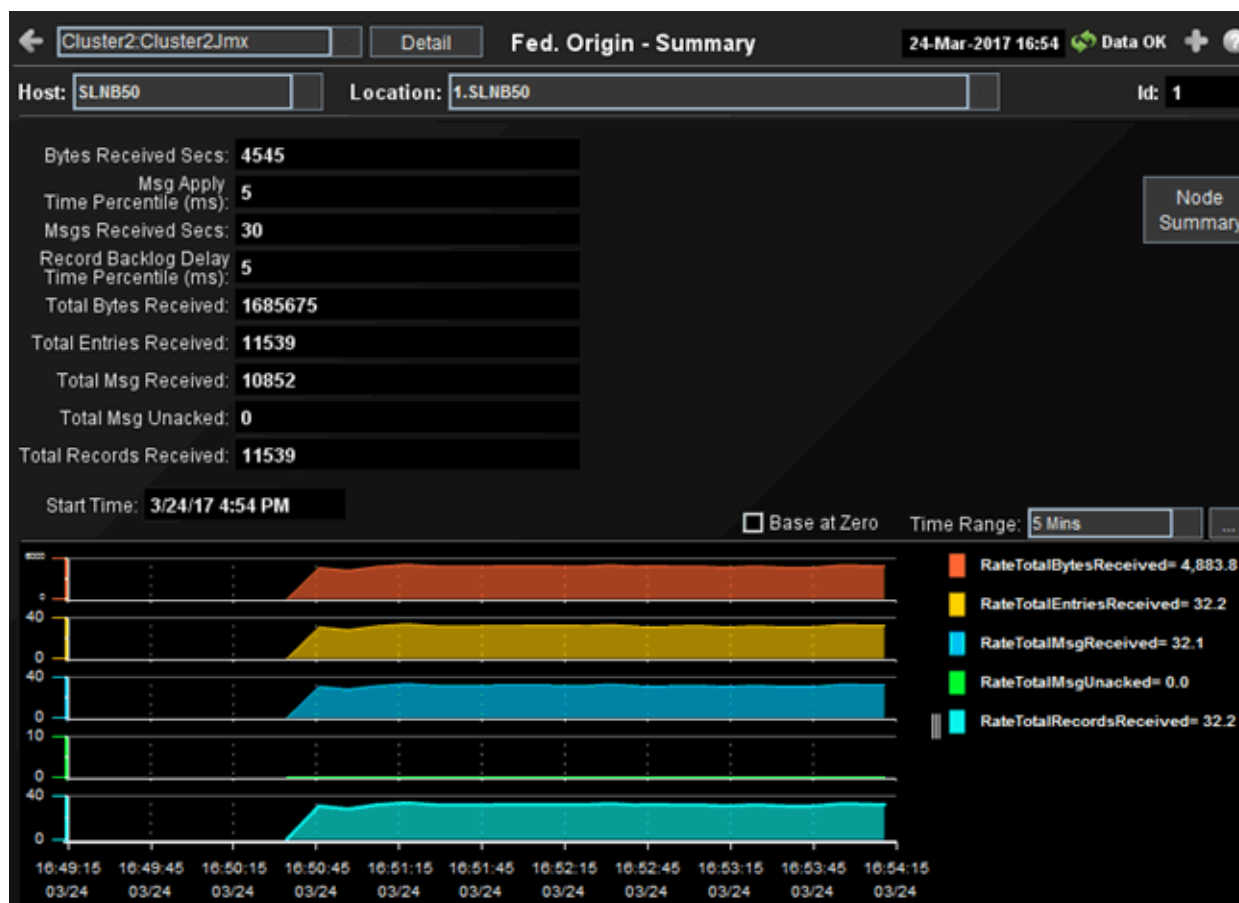
The Federated Service sub-type.

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.

**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.

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.


ReplicateAllPercentComplete: Traces the percent of completed **replicateAll** requests.

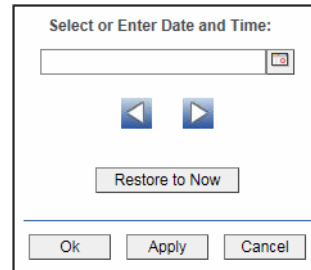
Start Time The start date and time.


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 .



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

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

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

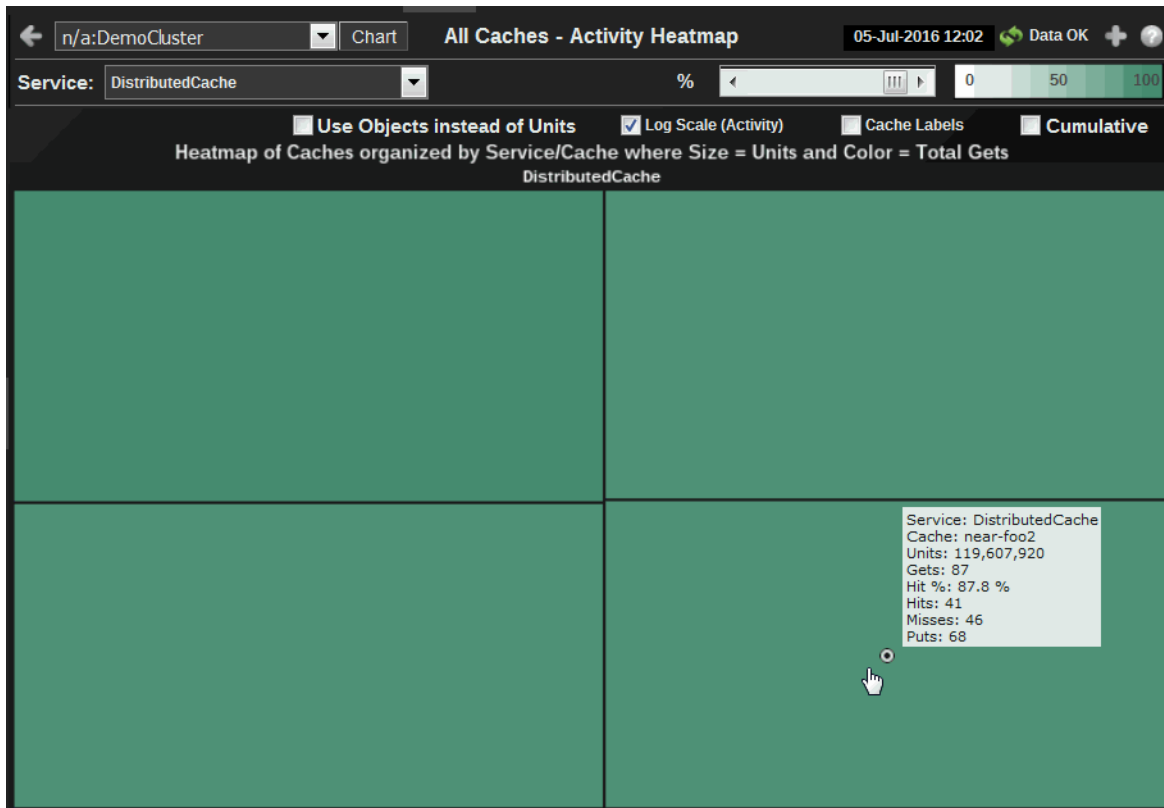
All Caches

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 588](#): Heatmap of caches by service where size represents Units and color represents Total Gets%.
- [“Storage Nodes Cache Map” on page 589](#): Heatmap of memory usage on storage nodes by service where size represents Units and color represents Units Used%.
- [“Current Size Chart” on page 591](#): Bar chart/table sorted by caches with largest size displays current size/capacity metrics.
- [“Current Activity Chart” on page 592](#): 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.



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.
- 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.

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

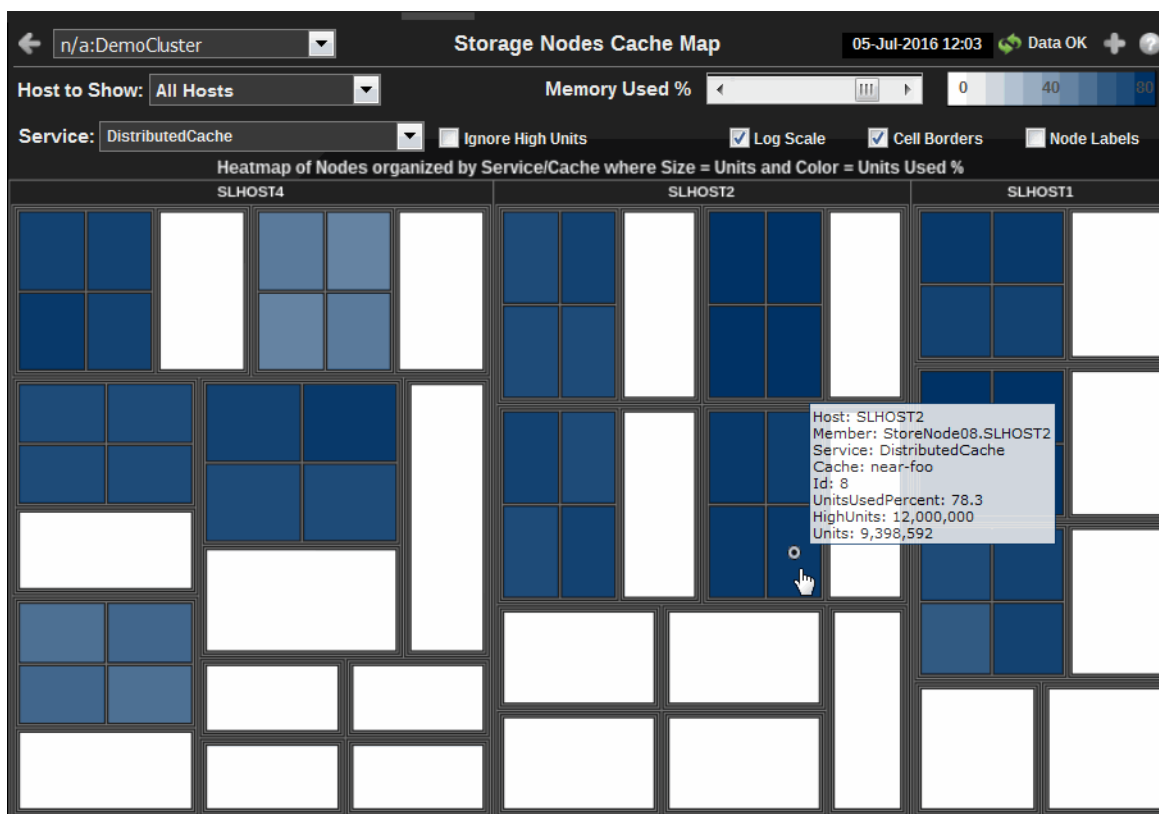
Select to show cumulative statistics for each cache.


**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.

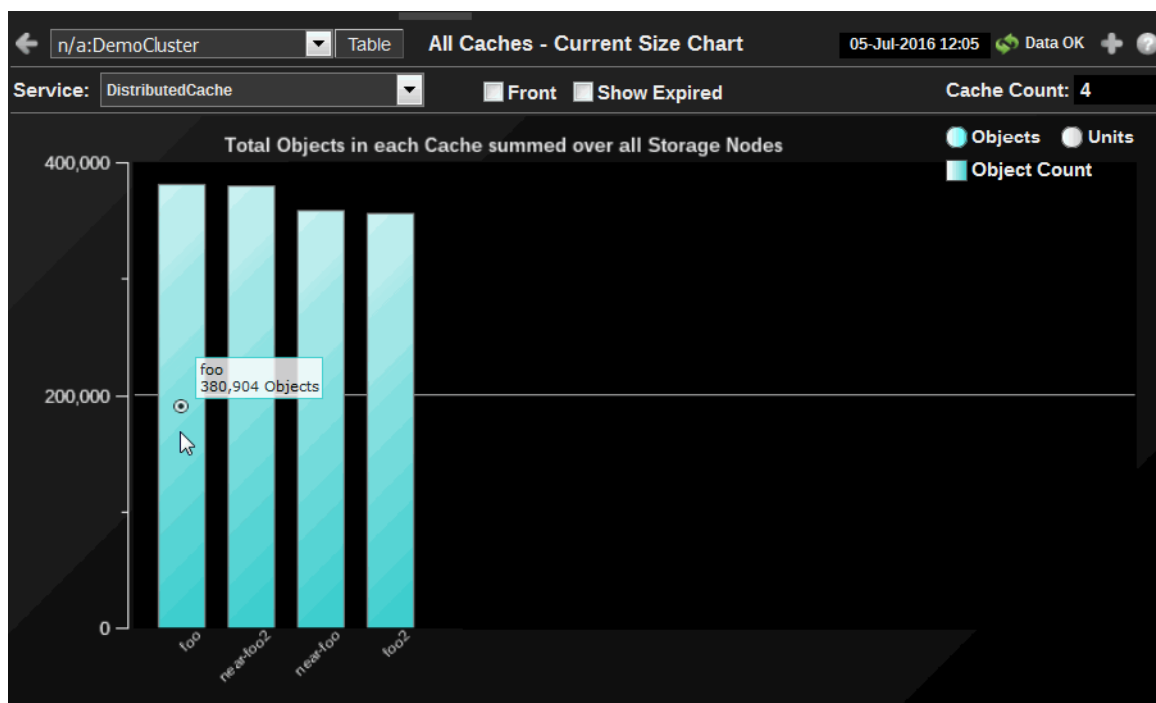




Cluster	Select a cluster to display.
Host to Show	Select a host to display.
Memory Used%	Set the memory used percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
Service	Select a service to display, or select All Services. NOTE: When you 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.

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

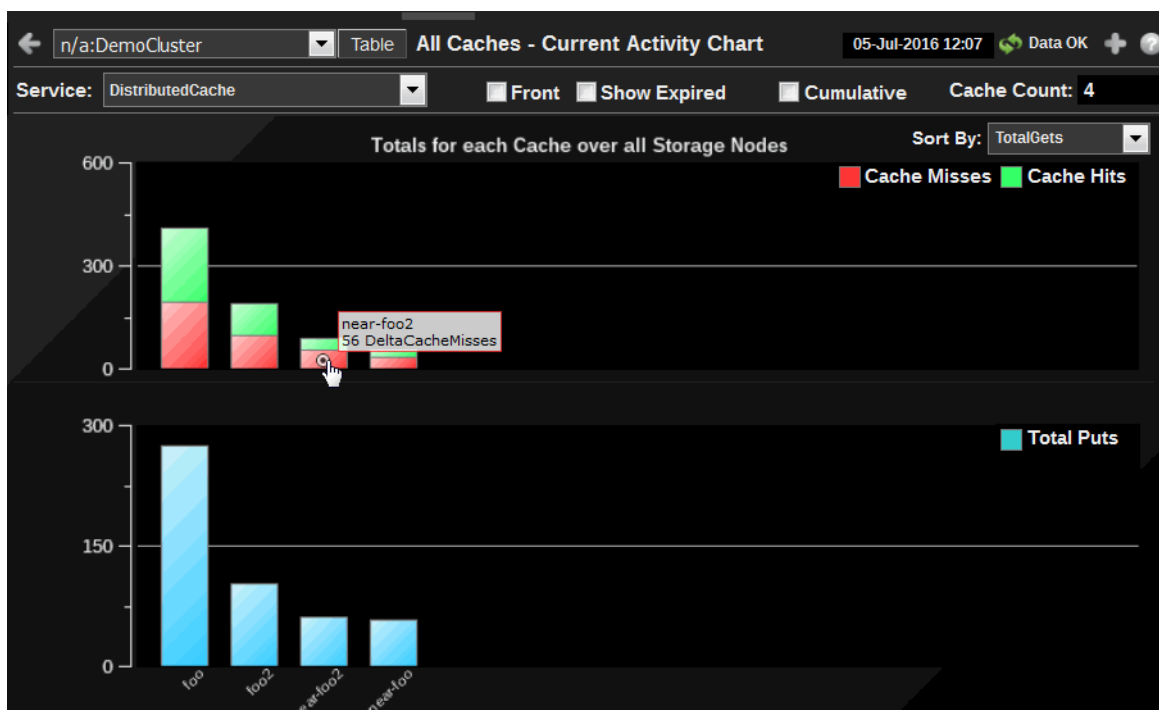
HighUnits Highest number of units before evictions occur

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.





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	<p>Totals for Cache summed over all Storage Nodes. This is the default view. Toggle to Table view.</p> <p>Sort by:</p> <p>Objects shows the total number of objects in this cache (Object Count).</p> <p>Units shows the highest number of units before evictions occur.</p> <p>Ignore High Units removes High Units bars from view.</p>
Current Activity Table	<p>Totals for each Cache over all Storage Nodes. Toggle to Chart view. Sort by:</p> <p>Cache Abbreviated name of cache</p> <p>tier Front or back</p> <p>Hits Total number of successful gets</p> <p>Misses Total number of failed gets</p> <p>Gets Total requests for data from this cache</p> <p>Puts Total data stores into this cache</p> <p>Hit% Ratio of hits to gets</p> <p>Service Service Name</p> <p>Cache Full Name Full name of cache</p>


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 594](#): Perform low level utilization analysis on a single cache.
- [“Size Trends” on page 597](#): Trend chart displays size/capacity metrics.
- [“Activity Trends” on page 600](#): Trend chart displays activity metrics.
- [“Cache Detail Tables” on page 602](#): Table showing current detailed cache statistics by node.
- [“Storage Manager Detail” on page 604](#): Table showing store manager metrics.
- [“Node/Group Distribution” on page 606](#): Bar chart displays metrics showing distribution across cluster nodes or groups.
- [“Front/Back Analysis” on page 608](#): 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.




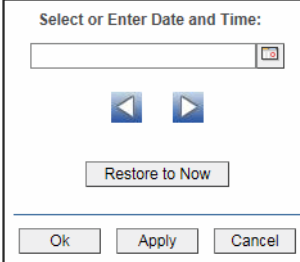


Cluster	Select a cluster to display.
Service	Select a service to display.
Cache	Select a cache. Click the Detail button to get information specific to the selected cache.
Front	Select for front tier, deselect for back tier.
Type	The type identifier string from the ServiceMBean (ReplicatedCache, DistributedCache, 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.
Type	The type of cache.
Storage Nodes	The number of storage nodes in the cache.
Status	<p>The high availability status of the service:</p> <ul style="list-style-type: none"> ● 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	<p>Units indicates memory usage for the back tier and number of objects for the front tier.</p> <p>Objects The number of objects in the selected cache. The value is the total across all storage nodes.</p> <p>Avg Size The average size of objects in the selected cache (in bytes if it is the back tier).</p> <p>Units The memory usage if back tier, or number of objects if front tier. The value is the total across all storage nodes.</p> <p>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.</p> <p>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.</p>
Persistence	<p>Type The persistence type for the cache. Possible values include: NONE, READ-ONLY, WRITE-THROUGH, and WRITE-BEHIND.</p> <p>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.</p> <p>Wr Queue The size of the queue, in kilobytes, that holds data scheduled to be written to the cache store.</p> <p>Writes The number of objects (cache entries) written to the cache store.</p>
Activity	<p>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.</p> <p>Cumulative: The total since the service was started for the selected cache, or since statistics were reset.</p> <p>Gets The number of requests for data from this cache.</p> <p>Hits The number of successful gets.</p> <p>Misses The number of failed gets.</p> <p>Puts The number of data stores into this cache.</p> <p>Evictions The number of objects removed to make room for other objects.</p>
Use Rates	<p>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.</p> <p>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.</p>
Log Scale	<p>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.</p>
Base at Zero	<p>Use zero for the Y axis minimum for all graphs.</p>



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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right) and a "Restore to Now" button. At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

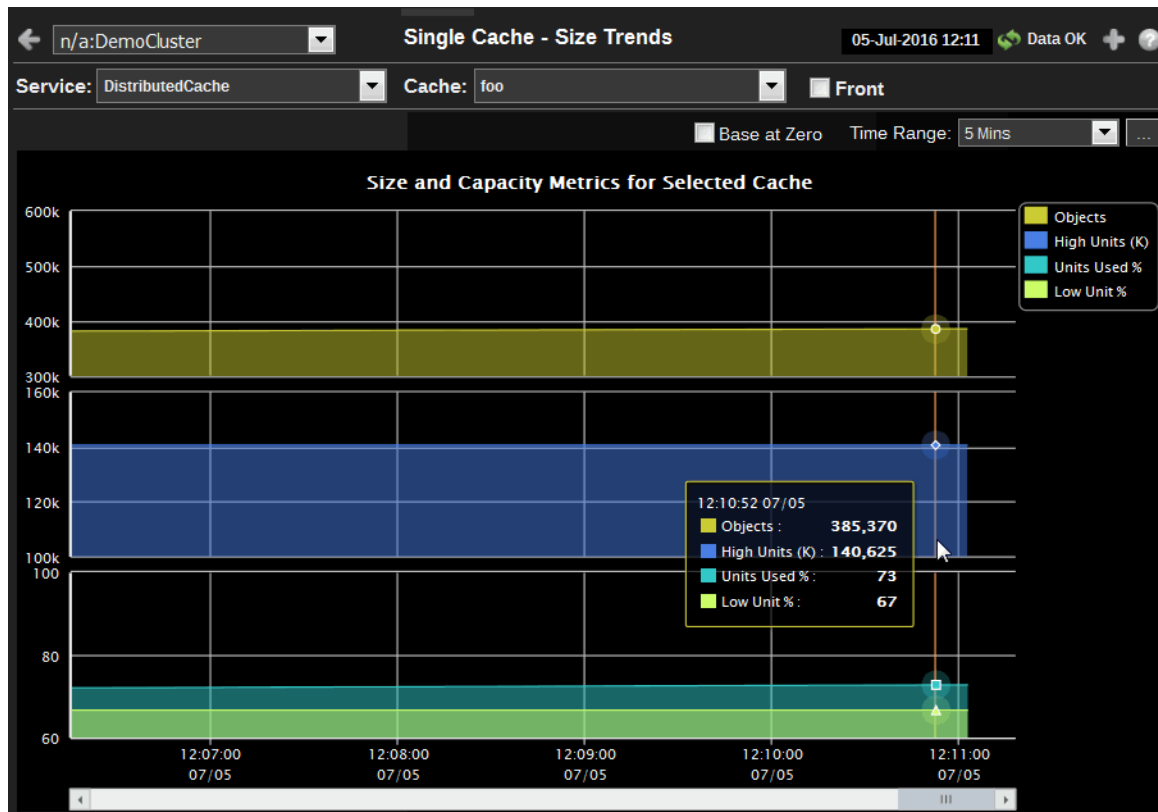
TotalPuts Total data stores into this cache.

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.



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.

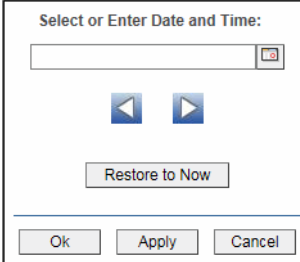
- Cluster** Select a cluster to display.
- Service** Select a service to display.
- Cache** Select a cache. Click the Detail button to get information specific to the selected cache.
- Front** Select for front tier, deselect for back tier.

Base at Zero


Use zero for the Y axis minimum for all graphs.



Time 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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.
- 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.

- 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.


Log Scale

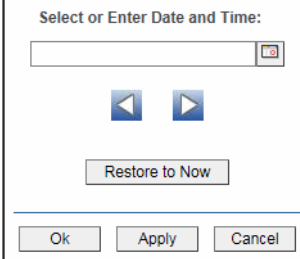
Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in 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 7 Days**, or display **All Data**. To specify a time range, click Calendar .



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

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

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

Activity Metrics for Selected Cache

Hits The number of successful gets from this cache.

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 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.

Objects	Units	LowUnits	HighUnits	Hits	Delta	Misses	Delta	Gets	Delta	Puts	Delta
389,542	105,955,424	96,000,000	144,000,000	4,366,039	282	1,383,079	193	5,749,118	475	3,500,956	310

AvgSize	Units	LowUnits	HighUnits	Hits	Delta	Misses	Delta	Gets	Delta	Puts	Delta
272	8,829,618	8,000,000	12,000,000	363,836	23	115,256	16	479,093	39	291,746	25

Location	tier	Objects	AvgSize	Units	LowUnits	HighUnits	Hits
StoreNode01.SLHOST1	back	34,541	272	9,395,152	8,000,000	12,000,000	1,669
StoreNode01.SLHOST2	back	37,492	272	10,197,824	8,000,000	12,000,000	183,778
StoreNode01.SLHOST4	back	21,943	272	5,968,496	8,000,000	12,000,000	734,269
StoreNode04.SLHOST1	back	33,379	272	9,079,088	8,000,000	12,000,000	1,628
StoreNode04.SLHOST4	back	31,346	272	8,526,112	8,000,000	12,000,000	658,192
StoreNode05.SLHOST2	back	32,555	272	8,854,960	8,000,000	12,000,000	202,794
StoreNode05.SLHOST4	back	32,889	272	8,945,808	8,000,000	12,000,000	770,567
StoreNode05n.SLHOST1	back	39,918	272	10,857,696	8,000,000	12,000,000	1,650
StoreNode05n.SLHOST2	back	33,736	272	9,176,192	8,000,000	12,000,000	200,245
StoreNode05n.SLHOST4	back	31,224	272	8,492,928	8,000,000	12,000,000	764,372
StoreNode08.SLHOST2	back	35,390	272	9,626,080	8,000,000	12,000,000	151,865
StoreNode08.SLHOST4	back	25,129	272	6,835,088	8,000,000	12,000,000	695,010

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.
- 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.

- 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

Storage Manager Detail

05-Jul-2016 12:16 Data OK + ?

Service: DistributedCache Cache: foo

(Select a Node 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

⏪ ||| ⏩

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.

- 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	<p>Location A unique identifier for each node. It is defined as member_name.machine.rack.site.</p> <p>EventsDispatched The total number of events dispatched by the Storage Manager since the last time the statistics were reset.</p> <p>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.</p> <p>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`.</p> <p>ListenerFilterCount The number of filter-based listeners currently registered with the Storage Manager.</p> <p>ListenerKeyCount The number of key-based listeners currently registered with the Storage Manager.</p> <p>ListenerRegistrations The total number of listener registration requests processed by the Storage Manager since the last time the statistics were reset.</p> <p>LocksGranted The number of locks currently granted for the portion of the partitioned cache managed by the Storage Manager.</p> <p>LocksPending The number of pending lock requests for the portion of the partitioned cache managed by the Storage Manager.</p> <p>RemoveCount The number of removes from the backing map managed by this Storage Manager caused by operations such as clear, remove or invoke.</p>

Storage IndexInfo View

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



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:

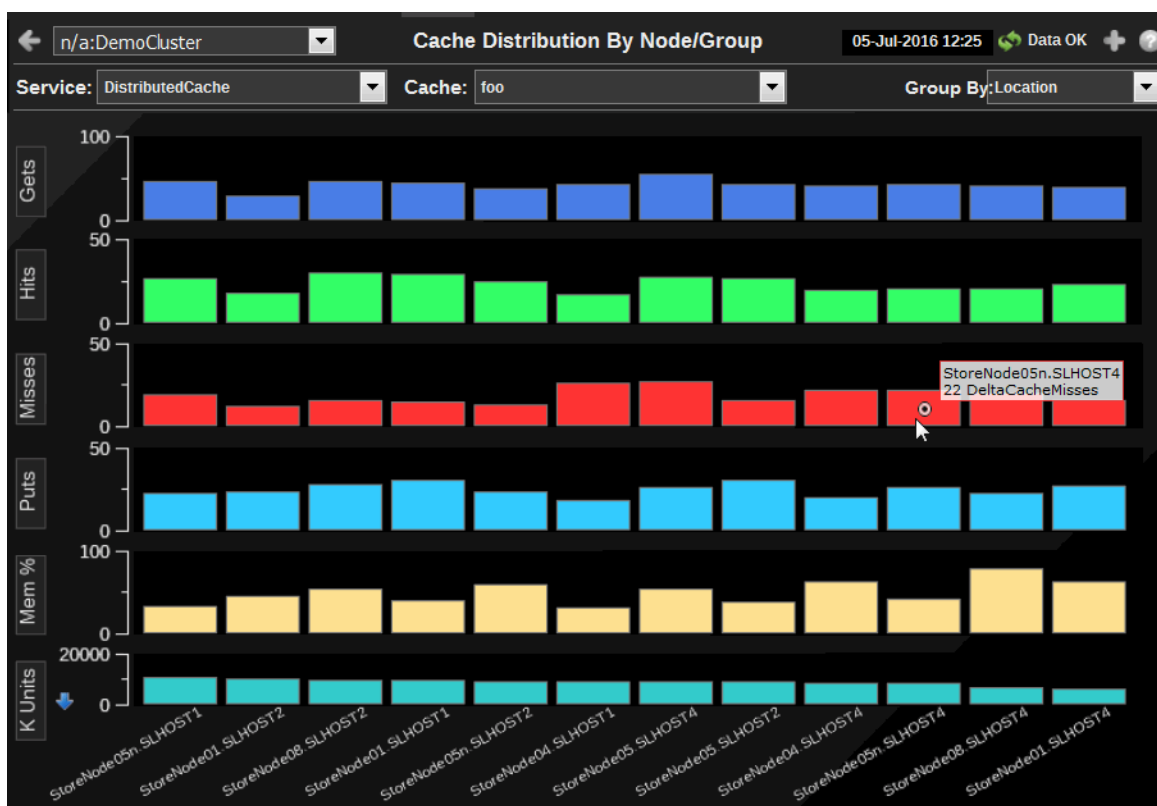
Extractor = the index name.


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

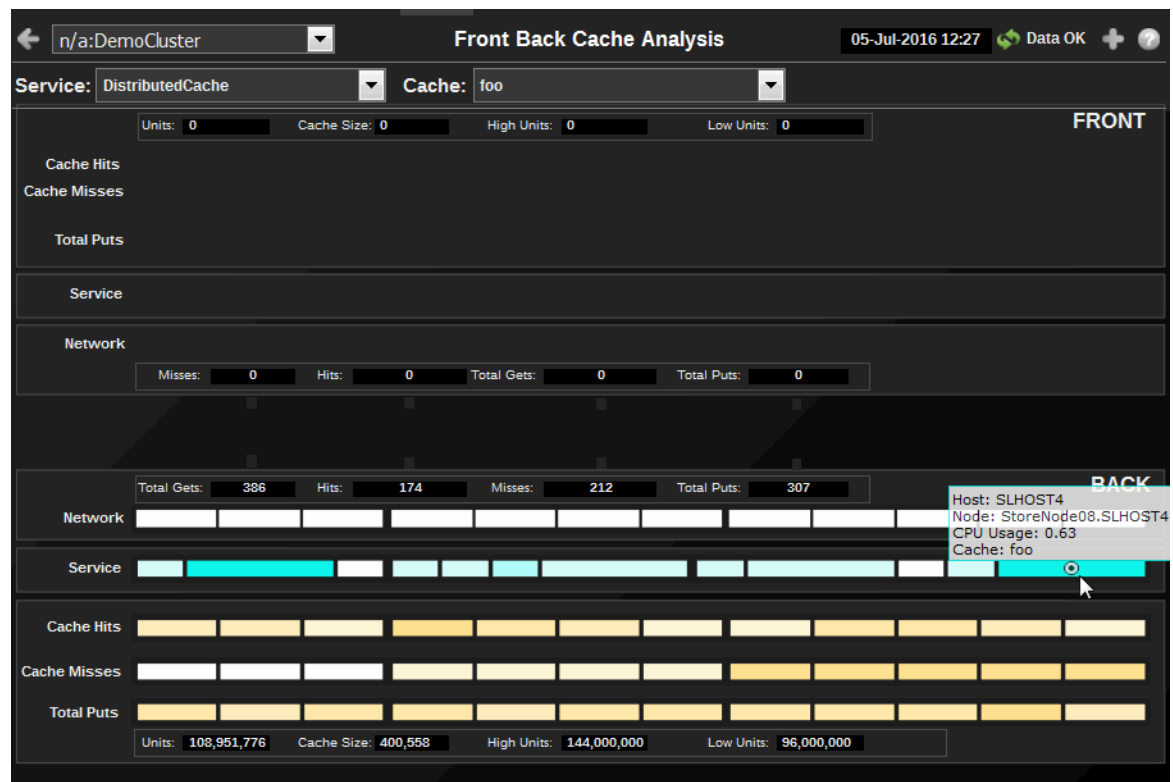
This display presents the distribution of cache activity across all storage nodes in the cluster. The buttons on the left may be used to select the metric by which all six bar charts are to be sorted. Note that the Gets, Hits, Misses, and Puts are shown in the same color as on the other Cache Analysis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





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	<p>Select the node group by which the data are totaled.</p> <p>Location A unique identifier for each node, defined as member_name.machine.rack.site. This is the default setting.</p> <p>Gets Requests for data from this cache.</p> <p>Hits Number of successful gets.</p> <p>Misses Number of failed gets.</p> <p>Puts Data stores into this cache.</p> <p>Mem% Calculated percent of memory used divided by total memory.</p> <p>K Units Units in thousand bytes.</p>

Front/Back Analysis



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 [dropdown], 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.

- 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.

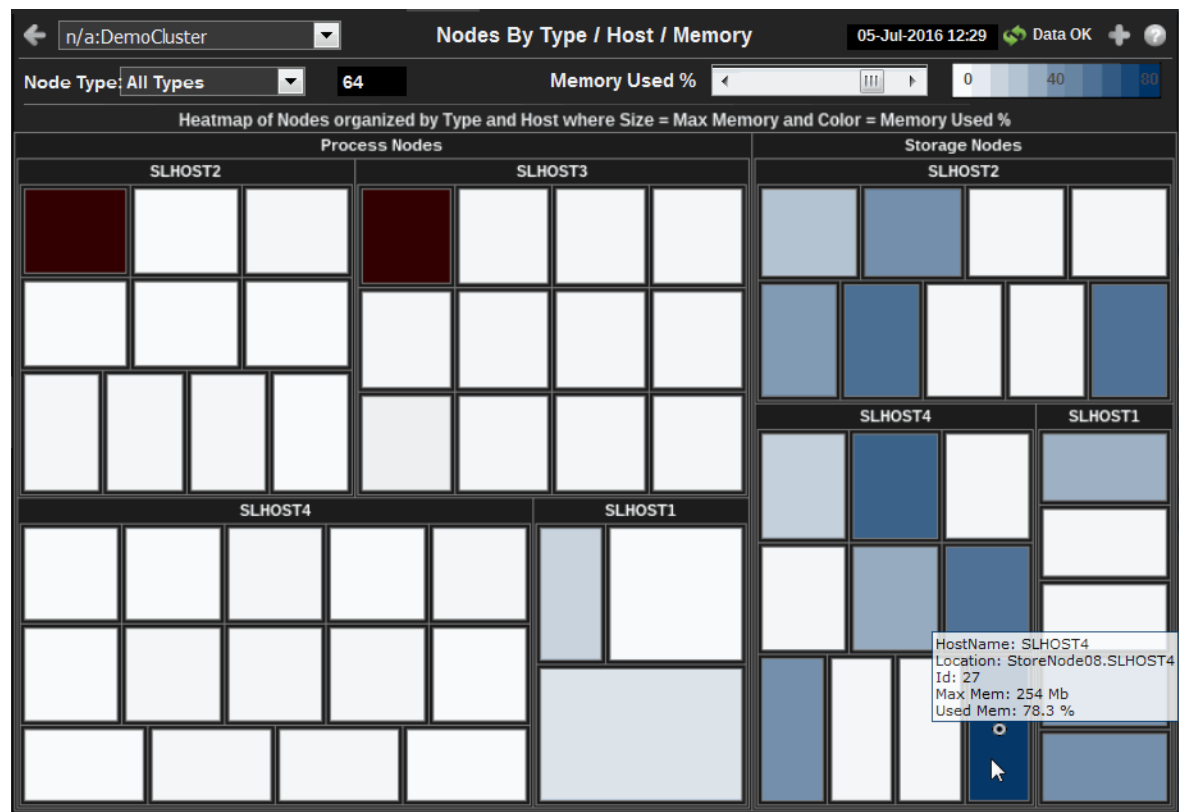
All Nodes

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 610](#): Heatmap of caches by service where size represents Max Memory and color represents percent of Memory Used.
- [“All Nodes CPU” on page 611](#): Heatmap shows CPU utilization for all nodes in the cluster.
- [“All Nodes Grid View” on page 612](#): Grid view showing information about all nodes.
- [“Communication Issues” on page 613](#): Bar chart displays current communication issues for all nodes.
- [“All Nodes - Detail” on page 615](#): Table shows current detailed statistics for all nodes.
- [“Invocation Service Detail” on page 617](#): 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.

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.

- 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):

- ← ↑ 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.

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.



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.

- 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 Nodes organized by Type/Host

A heatmap of CPU usage per host.

The following icon is shown for each node in the cluster:



The icon describes the node:


- Location (**StoreNode04.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.





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 Rate/Received Failure Rate by Node/Group	Packets failed to be sent by each node. 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.

Location	Id	Avail MB	Max MB	Pkts Sent	Delta	Pkts Rcvd	Delta	Pkts
StoreNode08.SLHOST4	27	56	254	20,680,021	1,881	19,769,956	1,759	13
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	4
StoreNode07.SLHOST2	7	249	254	2,954,996	811	2,652,154	693	3
StoreNode06.SLHOST4	15	245	254	14,413,751	1,289	13,550,004	1,187	2
StoreNode06.SLHOST2	6	246	254	4,998,465	1,519	4,697,421	1,386	2
StoreNode05n.SLHOST4	13	144	254	9,240,557	832	8,292,349	717	2
StoreNode05n.SLHOST2	4	106	254	3,215,830	818	2,917,402	698	2
StoreNode05n.SLHOST1	17	145	247	145,793	890	146,745	749	1
StoreNode05.SLHOST4	14	112	254	9,045,804	734	8,096,728	619	6
StoreNode05.SLHOST2	5	149	254	3,130,043	799	2,826,880	694	3
StoreNode04.SLHOST4	12	163	254	20,767,959	1,825	19,859,593	1,720	6
StoreNode04.SLHOST1	19	141	247	351,071	1,934	350,140	1,798	1
StoreNode03.SLHOST4	11	249	254	8,452,999	915	7,564,521	800	4
StoreNode03.SLHOST2	3	249	254	2,911,819	901	2,606,105	784	3
StoreNode03.SLHOST1	22	236	247	151,367	890	132,011	751	1
StoreNode02.SLHOST4	10	247	254	14,330,435	1,265	13,452,552	1,148	2
StoreNode02.SLHOST2	2	245	254	4,873,858	1,357	4,569,689	1,226	1
StoreNode02.SLHOST1	23	234	247	243,086	1,248	223,462	1,109	1
StoreNode01.SLHOST4	9	97	254	9,003,123	710	8,202,539	600	5
StoreNode01.SLHOST2	1	139	254	3,074,838	682	2,822,646	572	2
StoreNode01.SLHOST1	18	169	247	145,132	848	147,525	718	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.
- 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.

- 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-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.
- **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.

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	2	0	
ExtendNode01.SLHOST4	Management	✓	0.3	558,849	48	2	0	
ExtendNode02.SLHOST2	Management	✓	0.5	193,035	50	2	0	
ExtendNode02.SLHOST4	Management	✓	0.8	559,564	50	2	0	
JmxNode01.SLHOST1	Management	✓	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	✓	0.2	599,986	52	2	0	
OLS-StoreNode-1.SLHOST2	Management	✓	0.1	263,598	69	2	0	
OLS-StoreNode-1.SLHOST4	Management	✓	0.2	763,897	69	2	0	
ProcessNode01.SLHOST2	Management	✓	0.6	207,681	52	2	0	
ProcessNode01.SLHOST3	Management	✓	0.2	691,792	62	2	0	
ProcessNode01.SLHOST4	Management	✓	0.1	600,603	53	2	0	
ProcessNode02.SLHOST2	Management	✓	0.2	207,370	53	2	0	
ProcessNode02.SLHOST3	Management	✓	0	691,971	59	2	0	
ProcessNode02.SLHOST4	Management	✓	0.2	600,699	53	2	0	
ProcessNode03.SLHOST2	Management	✓	0.3	221,989	57	2	0	
ProcessNode03.SLHOST3	Management	✓	0.3	733,684	66	2	0	
ProcessNode03.SLHOST4	Management	✓	0.5	642,766	59	2	0	
ProcessNode04.SLHOST2	Management	✓	0.3	235,858	61	2	0	
ProcessNode04.SLHOST3	Management	✓	0.5	775,442	67	2	0	
ProcessNode04.SLHOST4	Management	✓	0.2	683,213	62	2	0	
ProcessNode05.SLHOST3	Management	✓	0.2	689,541	59	2	0	
ProcessNode05.SLHOST4	Management	✓	0.3	599,289	52	2	0	
ProcessNode05n.SLHOST2	Management	✓	0.5	235,663	61	2	0	
ProcessNode05n.SLHOST3	Management	✓	0.2	774,342	67	2	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.

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

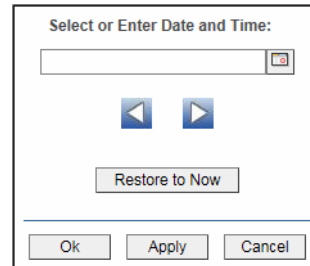
This display presents summary information about an individual node.








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 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 in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

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):


- 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.

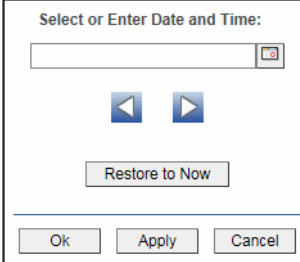
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

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.


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



The dialog box titled "Select or Enter Date and Time:" contains a text input field at the top with a calendar icon on its right. Below the field are two blue navigation arrows, a "Restore to Now" button, and at the bottom, three buttons: "Ok", "Apply", and "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   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 Location

Trend 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**.

CPU% CPU Utilization (as a percent) on the selected **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.

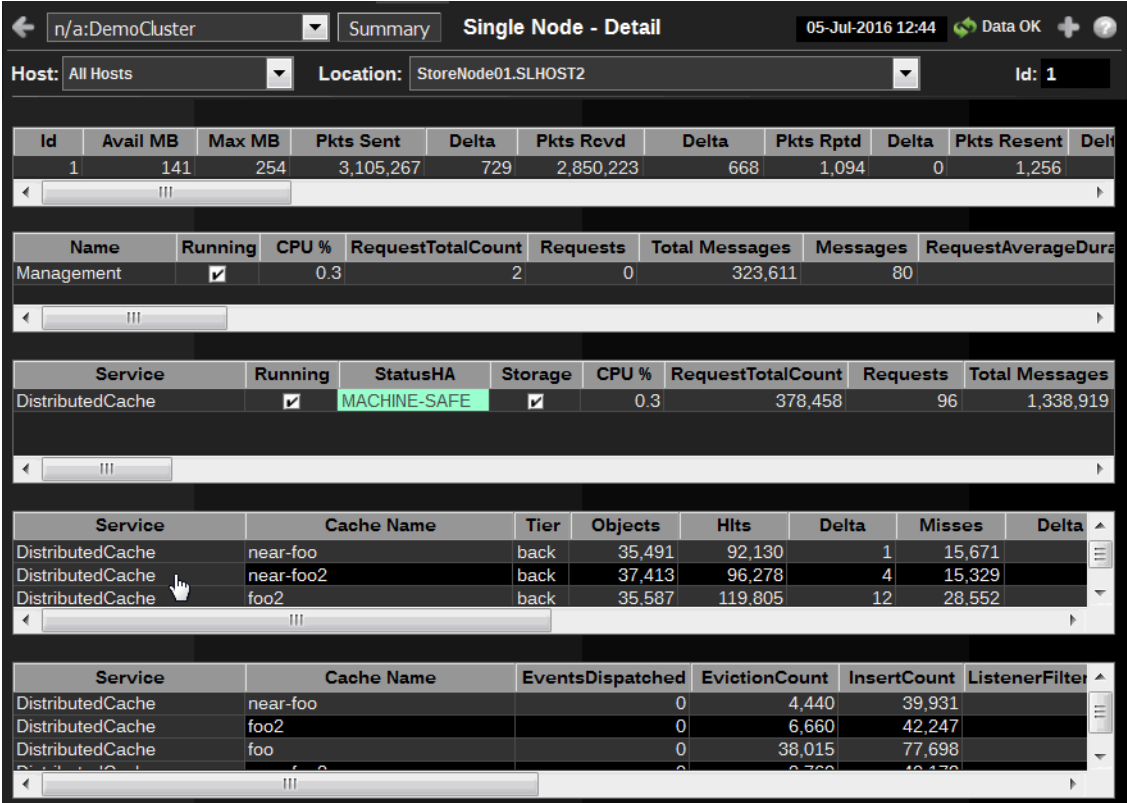
Request Pending Count Number of pending requests issued by the service.

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.



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.

- 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
Data**


This table contains data from the Invocation Services MBean for the selected node.


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**

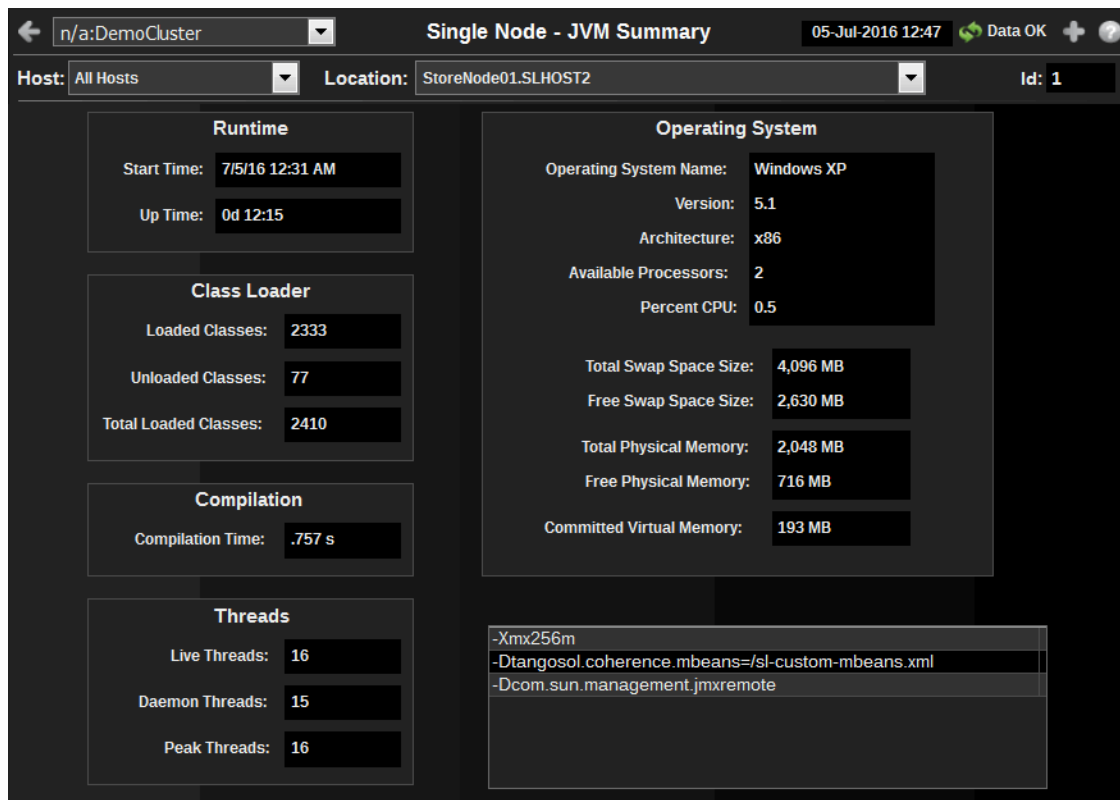
This table contains data from the Cache MBeans associated with the selected node.

**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.



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.

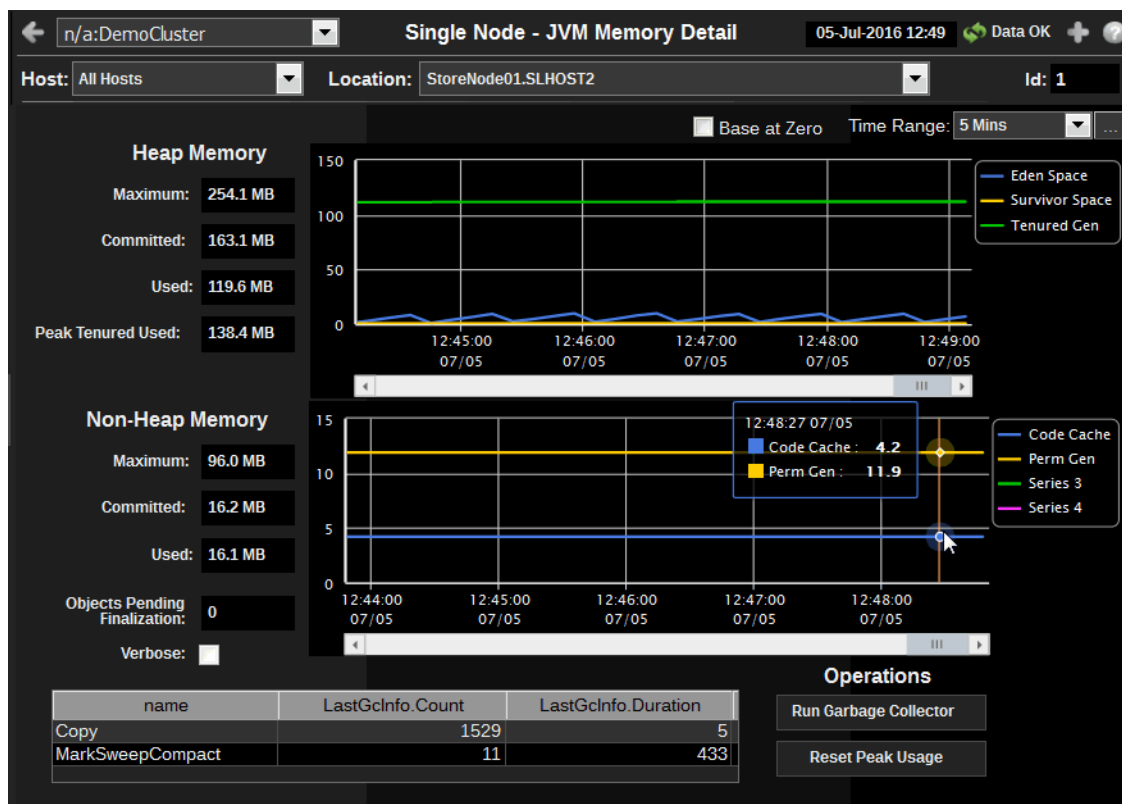
23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

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	<p>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.</p> <p>NOTE: Compilation Time monitoring may not be supported depending on the platform (for example, a Java virtual machine implementation).</p>
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
Input Arguments	Free Physical Memory The value of the OperatingSystemMXBean's FreePhysicalMemorySize attribute
	Committed Virtual Memory The value of the OperatingSystemMXBean's CommittedVirtualMemorySize attribute
	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/package-summary.html#package_description.



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.

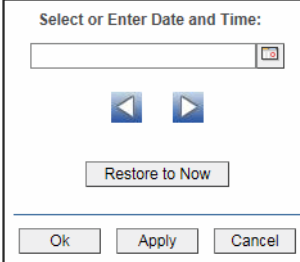
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

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.


- 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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field at the top with a calendar icon on its right. Below the field are two blue navigation arrows, a "Restore to Now" button, and at the bottom, three buttons: "Ok", "Apply", and "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   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 HeapMemoryUsage attribute.

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 Collection

name Name of the Garbage Collector MBean.


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.



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.


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 Range

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



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

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

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

Memory Usage (in MB) Before and After Garbage Collection

Max The maximum amount of memory used by the node or nodes.

Committed The amount of memory guaranteed to be available for use by the JVM.

Used - Before The amount of memory used by the node or nodes before garbage collection.

Used - After The amount of memory used by the node or nodes after 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.

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_57c1	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\Java	DemoCluste
java.home	C:\Program Files\Java\jre6	DemoCluste
java.io.tmpdir	C:\DOCUME~1\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

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.

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.version	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 633](#): Analyze service data for two sets of time ranges.
- [“Cache Comparison” on page 634](#): 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.

TRA - Service Comparison 05-Jul-2016 12:55 Data OK

Service: DistributedCache Storage Nodes Process Nodes

Time Range Analysis:

Time Range 1: Start: End:

Time Range 2: Start: End:

Time Range 1: Data Bucket Timestamp Start Jun 20, 2016 12:55:00 PM End Jul 5, 2016 12:54:22 PM

Time Range 2: Data Bucket Timestamp Start Jun 20, 2016 12:55:00 PM End Jul 5, 2016 12:54:22 PM

Metric Name	Time Range 1 Value	Time Range 2 Value	Percentage Change
DeltaMessages	444,734,253	444,734,253	0.00
DeltaRequestTotalCount	123,883,421	123,883,421	0.00
DeltaTaskCount	0	0	0.00
TaskBacklog	22,701	22,701	0.00

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.

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.

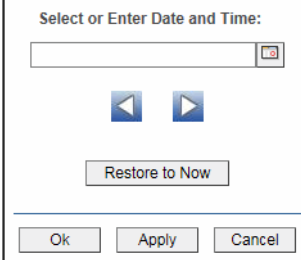
Process Nodes Select to display process node data in the trend graphs of this display.

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 .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right) and a "Restore to Now" button. At the bottom are "Ok", "Apply", and "Cancel" buttons.

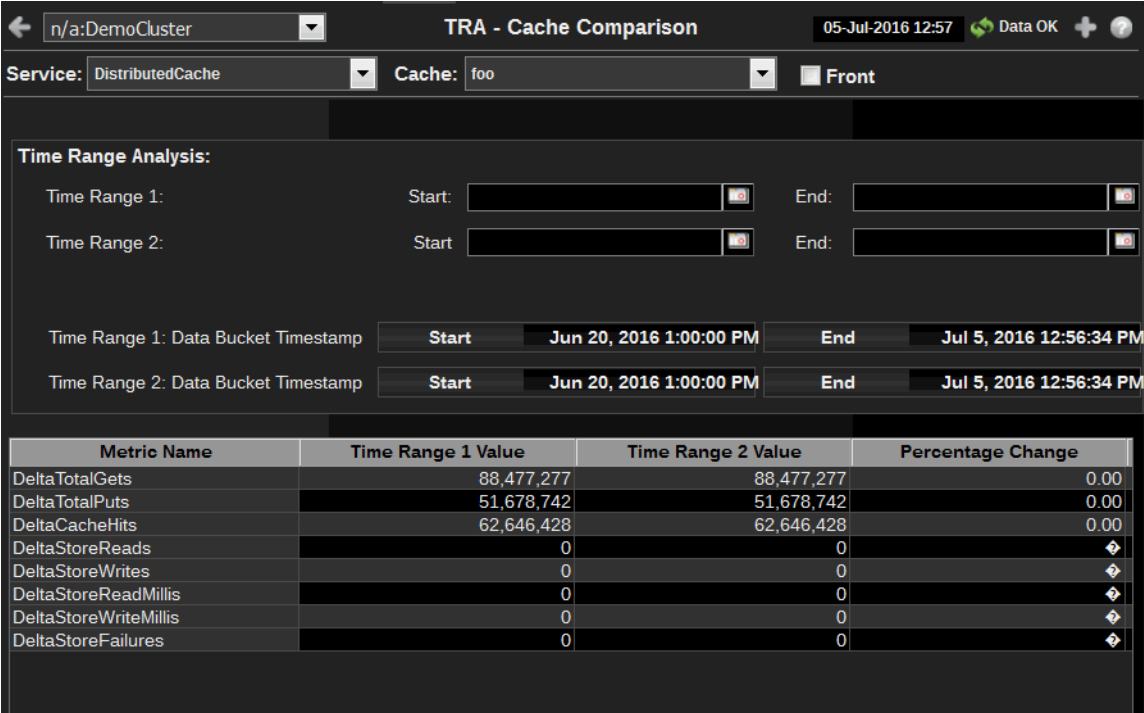
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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






Cache Comparison


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




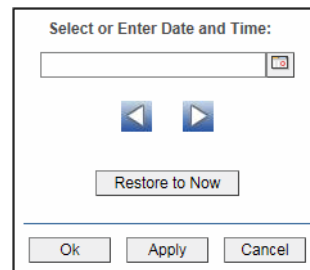
The screenshot shows the "TRA - Cache Comparison" window. At the top, it displays the cluster name "n/a:DemoCluster", the title "TRA - Cache Comparison", and the timestamp "05-Jul-2016 12:57" with a "Data OK" status. Below this, the "Service" is set to "DistributedCache" and the "Cache" is set to "foo". The "Front" checkbox is checked.


The "Time Range Analysis" section contains two rows for "Time Range 1" and "Time Range 2". Each row has "Start" and "End" fields with calendar icons. Below these, the "Data Bucket Timestamp" for each range is displayed, showing a start time of "Jun 20, 2016 1:00:00 PM" and an end time of "Jul 5, 2016 12:56:34 PM".



Metric Name	Time Range 1 Value	Time Range 2 Value	Percentage Change
DeltaTotalGets	88,477,277	88,477,277	0.00
DeltaTotalPuts	51,678,742	51,678,742	0.00
DeltaCacheHits	62,646,428	62,646,428	0.00
DeltaStoreReads	0	0	
DeltaStoreWrites	0	0	
DeltaStoreReadMillis	0	0	
DeltaStoreWriteMillis	0	0	
DeltaStoreFailures	0	0	



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	<p>Time Range 1: Set Start and End times for Time Range 1</p> <p>Time Range 2: Set Start and End times for Time Range 2</p> <p>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.</p>
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click Calendar  .



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

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

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


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 636](#): Monitor information on metrics acquisition. Permits user to reset system metrics.
- [“Cluster MBean Servers” on page 638](#): 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 640](#): Monitor information about Coherence JMX management settings.
- [“Node Administration” on page 642](#): Permits user to modify node parameters.
- [“Cache Administration” on page 644](#): 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	<p>Node MBeans Total number of node MBeans queried.</p> <p>Service MBeans Total number of service MBeans queried.</p> <p>Cache MBeans Total number of cache MBeans queried.</p> <p>Storage MBeans Total number of storage MBeans queried.</p> <p>JVM Platform MBeans Total number of JVM platform MBeans queried.</p> <p>Total Total number of MBeans queried.</p>
RTView Memory Utilization for Metrics History	<p>By default, the Oracle Coherence Monitor stores several hours of data using in-memory tables.</p> <p>Cache Metrics, # Rows The number of table rows used by the Monitor to store cache metrics data.</p> <p>Cache Metrics, Memory Used (KB) The amount of memory (KB) used by the Monitor to store cache metrics data.</p> <p>Node Metrics, # Rows The number of table rows used by the Monitor to store node metrics data.</p> <p>Node Metrics, Memory Used (KB) The amount of memory (KB) used by the Monitor to store node metrics data.</p>
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.

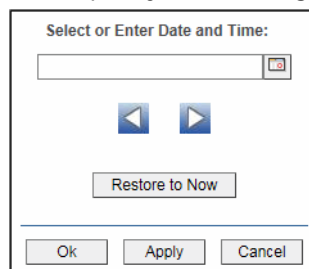


Total RTView JMX Query Time


Traces the total amount of time, in milliseconds, to query the monitoring MBeans from Coherence.



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 .



A dialog box titled "Select or Enter Date and Time:". It contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). At the bottom of the dialog is a "Restore to Now" button. At the very bottom are three buttons: "Ok", "Apply", and "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**.

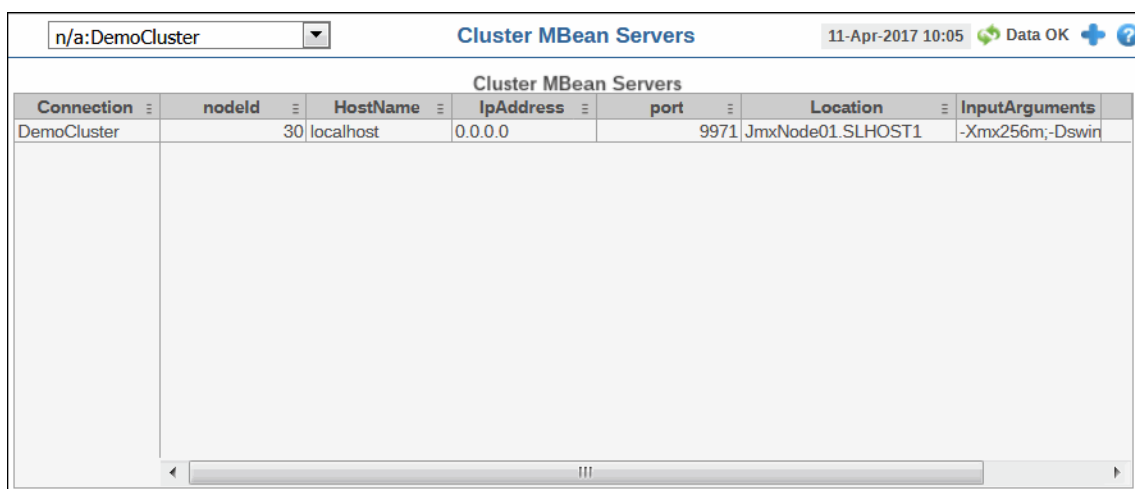
Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

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.



The screenshot shows a web application window titled "Cluster MBean Servers". At the top left is a dropdown menu showing "n/a: DemoCluster". At the top right is a timestamp "11-Apr-2017 10:05", a "Data OK" status with a green icon, and a help icon. Below the header is a table with the following data:

Connection	nodeId	HostName	IpAddress	port	Location	InputArguments
DemoCluster	30	localhost	0.0.0.0	9971	JmxNode01.SLHOST1	-Xmx256m;-Dswin


Below the table is a large empty rectangular area, likely for a detailed view or logs. At the bottom of the window is a horizontal scrollbar.



Cluster	Select a cluster to display.
Connection	The name of the JMX connection used to access the cluster data.
nodeId	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.



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.

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.

- Cluster

Select a cluster to display.
- JMX Connection

The name of the JMX connection used to access the cluster data.

JMX Management Settings**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.

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.

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 Settings for all Nodes
(Select a Node to change settings)

Location	BufferPublishSize	BufferReceiveSize	BurstCount	BurstDelay	MulticastThreshold
ExtendNode01.SLHOST4	32	1428	0	10	25
ExtendNode02.SLHOST2	32	1428	0	10	25
ExtendNode02.SLHOST4	32	1428	0	10	25
JmxNode01.SLHOST1	32	1428	0	10	25
OLS-ProcessNode-1.SLHOST2	32	1428	0	10	25
OLS-ProcessNode-1.SLHOST4	32	1428	0	10	25
OLS-StoreNode-1.SLHOST2	32	1428	0	10	25
OLS-StoreNode-1.SLHOST4	32	1428	0	10	25
ProcessNode01.SLHOST2	32	1428	0	10	25
ProcessNode01.SLHOST3	32	1428	0	10	25
ProcessNode01.SLHOST4	32	1428	0	10	25

Selected Node(s): All

Select All Nodes

Burst Count: 0

Burst Delay: 0

Logging Level: 0

Logging Limit: 0

Logging Format:

Shutdown Node

Ensure Cache Service: DistributedCache

Ensure Invocation Service:

Multicast Threshold: 0

Resend Delay: 0

Send Ack Delay: 0

Traffic Jam Count: 0

Traffic Jam Delay: 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.

- Cluster

Select a cluster to display.
- JMX Connection

The name of the JMX connection used to access the cluster data.

**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 than 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.

nodeId 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 not 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.

Location	LowUnits	HighUnits	BatchFactor	ExpiryDelay	FlushDelay	Queue
StoreNode01.SLHOST1	8,000,000	12,000,000	0.0	0	0	0
StoreNode01.SLHOST2	8,000,000	12,000,000	0.0	0	0	0
StoreNode01.SLHOST4	8,000,000	12,000,000	0.0	0	0	0
StoreNode04.SLHOST1	8,000,000	12,000,000	0.0	0	0	0
StoreNode04.SLHOST4	8,000,000	12,000,000	0.0	0	0	0
StoreNode05.SLHOST2	8,000,000	12,000,000	0.0	0	0	0
StoreNode05.SLHOST4	8,000,000	12,000,000	0.0	0	0	0
StoreNode05n.SLHOST1	8,000,000	12,000,000	0.0	0	0	0
StoreNode05n.SLHOST2	8,000,000	12,000,000	0.0	0	0	0
StoreNode05n.SLHOST4	8,000,000	12,000,000	0.0	0	0	0

Selected Node(s): All Select All Nodes

High Units: Batch Factor:
 Low Units: Expiry Delay:
 Flush Delay:
 Queue Delay:
 Refresh Factor:
 Requeue Threshold:

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.
- 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.

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 Selected Cache	<p>Location A unique identifier for each node. It is defined as: member_name.machine.rack.site.</p> <p>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.</p> <p>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.</p> <p>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].</p> <p>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.</p> <p>FlushDelay The number of milliseconds between cache flushes. Value of zero indicates that the cache will never flush.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>nodeId The node ID.</p>
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.

RTView Cache Tables 05-Jul-2016 13:09 Data OK + ?

DataServer: <Default> RTView Cache Tables Max Rows: 4000 History Tables

CacheTable	TableType	Rows	Columns	Memory
JmxStatsTotals	current	1	4	441
OcBadCommunicationNodes	current	140	6	14,999
OcCacheServiceStats	current	88	58	62,666
OcCacheServiceTotals	current	8	26	4,441
OcCacheStats	current	172	80	206,148
OcCacheTotals	current	17	52	13,406
OcClusterOverview	current	1	7	791
OcClusterStats	current	1	19	14,103
OcExtendConnections	current	112	30	68,304
OcInvocationServiceStats	current	63	60	62,252
OcInvocationServiceTotals	current	1	26	2,841
OcJmxConnection	current	2	7	1,254
OcJmxHostData	current	1	15	1,754

OcInvocationServiceStats Rows: 63

TIME_STAMP	BackupCou	OwnedParti	OwnedParti	RefreshTim	RequestAve	RequestMa	RequestPei	RequestPei	RequestPei
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	85.5	109	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	7.5	15	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	16	16	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	15.5	16	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	24	32	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	31.5	47	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	39.5	63	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	16	16	0	0	0
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	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 [v], 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.

About

This display shows details about the Solution Package version and data sources available to your system.



Oracle Database

Use the Solution Package for Oracle® Database to monitor the health of your Oracle databases. The following Oracle Database Views can be found under **Components** tab > **Databases** > **Oracle Database**:

- ["Database Instances View"](#)
- ["Database Details View"](#)

Database Instances View

These displays present performance metrics and alert statuses for all Oracle databases and instances. Displays are:

- [“All Databases Heatmap”](#): Heatmap shows alert status for all Oracle databases.
- [“All Databases Table”](#): List of all Oracle databases with detailed utilization metrics.
- [“Database Summary”](#): Detailed utilization metrics and configurations for a single Oracle database.
- [“All Instances Heatmap”](#): Heatmap shows alert status of all instances on a single Oracle database.
- [“All Instances Table”](#): List of all instances on a single Oracle database with detailed utilization metrics.
- [“Instance Summary”](#): Detailed utilization metrics and configurations for a single Oracle instance.

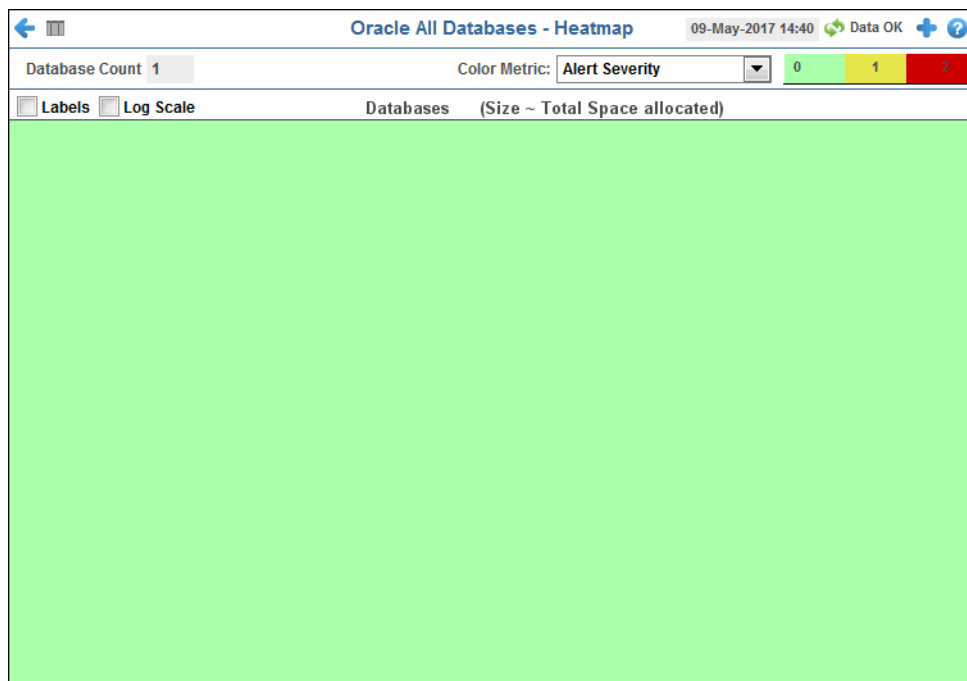
All Databases Heatmap


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

Each heatmap rectangle represents a different database. The rectangle color indicates the most critical alert state for the selected metric.

Choose a **Color Metric** from the drop-down menu: **Alert Severity**, **Alert Count**, **Response Time** or **Database Space Usage**. Use the check-boxes ☒ to include labels and apply log scale in the heatmap. Move your mouse over a rectangle to see additional information. Toggle between the commonly accessed **Table** and this **Heatmap** display by clicking the icon in the upper left-hand corner.

Investigate a database by clicking a heatmap rectangle to view details in the [“Database Summary”](#) display.





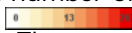






Database Count The number of databases in the display.

Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps current relative values to colors:

Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Response Time	<p>The amount of time, in milliseconds, since the last execution of the database. The color gradient  bar numerical values range from 0 to the maximum amount of time in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Database Space Usage	<p>The amount of space used, in megabytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum amount of space used in the heatmap. The middle value in the gradient bar indicates the average amount.</p>

Labels Select to include labels in the heatmap.

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

All Databases Table

View a list of all your Oracle databases with data from the [“All Databases Heatmap”](#) display (**Alert Severity**, **Alert Count**, **Response Time**, **Database Space Usage**) in a tabular format. Each row in the table is a different Oracle database. All values refer to the database except where noted.

Light red rows indicate the database is expired or not connected.

The **Severity** column indicates the most critical alert on the database, where:

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

Investigate a database by clicking a row to view details in the “Database Summary” display.

All Oracle Databases - Table09-May-2017 14:42Data OK

Database Count: 1

All Oracle Databases												
DBNAME	Expired	Severity	Alert Count	Instance Count	Response Time (ms)	% Used Space	Used Space	Free Space	Total Space	Temp Space	Log Space	
orclStandard11gR2	<input type="checkbox"/>	●	0	1	63	85	2,863	535	3,385	20		

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.

Database Count The number of databases in the display.

All Oracle Databases Table

Each row is a different database.

- DBName**

The database name.
- Expired**

When checked, the database is expired due to a connection issue.
- Severity**

The maximum level of alerts.
 - Red indicates that one or more metrics have reached their alarm threshold.
 - Yellow indicates that one or more metrics have reached their alarm threshold.
 - Green indicates that no metrics have reached their alert thresholds.
- Alert Count**

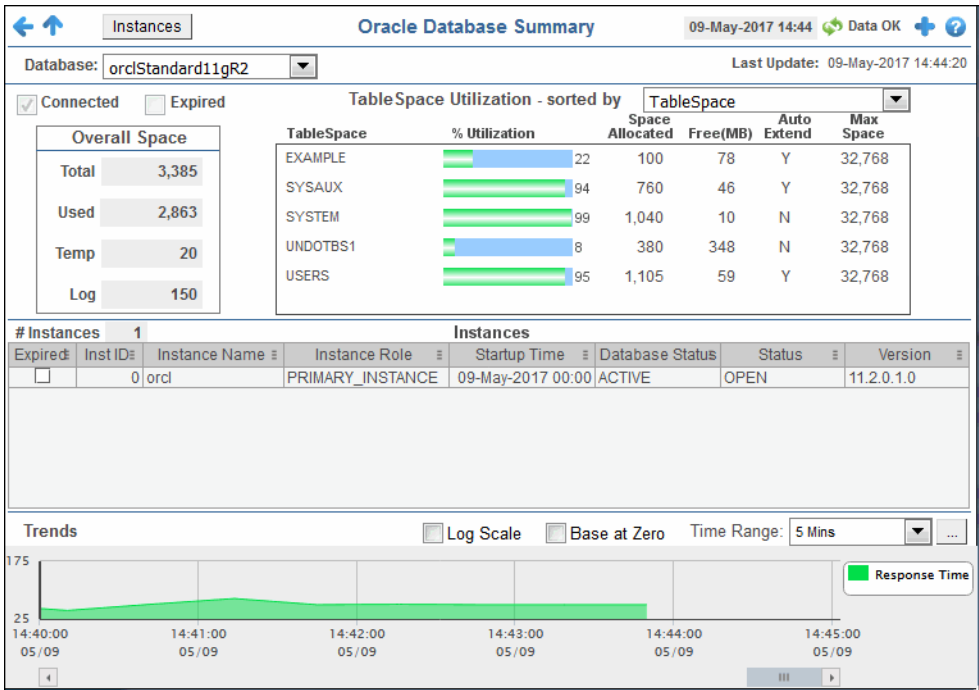
The total number of critical and warning alerts.

Instance Count	The total number of Oracle Database instances.
Response Time (ms)	The amount of time, in milliseconds, since the last execution of the database.
% Used Space	The percent space used.
Used Space	The space used by the database, in megabytes.
Free Space	The amount of available space, in megabytes.
Total Space	The total amount of used and available space, in megabytes.
Temp Space	The amount of temporary space, in megabytes.
Log Space	The amount of log space, in megabytes.
Timestamp	The date and time of the last data update.

Database Summary

View detailed performance data for a single Oracle Database, including tablespace utilization, instances on the database and response time trends. Use this display to investigate database health.

Choose a **Database** and select one metric from the **TableSpace Utilization** drop-down menu to order the tablespace rows by that metric. The order is ascending. Mouse-over the trend graph to see additional information. Investigate an instance by clicking a row in the **Instances** table to view details in the “[Instance Summary](#)” display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.

Menu, Table open commonly accessed displays.

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.

- Database**
- Choose a database to display.
- Connected**
- When checked, the database is connected.
- Expired**
- When checked, the database is expired due to a connection issue.
- Overall Space**
- Values refer to the selected database.
- Total**
- The total amount of space, in megabytes.
- Used**
- The amount of space used, in megabytes.
- Temp**
- The total amount of temporary space, in megabytes.
- Log**
- The total amount of log space, in megabytes.

TableSpace Utilization

Choose a tablespace utilization metric to show data for (in the box below the drop-down menu). Values refer to the tablespaces on the selected database:

- **TableSpace:** Tablespace names in alphabetical order.
- **% Utilization:** The percent utilization of each tablespace in ascending order.
- **Allocated:** The allocated space for each tablespace in ascending order.

TableSpace	The tablespace name.
% Utilization	The percent utilization.
Space Allocated	The amount of allocated space, in megabytes.
Free(MB)	The amount of free space, in megabytes.
Auto Extend	Indicates whether auto extend is enabled (Y/N).
Max Space	The maximum amount of allocated space for the tablespace, in megabytes.

Instances

Each table row is an instance on the selected database. Click an instance to see details in the "[Instance Summary](#)" display.

Expired	When checked, the instance is expired due to a connection issue.
Inst ID	The unique identifier for the instance.
Instance Name	The name of the instance.
Instance Role	The role of the instance.
Startup Time	The date and time the instance began.
Database Status	The status of the database this instance is on.
Status	The status of this instance.
Version	The version number.

Trends


Traces response time of the selected database.

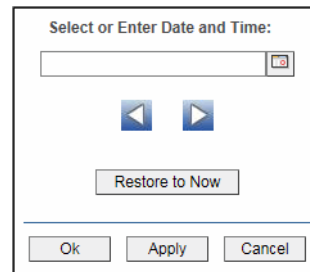
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log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


Base at Zero



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

Time Range

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



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

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

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

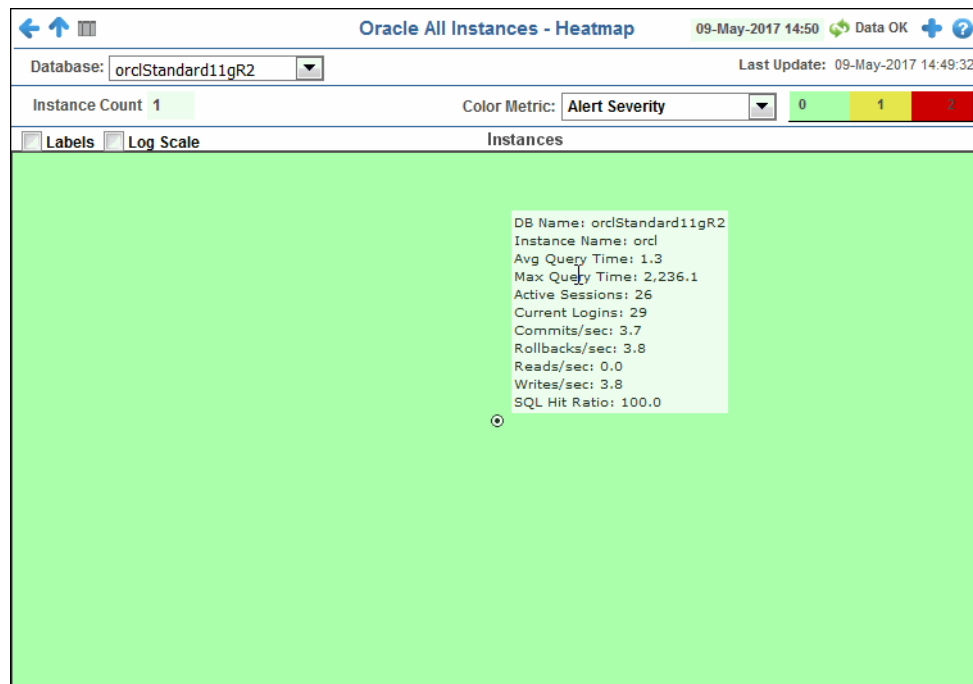
All Instances Heatmap

View status and alerts of all instances on an Oracle Database. Use this display to quickly identify instances with performance or utilization issues.

Each heatmap rectangle represents a different instance. The rectangle color indicates the most critical alert state for the selected metric.

Choose a **Color Metric** from the drop-down menu, such as **Alert Severity**, **Active Sessions** and **Disk Reads**. Use the check-boxes ☒ to include or exclude labels and use log scale in the heatmap. Move your mouse over a rectangle to see additional information. Toggle between the commonly accessed **Table** and this **Heatmap** display by clicking the icon in the upper left-hand corner. Mouse-over rectangles to view more instance performance details.

Investigate an instance by clicking a heatmap rectangle to view details in the “[Instance Summary](#)” display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Database Choose a database to display.

Instance Count The number of instances in the display.

Metric:

Choose the type of metric to show in the heatmap. Each rectangle is an instance. Each metric has its own gradient bar that maps current relative values to colors:

Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of 0.
Current Logins	The number of users logged on. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.
Active Sessions	The number of active sessions. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.
Avg Query Time	The average amount of time, in seconds, to perform a query. The color gradient  bar values range from 0 to the maximum average in the heatmap. The middle value in the gradient bar indicates the average amount.
Max Query Time	The maximum amount of time, in seconds, to perform a query. The color gradient  bar values range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average maximum amount.
Latch Hit Ratio	The ratio of the number of latch misses to the number of latch gets. The color gradient  bar values range from the lowest count to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
Data Dict Hit Ratio	The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.
SQL Hit Ratio	The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.
Commits/sec	The number of commits per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
Rollbacks/sec	The number of rollbacks per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
Disk Reads/sec	The number of disk reads per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
Disk Writes/sec	The number of disk writes per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
Labels	Select to include labels in the heatmap.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

All Instances Table

View data from the “[All Instances Heatmap](#)” display in a tabular format, as well as additional metrics and configuration information about instances.

Choose a database from the drop-down menu. Each row in the table is a different Oracle Database instance on the selected database. Investigate an instance by clicking a row to view details in the “[Instance Summary](#)” display.

Light red rows indicate the following:

- The database has an issue (the **Database Status** value is not **ACTIVE** or the **Active State** value is not **NORMAL**))
- The instance has an issue (the instance expired or the **Instance Status** value is not **OPEN**)

DBNAME	Instance ID	Instance Name	Expired	Severity	Host Name	Instance Role	Database Status	Instance Status
orclStandard11gR2	0	orcl	<input type="checkbox"/>		WIN44	PRIMARY_INSTANCE	ACTIVE	OPEN

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.
- 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.




Database Choose a database to display.

Instance Count The number of instances in the display.

Oracle Databases Instances Table

Each row is a different instance. Column values refer to the instance except where noted.

This data is obtained from the vendor. See vendor documentation for details.

DBName	The name of the database the instance is using.
Instance ID	The unique identifier for the instance.
Instance Name	The instance name.
Expired	When checked, the database is expired due to a connection issue.
Severity	<p>The maximum level of alerts.</p> <p> Red indicates that one or more metrics have reached their alarm threshold.</p> <p> Yellow indicates that one or more metrics have reached their alarm threshold.</p> <p> Green indicates that no metrics have reached their alert thresholds.</p>
Host Name	The host for the instance.
Instance Role	<p>The instance role. Values are:</p> <ul style="list-style-type: none"> • PRIMARY_INSTANCE • SECONDARY_INSTANCE • UNKNOWN
Database Status	<p>The current database status. Values are:</p> <ul style="list-style-type: none"> • ACTIVE • SUSPENDED • INSTANCE_RECOVERY
Instance Status	<p>The current state of the instance. Values are:</p> <ul style="list-style-type: none"> • STARTED • MOUNTED • OPEN • OPEN MIGRATE
Active State	<ul style="list-style-type: none"> • The database state when active. Values are: • NORMAL • QUIESCING • QUIESCED
Archiver	This data is obtained from the vendor. See vendor documentation for details.
Blocked	Is the instance in block mode? Yes/No
Parallel	Is the instance in parallel mode? Yes/No
Shutdown Pending	Is a shutdown pending? Yes/No
Start Time	The date and time the instance started.
Edition	The edition number of the instance.
Version	The instance software version.

Host OS	The host operating system.
CPU Count	The database CPU count.
Avg Query Time	The average amount of time, in seconds, to perform a query.
Min Query Time	The minimum amount of time, in seconds, to perform a query.
Max Query Time	The maximum amount of time, in seconds, to perform a query.
Buffer Hit Ratio	The ratio of the number of latch misses to the number of latch gets.
DD Hit Ratio	The ratio of logical reads to physical disk reads.
SQL Hit Ratio	The ratio of logical SQL reads to physical disk reads.
Logins	This data is obtained from the vendor. See vendor documentation for details.
Current Logins	The number of users currently logged in.
Cum Logins	The total number of logged in users since the instance started.
Logins Highwater	The login limit for the database.
Active Sessions	The number of logins that are currently active.
Inactive Sessions	The number of logins that are currently inactive.
Other Sessions	This data is obtained from the vendor. See vendor documentation for details.
Block Gets	The number of gets in block mode.
Consistent Gets	The number of gets in consistent mode.
Physical Get Reads	The number of physical gets reads.
Latch Hit%	The amount of latch gets, in percent.
Latch Miss%	The amount of latch misses, in percent.
Latch Spin%	This data is obtained from the vendor. See vendor documentation for details.
Parse Time%	This data is obtained from the vendor. See vendor documentation for details.
Recursive%	This data is obtained from the vendor. See vendor documentation for details.
Other%	This data is obtained from the vendor. See vendor documentation for details.
Deadlocks	The number of deadlocks since the database started.
Commits/sec	The number of commits per second.
Rollbacks/sec	The number of rollbacks per second.
Block Reads/sec	The number of block reads per second.

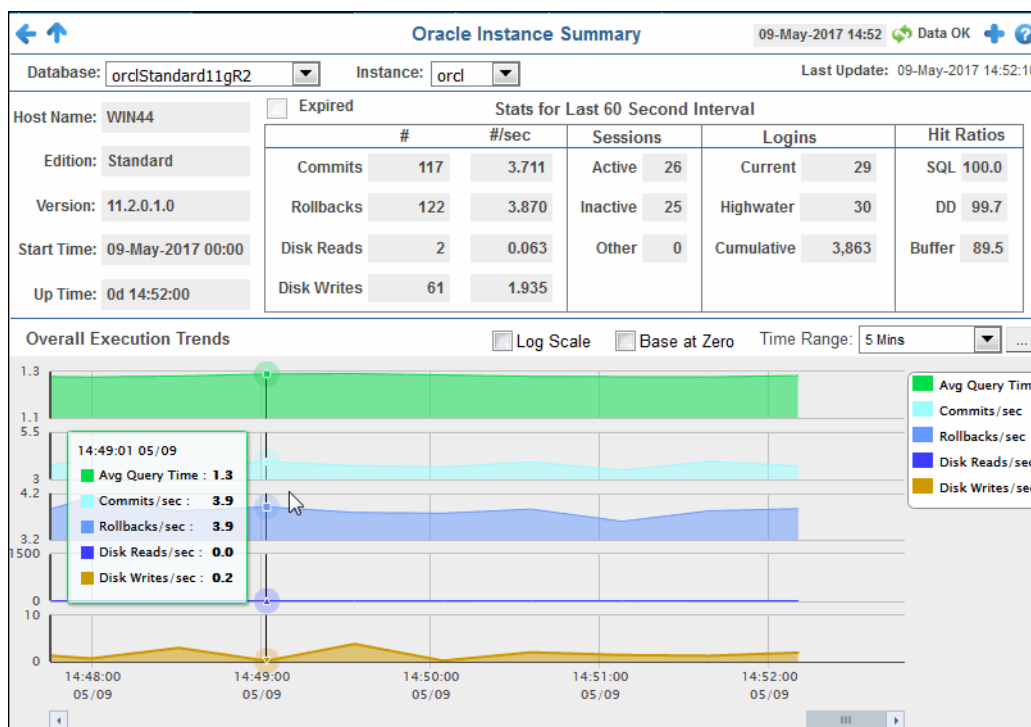
Block Writes/sec The number of block writes per second.

Timestamp The data and time of the last data update.

Instance Summary

View detailed data for a single Oracle Database instance, such as performance statistics for last 60 seconds, overall usage and execution trends. Use this display to closely investigate the health of an instance.

Choose a **Database** and **Instance** metric from the drop-down menus. Mouse-over the trend graph to see additional information.



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.

Database Choose a database to display.

Instance Choose an instance to display.

Host Name The name of the host.

Edition This data is obtained from the vendor. See vendor documentation for details.

Version	This data is obtained from the vendor. See vendor documentation for details.
Start Time	The date and time the instance started.
Up Time	The instance session duration.
Expired	When checked, the database is expired due to a connection issue.

Stats for Last 60 Second Interval

Values refer to the selected instance.

	#	#/sec
Commits	The number of commits since the instance started.	The current number of commits per second.
Rollbacks	The number of rollbacks since the instance started.	The current number of rollbacks per second.
Disk Reads	The number of disk reads since the instance started.	The current number of disk reads per second.
Disk Writes	The number of disk writes since the instance started.	The current number of disk writes per second.
Sessions	Active	The current number of active sessions on the instance.
	Inactive	The current number of inactive sessions on the instance.
	Other	This data is obtained from the vendor. See vendor documentation for details.
Logins	Current	The current number of logins on the instance.
	Highwater	The login limit for the instance.
	Cumulative	This data is obtained from the vendor. See vendor documentation for details.
Hit Ratios	SQL	The ratio of logical SQL reads to physical disk reads.
	DD	The ratio of logical reads to physical disk reads.
	Buffer	The ratio of the number of latch misses to the number of latch gets.

Overall Execution Trends


Traces the following for the selected instance:

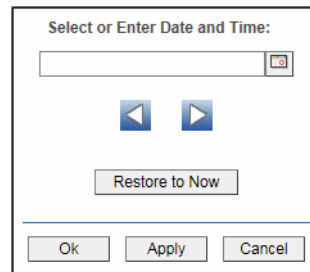
Avg Query Time: The average amount of time to perform a query, in seconds.


- **Commits/sec:** The number of commits per second.
- **Rollbacks/sec:** The number of rollbacks per second.
- **Disk Reads/sec:** The number of disk reads per second.
- **Disk Writes/sec:** The number of disk writes per second.



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

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

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



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

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

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

Database Details View

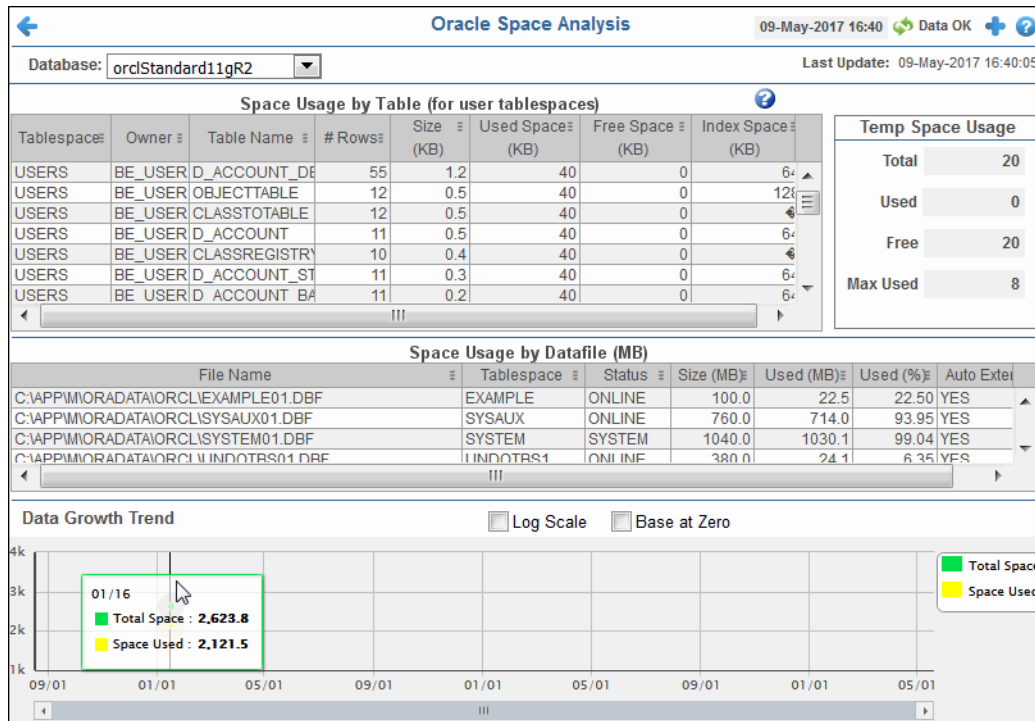
These displays present detailed database utilization metrics for an Oracle database. Displays are:

- **"Space"**: Shows detailed space utilization metrics for a single database.
- **"Contention"**: Shows detailed contention data for a single database.
- **"Processes"**: Shows detailed CPU and memory utilization per process for a single database.
- **"Sessions"**: Shows detailed CPU and memory utilization per session and user on a single database.
- **"Queries"**: Shows detailed query performance data such as disk utilization, execution and wait times and full SQL texts.
- **"Reports"**: Shows the following detailed reports: Audited Objects, Audited Privileges, Database Initialization, Feature Usage and Invalid Objects.

Space

Analyse tablespace utilization and trends at the data file level. Choose a database from the drop-down menu to view table and data file utilization data in a tabular format. Each row in the upper table is a different tablespace on the selected database. Each row in the lower table is a different data file on the selected database.

Mouseover the trend graph to see metrics for a specific day or time.



Title Bar (possible features are):

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

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.

Database Choose a database to display.

Space Usage by Table (for user tablespaces)

Each row is a different table on the selected database. Column values refer to the table.

Tablespace	The tablespace in which the table resides.
Owner	The table owner.
Table	The table name.
#Rows	The current number of rows in the table.
Size (KB)	The table size, in kilobytes.
Used Space (KB)	The amount of space used by the table, in kilobytes.
Free Space (KB)	The amount of free space for the table, in kilobytes.

Index Space (KB)	The amount of index space for the table, in kilobytes.
-------------------------	--

Temp Space Usage

Total	The total amount of temp space for the table, in kilobytes.
--------------	---

Used	The amount of used space for the table, in kilobytes.
-------------	---

Free	The amount of free temp space for the table, in kilobytes.
-------------	--

Max Used	The maximum amount of space used by the table, in kilobytes.
-----------------	--

Space Usage by Datafile (MB)

File Name	The name of the data file.
------------------	----------------------------

Tablespace	The name of the tablespace in which the data file resides.
-------------------	--

Status	The data file status:
---------------	-----------------------

- ONLINE
- SYSTEM

Size (MB)	The data file size, in megabytes.
------------------	-----------------------------------

Used (MB)	The amount of space used, in megabytes.
------------------	---

Auto Extend	Describes whether auto extend is configured (Yes/No).
--------------------	---

Data Growth Trend

Traces the following for the selected database:


- **Total Space:** The total amount of space for the table, in kilobytes.
- **Space Used:** The total amount of space used for the table, in kilobytes.


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


Base at Zero

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

Time Range

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



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

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

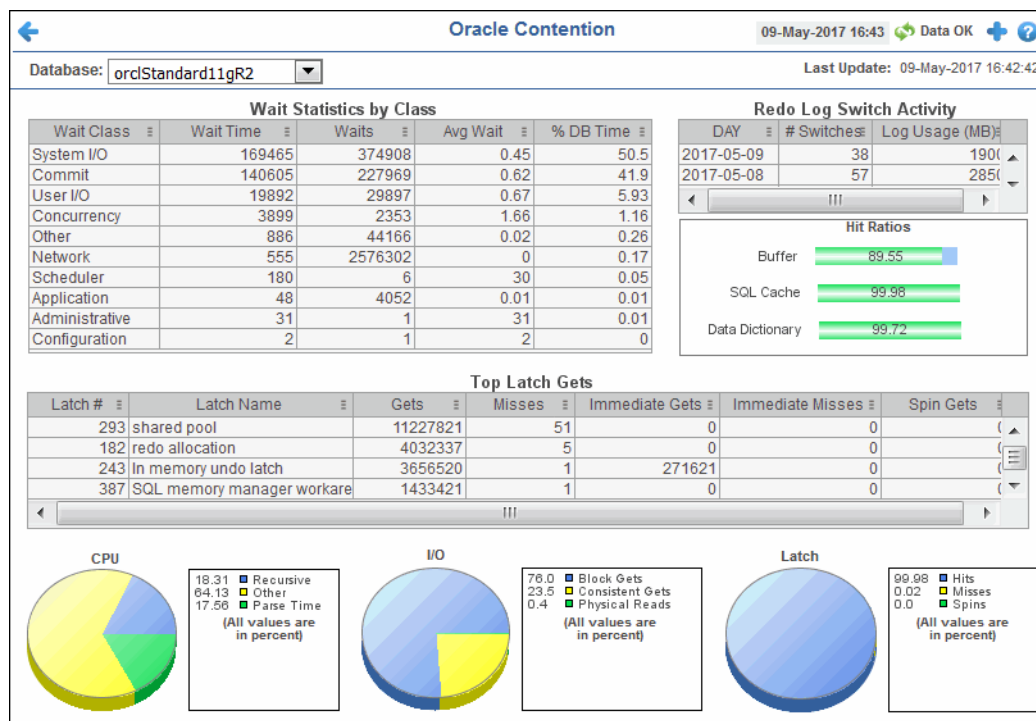
Click **Restore to Now** to reset the time range end point to the current time.

Contention

Analyse contention issues impacting database performance, including wait statistics per class, log switch activity, hit ratios and top latch gets.

Choose a database from the drop-down menu. Each row in the upper table is a different wait class on the selected database. Each row in the lower table is a different top latch on the selected database.

Mouseover the pie graph to see more information about database utilization.



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 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.

Database Choose a database to display.

Wait Statistics by Class

Each row is a different wait class on the selected database. Column values refer to the wait class.

Wait Class	The wait class name.
Wait Time	The amount of time, in milliseconds, for the class to process an event.
Waits	The number of wait events.
Avg Wait	The average amount of time, in seconds, for the class to process an event.
% DB Time	The amThis data is obtained from the vendor. See vendor documentation for details.

Redo Log Switch Activity

Each row is a different day. Column values refer to the selected database.

DAY	The date of the data update.
------------	------------------------------

#Switches	The number of switch processes performed by the database for the day.
Log Usage (MB)	This data is obtained from the vendor. See vendor documentation for details.

Hit Ratios


Values refer to the selected database.

Buffer	The ratio of the number of latch misses to the number of latch gets.
SQL Cache	The ratio of logical SQL reads to physical disk reads.
Data Dictionary	The ratio of logical reads to physical disk reads.

Top Latch Gets

Each row is a different latch on the selected database. Column values refer to the latch.

Latch #	The unique identifier for the latch.
Latch Name	The latch name.
Gets	The number of gets for the latch.
Misses	The number of misses for the latch.
Immediate Gets	The number of immediate gets for the latch.
Immediate Misses	The number of immediate misses for the latch.
Spin Gets	The number of spin gets for the latch.



CPU

Pie graph represents the selected database. Values are the percent utilization of **Recursive** calls, **Parse** time, and **Other** background processes (such as looking for buffers and fetching).

I/O

Pie graph represent the selected database. Values are the percent utilization of **Block Gets** calls, **Consistent Gets** and **Physical Reads**.

Latch

Pie graph represents the selected database. Values are the percent utilization of **Hits**, **Misses** and **Spins**.

Processes

Find out which persons and processes are using the most CPU on a database.

Choose a database from the drop-down menu. Each row in the upper table is a different user on the selected database. Each row in the lower table is a different process on the selected database.

PID	The unique process ID.
SID	The unique SQL process ID.
Serial #	This data is obtained from the vendor. See vendor documentation for details.
Physical Reads	The number of physical gets reads.
Block Gets	The number of gets in block mode.
Consistent Gets	The number of gets in consistent mode.

Top Sessions by CPU Utilization

Each row is a different session on the selected database. Column values refer to the session.

User Name	The name of the user.
SID	The unique SQL process ID.
CPU Usage (sec)	The amount of CPU used per second, in megabytes.
CPU Recursive (sec)	The amount of CPU used for recursive calls per second, in megabytes.
CPU Parsetime (sec)	The amount of CPU used for parsetime per second, in megabytes.

Sessions Holding Locks

Each row is a different lock on the selected database. Column values refer to the session.

User Name	The name of the user with the lock.
SID	The unique SQL process ID.
Lock Mode	The lock mode.
Lock Type	The lock type.

Session Memory Allocation

Each row is a different session on the selected database. Column values refer to the session.

User Name	The name of the user with the blocked SQL process.
OS User	This data is obtained from the vendor. See vendor documentation for details.
Machine	The unique identifier of the user's machine.
Program	The name of the application.
Memory (MB)	The amount of memory used, in megabytes.

Blocked Sessions

Each row is a different session on the selected database. Column values refer to the session with a blocked SQL process and the SQL process causing the block.

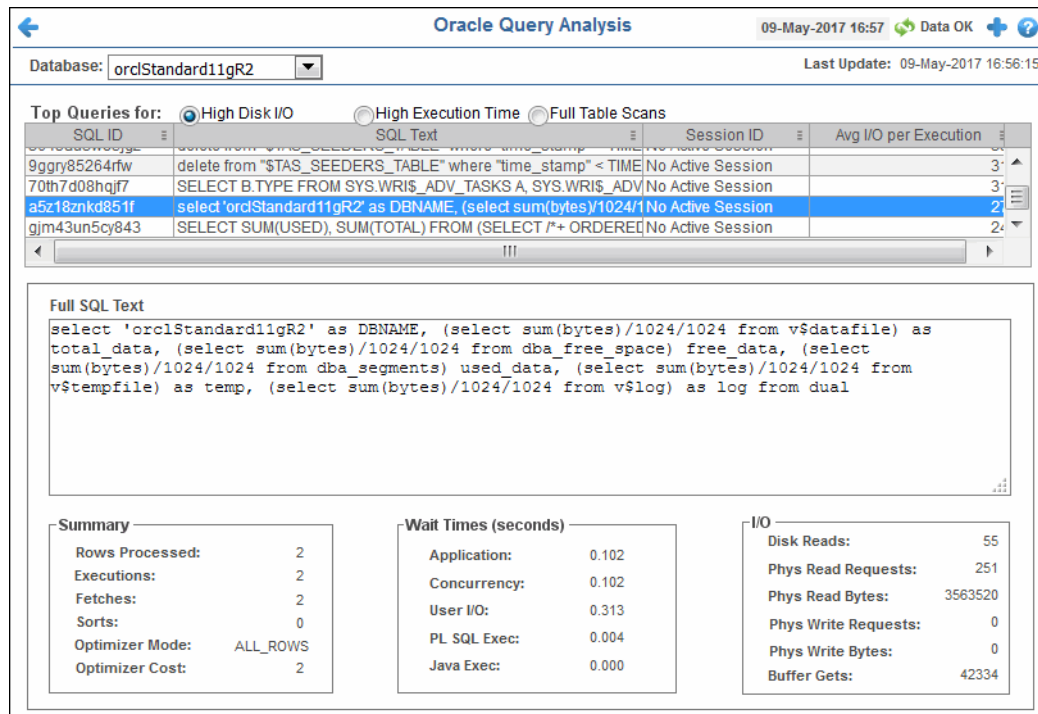
Blocking SID	The process ID of the SQL process that is blocking.
Blocked SID	The ID of the SQL process that is blocked.
Serial #	This data is obtained from the vendor. See vendor documentation for details.

User Name	The name of the user with the blocked SQL process.
OS User	This data is obtained from the vendor. See vendor documentation for details.
Machine	The unique identifier of the user's machine.
Wait Class	The wait class.
Seconds in Wait	The length of the wait, in seconds.

Queries

Analyse database query disk utilization, execution times and full table scans per SQL process.

Choose a database from the drop-down menu, then toggle between **High Disk I/O**, **High Execution Time** or **Full Table Scans**. Click a row to populate fields below the table, such as wait time, rows processed and disk reads.



Oracle Query Analysis 09-May-2017 16:57 Data OK

Database: orclStandard11gR2 Last Update: 09-May-2017 16:56:15

Top Queries for: ☒ High Disk I/O ☐ High Execution Time ☐ Full Table Scans

SQL ID	SQL Text	Session ID	Avg I/O per Execution
9ggry85264rfw	delete from "STAS_SEEDERS_TABLE" where "time_stamp" < TIME	No Active Session	3
70th7d08hqf7	SELECT B.TYPE FROM SYS.WRIS_ADV_TASKS A, SYS.WRIS_ADV	No Active Session	3
a5z18znkd851f	select 'orclStandard11gR2' as DBNAME, (select sum(bytes)/1024/1	No Active Session	2
gjm43un5cy843	SELECT SUM(USED), SUM(TOTAL) FROM (SELECT /*+ ORDERED	No Active Session	2

Full SQL Text

```
select 'orclStandard11gR2' as DBNAME, (select sum(bytes)/1024/1024 from v$datafile) as
total_data, (select sum(bytes)/1024/1024 from dba_free_space) free_data, (select
sum(bytes)/1024/1024 from dba_segments) used_data, (select sum(bytes)/1024/1024 from
v$tempfile) as temp, (select sum(bytes)/1024/1024 from v$log) as log from dual
```

Summary

Rows Processed:	2
Executions:	2
Fetches:	2
Sorts:	0
Optimizer Mode:	ALL_ROWS
Optimizer Cost:	2

Wait Times (seconds)

Application:	0.102
Concurrency:	0.102
User I/O:	0.313
PL SQL Exec:	0.004
Java Exec:	0.000

I/O

Disk Reads:	55
Phys Read Requests:	251
Phys Read Bytes:	3563520
Phys Write Requests:	0
Phys Write Bytes:	0
Buffer Gets:	42334

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.

Database Choose a database to display.

Top Queries for:

Each row is a different SQL process on the selected database. Column values refer to the SQL process.

High Disk I/O

SQL ID	The SQL Id.
SQL Text	The textual content.
Session ID	The unique identifier for the session.
Avg I/O per Execution	The average amount of data input and output per execution, in megabytes.

High Execution Time

SQL ID	The SQL Id.
SQL Text	The textual content.
Session ID	The unique identifier for the session.
Execution Time Sec	The amount time to execute per second, in megabytes.

Full Table Scans

SQL ID	The SQL Id.
SQL Text	The textual content.
Operation Cost	This data is obtained from the vendor. See vendor documentation for details.
CPU Cost	This data is obtained from the vendor. See vendor documentation for details.
IO Cost	This data is obtained from the vendor. See vendor documentation for details.
Temp Space	This data is obtained from the vendor. See vendor documentation for details.

Full SQL Text The SQL textual content.

Summary

Values refer to the SQL process selected in the table.

Rows Processed	The total number of rows processed since the process started.
Executions	The total number of executions since the process started.
Fetches	The total number of fetches since the process started.
Sorts	The total number of sorts since the process started.
Optimizer Mode	The optimizer mode used.

Optimizer Cost	The total number of single block disk reads since the process started.
-----------------------	--

Wait Times

Values refer to the SQL process selected in the table.

Application	The name of the application.
Concurrency	This data is obtained from the vendor. See vendor documentation for details.
User I/O	This data is obtained from the vendor. See vendor documentation for details.
PL SQL Exec	This data is obtained from the vendor. See vendor documentation for details.
Java Exec	This data is obtained from the vendor. See vendor documentation for details.

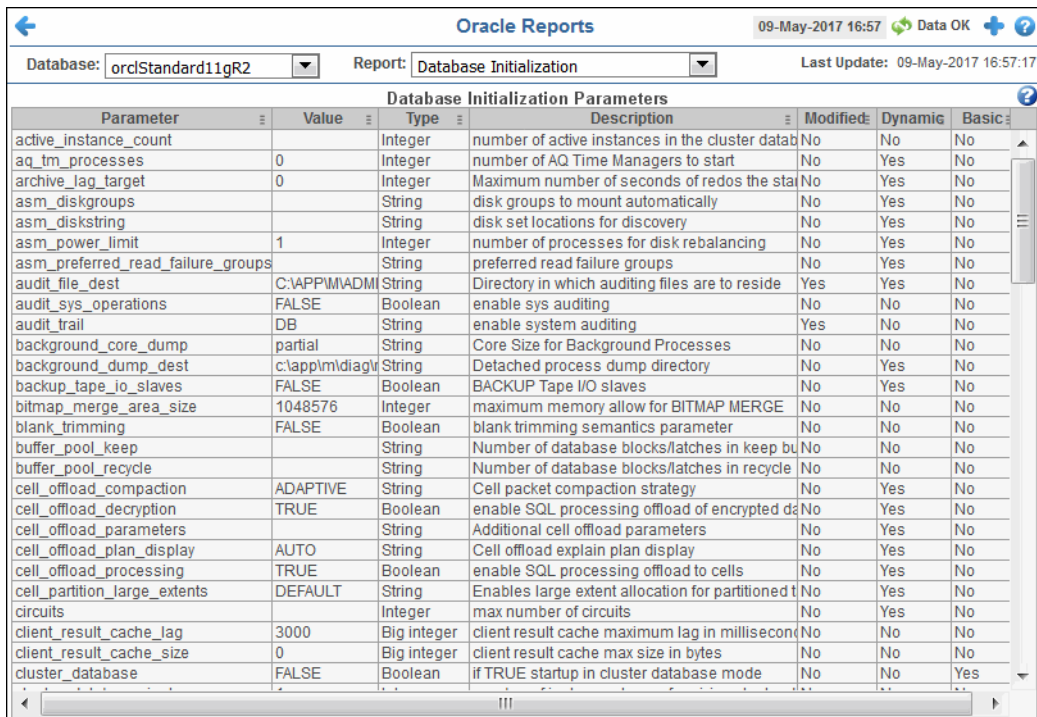
I/O

Values refer to the SQL process selected in the table.

Disk Reads	The total number of disk reads since the process started.
Phys Read Requests	The total number of physical disk reads since the process started.
Phys Read Bytes	The total number of physical disk reads since the process started.
Phys Write Requests	The total number of physical disk writes since the process started.
Phys Write Bytes	The total number of bytes written since the process started.
Buffer Gets	The total number of gets since the process started.

Reports


Analyse detailed database activity. Choose a database and a report from the drop-down menus. Reports available are: **Audited Objects**, **Audited Privileges**, **Database Initialization**, **Feature Usage** and **Invalid Objects**.




Parameter	Value	Type	Description	Modified	Dynamic	Basic
active_instance_count		Integer	number of active instances in the cluster database	No	No	No
aq_tm_processes	0	Integer	number of AQ Time Managers to start	No	Yes	No
archive_lag_target	0	Integer	Maximum number of seconds of redo the standby database can lag behind the primary	No	Yes	No
asm_diskgroups		String	disk groups to mount automatically	No	Yes	No
asm_diskstring		String	disk set locations for discovery	No	Yes	No
asm_power_limit	1	Integer	number of processes for disk rebalancing	No	Yes	No
asm_preferred_read_failure_groups		String	preferred read failure groups	No	Yes	No
audit_file_dest	C:\APP\ADMINISTRATOR\diag\audit	String	Directory in which auditing files are to reside	Yes	Yes	No
audit_sys_operations	FALSE	Boolean	enable sys auditing	No	No	No
audit_trail	DB	String	enable system auditing	Yes	No	No
background_core_dump	partial	String	Core Size for Background Processes	No	No	No
background_dump_dest	c:\app\ADMINISTRATOR\diag\background_dump	String	Detached process dump directory	No	Yes	No
backup_tape_io_slaves	FALSE	Boolean	BACKUP Tape I/O slaves	No	Yes	No
bitmap_merge_area_size	1048576	Integer	maximum memory allow for BITMAP MERGE	No	No	No
blank_trimming	FALSE	Boolean	blank trimming semantics parameter	No	No	No
buffer_pool_keep		String	Number of database blocks/latches in keep buffer pool	No	No	No
buffer_pool_recycle		String	Number of database blocks/latches in recycle buffer pool	No	No	No
cell_offload_compaction	ADAPTIVE	String	Cell packet compaction strategy	No	Yes	No
cell_offload_decryption	TRUE	Boolean	enable SQL processing offload of encrypted data	No	Yes	No
cell_offload_parameters		String	Additional cell offload parameters	No	Yes	No
cell_offload_plan_display	AUTO	String	Cell offload explain plan display	No	Yes	No
cell_offload_processing	TRUE	Boolean	enable SQL processing offload to cells	No	Yes	No
cell_partition_large_extents	DEFAULT	String	Enables large extent allocation for partitioned tables	No	Yes	No
circuits		Integer	max number of circuits	No	Yes	No
client_result_cache_lag	3000	Big integer	client result cache maximum lag in milliseconds	No	No	No
client_result_cache_size	0	Big integer	client result cache max size in bytes	No	No	No
cluster_database	FALSE	Boolean	if TRUE startup in cluster database mode	No	No	Yes

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.

Databases Choose a database to display.

Report Choose a report to display.

Audited Objects

Each row is a different audited object on the selected database. Column values refer to the object.

Success	This data is obtained from the vendor. See vendor documentation for details.
Failure	This data is obtained from the vendor. See vendor documentation for details.
Schema	This data is obtained from the vendor. See vendor documentation for details.

Object Name	This data is obtained from the vendor. See vendor documentation for details.
Statement	This data is obtained from the vendor. See vendor documentation for details.
Object Type	This data is obtained from the vendor. See vendor documentation for details.

Audited Privileges

Each row is a different audited privilege on the selected database. Column values refer to the privilege.

Privilege	The name of the privilege.
User	The name of the user associated with the audit.
Proxy	The name of the proxy. This data is obtained from the vendor. See vendor documentation for details.
Success	This data is obtained from the vendor. See vendor documentation for details.
Failure	This data is obtained from the vendor. See vendor documentation for details.

Database Initialization

Each row is a different parameter on the selected database. Column values refer to the parameter.

Parameter	The name of the parameter.
Value	The configuration setting of the parameter.
Type	The configuration setting type. Values are: <ul style="list-style-type: none"> • Integer • Boolean • String • Big Integer
Description	Text describing the parameter.
Modified	Indicates whether the parameter was modified. Values are Yes/No . This data is obtained from the vendor. See vendor documentation for details.
Dynamic	Indicates whether the parameter is dynamic. Values are Yes/No . This data is obtained from the vendor. See vendor documentation for details.
Basic	Indicates whether the parameter is basic. Values are Yes/No . This data is obtained from the vendor. See vendor documentation for details.

Feature Usage

Each row is a different feature on the selected database. Column values refer to the feature.

Feature Name	The name of the feature.
---------------------	--------------------------

Currently Used	Indicates whether the feature is currently being used. Values are True/False .
#Detected Uses	The total number of times the feature has been used since the database started.
Total Samples	This data is obtained from the vendor. See vendor documentation for details.
First Use	The date and time the feature was first used.
Last Use	The date and time the feature was last used.
Version	This data is obtained from the vendor. See vendor documentation for details.

Invalid Objects

Each row is a different invalid object on the selected database. Column values refer to the invalid object.

Owner	The object owner.
Object Type	The type of object.
Object Name	The name of the object.

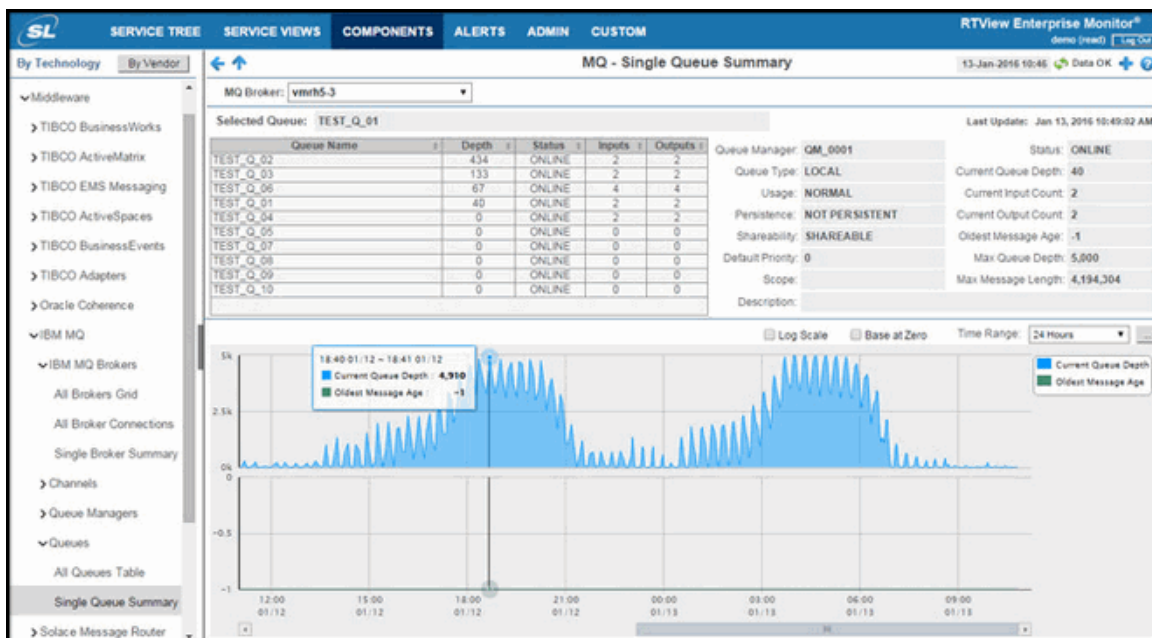
Connector for Oracle Enterprise Manager

The Connector for Oracle Enterprise Manager (OEM) allows RTView Enterprise Monitor® to connect to existing deployments of OEM and collect performance data for databases and hosts (physical servers) that have been designated as OEM targets.

When paired with the [“Oracle Database”](#) and [“RTView Host Agent”](#) solution packages, these performance metrics are then stored in the RTView Enterprise Monitor caches and available for summary views detailing the health of your OEM managed hosts and databases, including drill down views, correlation with services and other technologies, historical analysis, capacity planning and alert management.

Oracle WebLogic

Use the Solution Package for Oracle® WebLogic to monitor the health of your Oracle WebLogic servers and applications as well as your JMS destinations, bridges and connections.



Note: This document assumes familiarity with Oracle WebLogic. For additional details, refer to vendor documentation.

The following Solution Package for Oracle® WebLogic Views can be found under **Components** tab > **Application/Web Servers** > **Oracle WebLogic**:

- "WebLogic Servers View"
- "Single WebLogic Server View"
- "Application Views View"
- "JMS Servers View"
- "JMS Destinations View"
- "JMS Bridges View"
- "JMS Connections View"

WebLogic Servers View

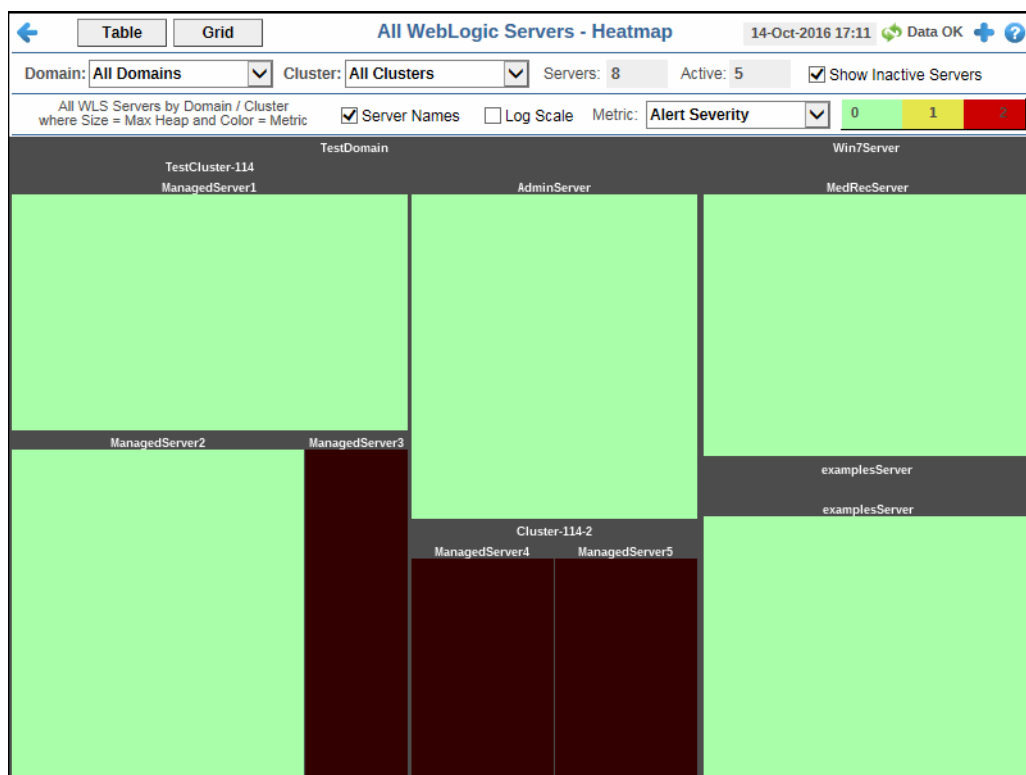
These displays present performance metrics and alert statuses for all Oracle WebLogic Servers and Clusters. The following displays are available:


- **"All Servers Heatmap"**: This display shows status and alerts for all Oracle® WebLogic servers in a heatmap.
- **"All Servers Table"**: This displays shows all available utilization metrics for all Oracle® WebLogic servers in a tabular format.
- **"All Servers Grid"**: This display enables you to track utilization and performance metrics and trend data for all WebLogic servers on a particular domain.
- **"All Clusters Table"**: This table enables you track utilization and performance metrics for all clusters on a particular domain, or on all domains

All Servers Heatmap

View status and alerts of all Oracle® WebLogic servers. Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **Jvm CPU %**, **Host CPU %**, **Jvm Memory %**, **Open Sockets**, **Thread Total Count**, and **Hogging Threads**.

The heatmap is organized by host, each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the **"All Servers Grid"** display and view metrics for a particular connection. You can toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the icon in the upper left-hand corner. Mouse-over rectangles to view more details about host performance and status.





Fields and Data

This display includes:

Domain	Select the domain (or All Domains) from the drop down list for which you want to view data.
Cluster	Select the cluster (or All Clusters) from the drop down list for which you want to view data.
Servers	The total number of active, inactive, and standby servers.
Active	The number of active servers listed in the display.
Show Inactive Servers	Select this check box to display inactive servers in the heatmap.
Server Names	Select this check box to display the names of the servers in the heatmap
Log Scale	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
Metric	Select the metric driving the heatmap display. The default is Alert Severity. Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the instances by host, where each rectangle represents an instance. Mouse-over any rectangle to display the current values of the metrics for the instance. Click on a rectangle to drill-down to the associated "All Servers Grid" display for a detailed view of metrics for that particular instance.

Alert Severity

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



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


2 -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

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


0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

Alert Count


The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

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

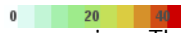
Jvm CPU %

The percentage of JVM CPU currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsServerCpuHigh**. The middle value in the gradient bar indicates the middle value of the range.

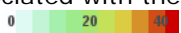
Host CPU %

The percentage of Host CPU currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsServerHostCpuHigh**. The middle value in the gradient bar indicates the middle value of the range.


JVM Memory %

The percentage of JVM Memory currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsMemoryUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

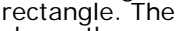
Open Sockets

The total number of open sockets currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsOpenSocketsHigh**. The middle value in the gradient bar indicates the middle value of the range.

Thread Total Count

The total number of threads in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsThreadsTotalHigh**. The middle value in the gradient bar indicates the middle value of the range.

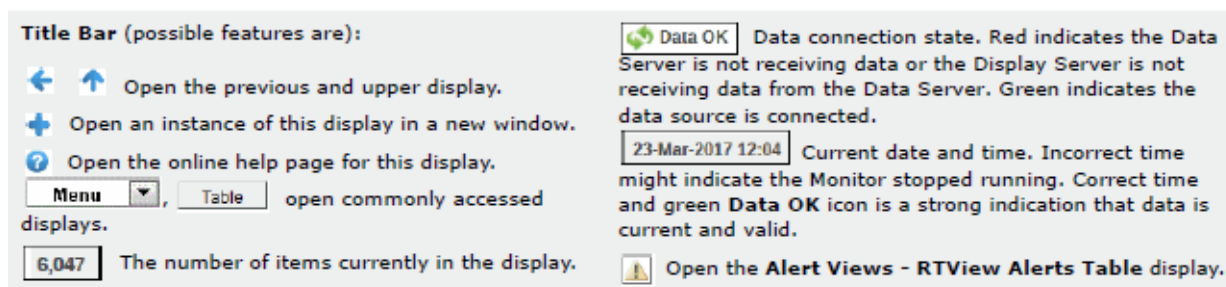
Hogging Threads

The total number of hogging threads currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsHoggingThreadsHigh**. The middle value in the gradient bar indicates the middle value of the range.

All Servers Table

This display provides utilization metrics for all WebLogic Servers for a particular domain in a tabular format. Each row in this table includes heap, processing, thread, and version metrics (among others) for a particular server. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill-down to the ["Server Summary"](#) display and view metrics for that particular server. You can click on one of the buttons in the upper left-hand corner to toggle between the commonly accessed **Grid** and **Heatmap** displays.

All WebLogic Servers - Table									
Domain: All Domains		Cluster: All Clusters		Servers: 8		Active: 5		<input checked="" type="checkbox"/> Show Inactive Servers	
Domain	Cluster	Server	State	Expired	Alerts	Port	JVM Proc Load %	Heap Free %	...
examplesServer		examplesServer	RUNNING	<input type="checkbox"/>		80	7.21	61.0	5
TestDomain		AdminServer	RUNNING	<input type="checkbox"/>		6001	1.60	5.0	5
TestDomain	Cluster-114-2	ManagedServer4	SHUTDO...	<input type="checkbox"/>		0	0	0.0	5
TestDomain	Cluster-114-2	ManagedServer5	SHUTDO...	<input type="checkbox"/>		0	0	0.0	5
TestDomain	TestCluster-114	ManagedServer1	RUNNING	<input type="checkbox"/>		6003	1.26	31.0	5
TestDomain	TestCluster-114	ManagedServer2	RUNNING	<input type="checkbox"/>		6004	0.00	47.0	5
TestDomain	TestCluster-114	ManagedServer3	SHUTDO...	<input type="checkbox"/>		0	0	0.0	5
Win7Server		MedRecServer	RUNNING	<input type="checkbox"/>		80	2.99	76.0	5



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

Fields and Data

This display includes:

Domain	Select the domain for which you want to view data, or select All Domains to view data for all domains.
Cluster	Select the cluster on the domain for which you want to view data, or select All Clusters to view data for all domains.
Servers	The total number of servers on the cluster.
Active	The total number of active servers on the cluster.
Show Inactive Servers	Select this check box to display inactive servers in the table.

All WebLogic Servers Table

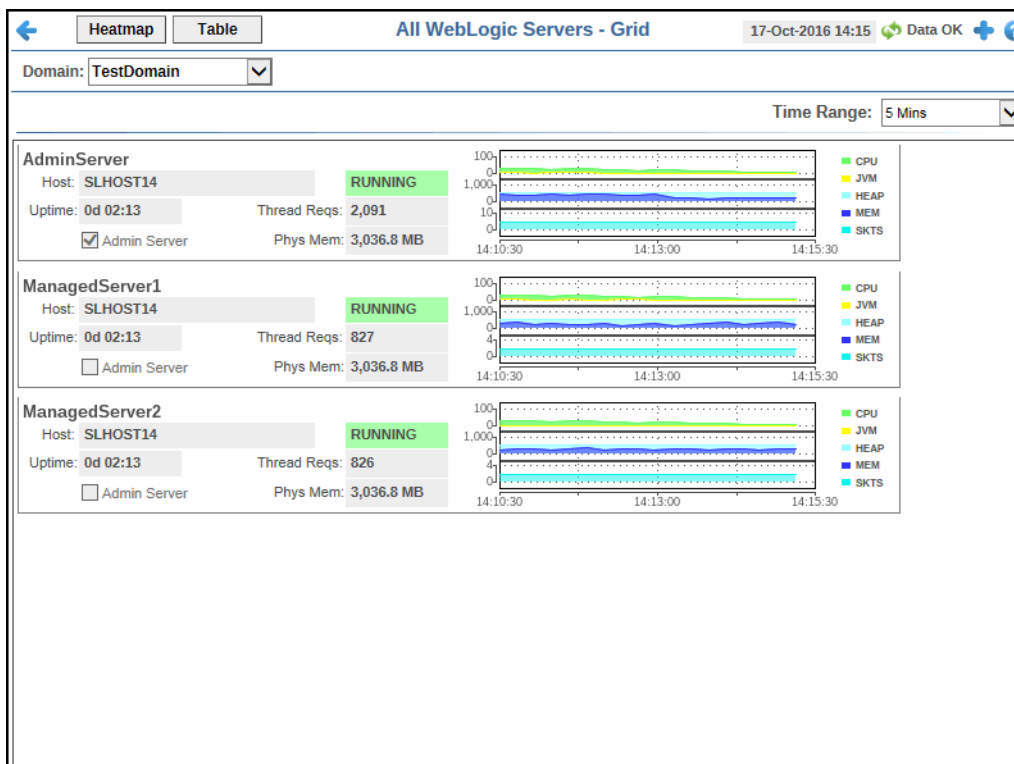
This table shows information for the selected domain/cluster(s) combination. Click on a table row to drill-down to the "Server Summary" display and view metrics for that particular server.

Domain	The name of the domain.
Cluster	The name of the cluster.
Server	The name of the server.
State	The current state of the server.
Expired	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
Alerts	<p>The current alert level.</p> <ul style="list-style-type: none"> -- One or more alerts have exceeded their specified ALARMLEVEL threshold. -- One or more alerts have exceeded their specified WARNINGLEVEL threshold. -- No alerts have exceeded an alert threshold.

Port	The port on which this server is listening for SSL connections.*
JVM Proc Load %	A snapshot of the load that the virtual machine is placing on all processors on the host computer.*
Heap Free %	The percentage of free heap memory on the server.*
Heap Max (bytes)	The maximum amount of heap, in bytes, available for use.*
Used Heap (bytes)	The total amount of heap used, in bytes.*
Heap Current (bytes)	The current size of the JVM heap, in bytes, being used.*
Open Sockets	The current number of sockets registered for socket muxing on this server.*
Hogging Threads	The number of hogging threads on the server.*
Execute Threads	The current number of execute threads.*
Idle Threads	The current number of idle threads.
Restarts Count	The total number of restarts for this server since the cluster was last started.*
All Procs Avg Load %	The average load percentage for all processors on the host computer.*
Shutting Down	When checked, denotes that the server is currently shutting down.*
Restart Required	When checked, denotes that the server needs to be restarted in order to activate configuration changes.*
Uptime	The length of time (in milliseconds) that the server has been up and running.*
Startup Time	The length of time (in milliseconds) that it took for the server to start up.*
WebLogic Version	The current version of WebLogic running on the server.*
JVM Type	The type of JVM currently being used on the server.*
Java Version	The current version of Java running on the server.*
JavaVendor	The name of the vendor of the Java version running on the server.*
OS Name	The name of the operating system running on the server.*
Time Stamp	The date and time this row of data was last updated.

All Servers Grid

Track utilization and performance metrics and trend data for all WebLogic servers on a particular domain.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

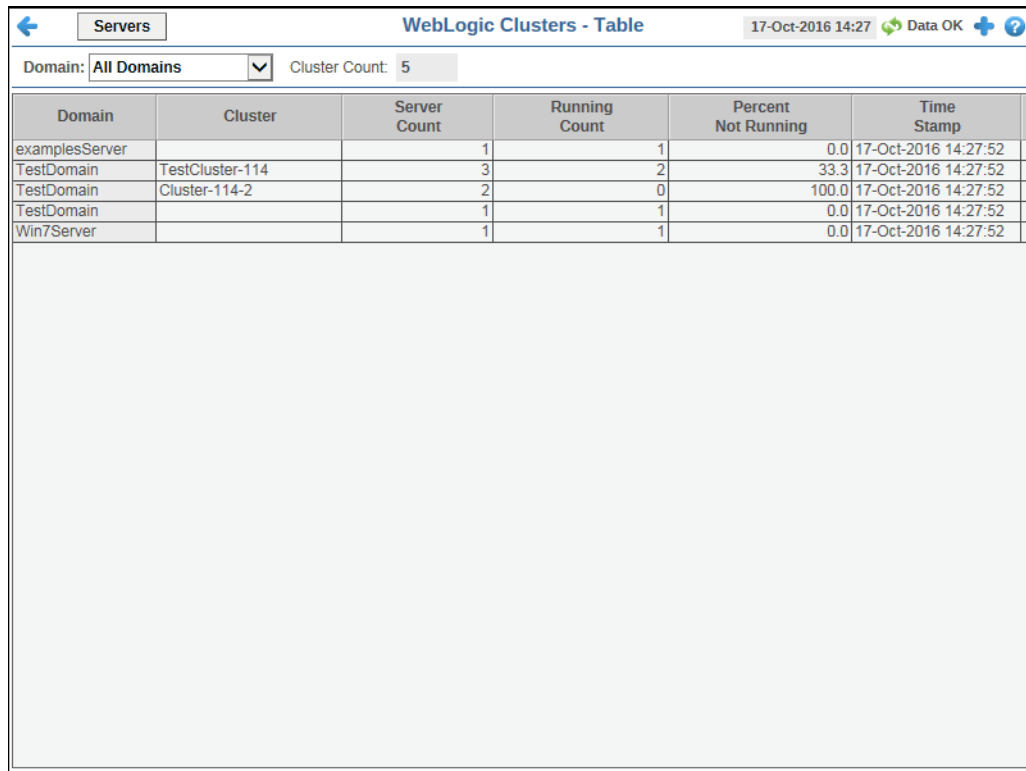
Domain

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

Time Range	Select a time range from this drop down menu to define the data displayed in the trend graph for a selected period of time. You can select from as little as the past 2 Minutes to the Last 7 Days , or you can display All Data .
Server Grid	Displays data and a trend graph for each server in your domain. The trend graph metrics are: <ul style="list-style-type: none"> CPU -- Traces the amount of CPU being used by the server. JVM -- Traces the JVM processing load that the virtual machine is placing on all processors on the host computer. HEAP -- Traces the total amount of heap used, in bytes. MEM -- Traces the amount of memory being used. SKTS -- Traces the current number of sockets registered for socket muxing on this server.
(Server Name)	Displays the name of the Server
Uptime	The amount of time since the server was last started, shown in days, hours, and minutes (for example, 1d 23:43).
Admin Server	Indicates whether the server is an Administration Server.*
(Status)	Displays the status of the server.*
Thread Reqs	The current number of thread requests.*
Phys Mem	Displays the available physical memory (in megabytes) for the server.








All Clusters Table


Track utilization and performance metrics for all clusters on a particular domain, or on all domains.





Domain	Cluster	Server Count	Running Count	Percent Not Running	Time Stamp
examplesServer		1	1	0.0	17-Oct-2016 14:27:52
TestDomain	TestCluster-114	3	2	33.3	17-Oct-2016 14:27:52
TestDomain	Cluster-114-2	2	0	100.0	17-Oct-2016 14:27:52
TestDomain		1	1	0.0	17-Oct-2016 14:27:52
Win7Server		1	1	0.0	17-Oct-2016 14:27:52

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.
-   open commonly accessed displays.
-  6,047 The number of items currently in the display.

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

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

 Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

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

Cluster Count The current number of clusters listed in the table.

Clusters Table

Lists the clusters in the currently selected domain, or lists all clusters in all domains.

Domain	The name of the domain
Cluster	The name of the cluster.*
Server Count	The total number of servers on the cluster.*
Running Count	The total number of servers running on the cluster.*
Percent Not Running	The percentage of servers not running on the cluster.*
Time Stamp	The date and time this row of data was last updated.

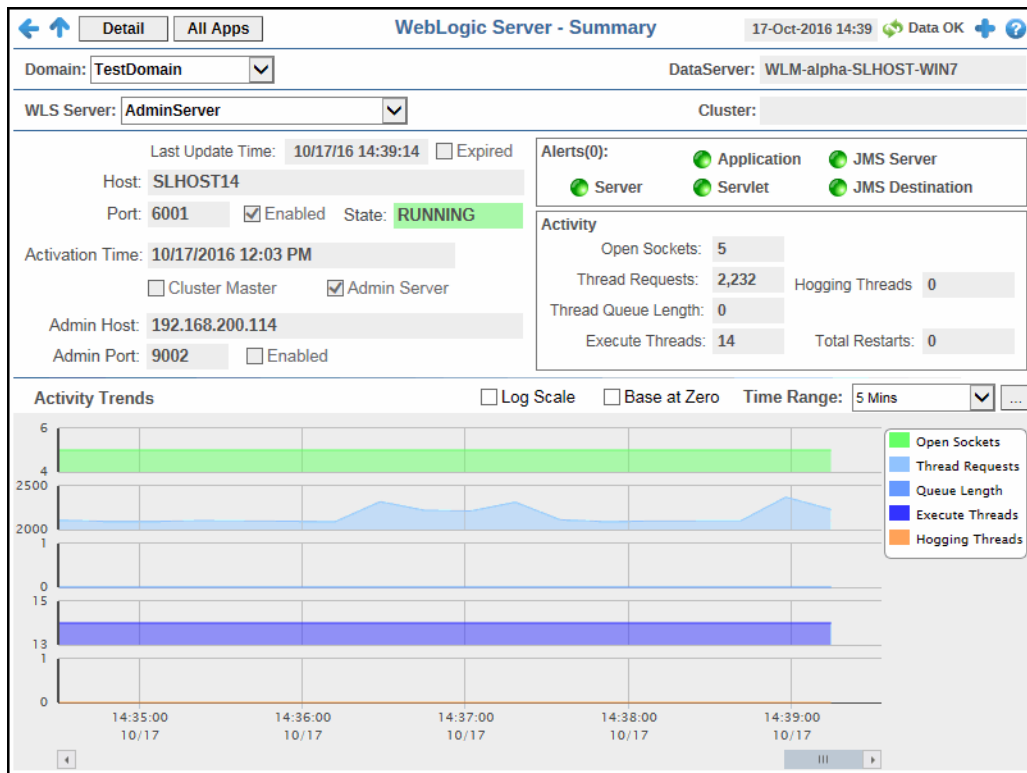
Single WebLogic Server View

These displays present detailed performance metrics and alert statuses for a particular WebLogic server.

- [“Server Summary”](#): Track utilization, performance, and trend data for a particular WebLogic server
- [“WLS JVM Summary”](#): Displays the JVM details for a particular WebLogic server on a specific domain.
- [“WLS Server Detail”](#): Displays server runtime data, threadpool runtime data, JRockit runtime data, and server version information for a specific WebLogic server
- [“WLS JDBC Summary”](#): Displays JDBC module utilization, performance, and trend data for a specific WebLogic server.
- [“WLS ThreadPool Summary”](#): Displays threadpool utilization, performance, and trend data for a specific WebLogic server.
- [“Work Manager”](#): Displays server runtime data for all work managers on a specific WebLogic Server.
- [“Persistent Stores”](#): Displays available utilization and performance data for all configurations on a specific domain.

Server Summary

Track utilization, performance, and trend data for a particular WebLogic server.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:




Domain Select the domain for which you want to view data.

Data Server	The name of the data server.
WLS Server	Select the WebLogic server for which you want to see data.
Cluster	The name of the cluster.

(Server Information)

Last Update Time	The date and time the data in the display was last updated.
Expired	This check box becomes automatically checked when the data has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
Host	The name of the host.*
Port/Enabled	The name of the port. The port is enabled when the associated Enabled check box displays as checked.*
State	The current state of the WebLogic server.*
Activation Time	The date and time in which the server was started.*
Cluster Master	When selected, denotes that the server is a cluster master.*
Admin Server	Indicates whether the server is an Administration Server.*
Admin Host	The IP address of the administration server's host.*
Admin Port/Enabled	The name of the administration server's port. The port is enabled when the associated Enabled check box displays as checked.*

Alerts (#) -- Displays the total number of alerts and the current status of the associated **Application, JMS Server, Server, Servlet, and JMS Destination**.

-  -- One or more alerts have exceeded their specified **ALARMLEVEL** threshold.
-  -- One or more alerts have exceeded their specified **WARNINGLEVEL** threshold.
-  -- No alerts have exceeded an alert threshold.

Activity

Open Sockets	The number of current open sockets for the server.
Thread Requests	The current number of thread requests.*
Hogging Threads	The current number of hogging threads.*
Thread Queue Length	The current thread queue length.*
Execute Threads	The current number of execute threads.*
Total Restarts	The total number of times the server has restarted since the last update time.*

Activity Trends

Displays data and a trend graph for the following:

Open Sockets-- Traces the number of open sockets of the server.

Thread Requests-- Traces the number of thread requests on the server.

Queue Length-- Traces the queue length on the server.

Execute Threads-- Traces the number of execute threads on the server.

Hogging Threads-- Traces the number of hogging threads on the server.


Log Scale


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


Base at Zero


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

Time Range

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



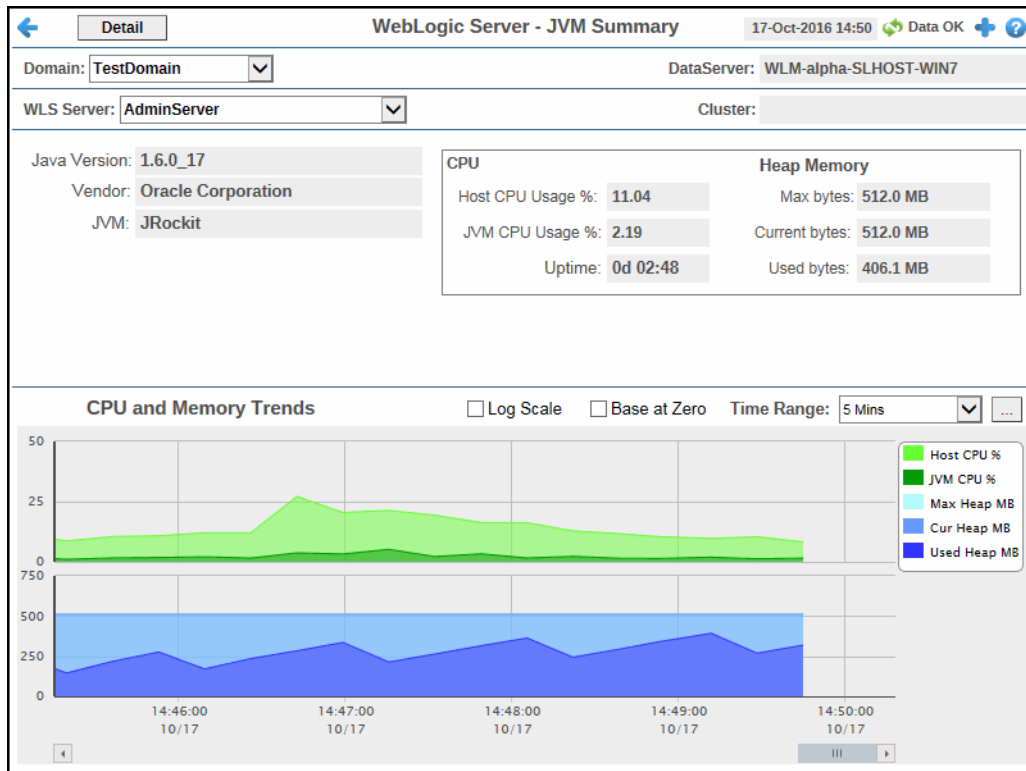
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

WLS JVM Summary

This display allows you to view the JVM details for a particular WebLogic server on a specific domain.



Title Bar (possible features are):

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

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

Domain Select the domain for which you want to view data.

Data Server	The name of the data server.
WLS Server	Select the WebLogic server for which you want to see data.
Cluster	The name of the cluster.

(JVM Information)

Java Version	The current version of Java running on the server.*
Vendor	The name of the vendor of the current version of Java running on the server.*
JVM	The type of JVM currently being used on the server.*




CPU

Host CPU Usage %	The current CPU usage percentage on the host.*
JVM CPU Usage %	The current JVM CPU usage percentage.*
Uptime	The amount of time since the server was last started, shown in days, hours, and minutes (for example, 1d 23:43).*

Heap Memory

Max Bytes	The maximum amount of available heap memory, in megabytes.*
Current Bytes	The current size of the JVM heap, in megabytes.*
Used Bytes	The amount of heap memory used, in megabytes.*

Alerts (#) -- Displays the total number of alerts and the current status of the associated **Application**, **JMS Server**, **Server**, **Servlet**, and **JMS Destination**.

-  -- One or more alerts have exceeded their specified **ALARMLEVEL** threshold.
-  -- One or more alerts have exceeded their specified **WARNINGLEVEL** threshold.
-  -- No alerts have exceeded an alert threshold.

CPU and Memory Trends

Displays data and a trend graph for the following:

- Host CPU %** -- Traces the percentage of the host CPU being used.
- JVM CPU %** -- Traces the percentage of the JVM CPU being used.
- Max Heap MB** -- Traces the maximum amount of heap memory available (in megabytes).
- Current Heap MB** -- Traces the current size of the JVM heap, in megabytes.
- Used Heap MB** -- Traces the total amount of heap used, in bytes.*


Log Scale


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


Base at Zero



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

Time Range

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



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

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

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

WLS Server Detail

View server runtime, threadpool runtime, JRockit runtime, and server version information for a specific WebLogic server.

WLS Server - Detail Tables 17-Oct-2016 14:59 Data OK

Domain: **TestDomain** DataServer: **WLM-alpha-SLHOST-WIN7**

WLS Server: **AdminServer** Cluster:

Server Runtime							
Connection	Location	Cluster	ActivationTime	AdminServer	AdminServerHost	AdminServerListenPort	Ad
TestDomain	AdminServer		1476731039316	<input checked="" type="checkbox"/>	192.168.200.114	6001	

ThreadPool Runtime					
Connection	CompletedRequestCount	DeltaCompletedRequestCount	ExecuteThreadIdleCount	ExecuteThreadTotalCount	Hog
TestDomain	1,353,303		2,098	8	14

JRockit Runtime						
Connection	AllProcessorsAverageLoad	Concurrent	FreeHeap	FreePhysicalMemory	GCHandlesCompaction	
TestDomain	6.934	<input type="checkbox"/>	215,658,408	929,218,560	<input checked="" type="checkbox"/>	Dynamic GC

Server Version Info

WebLogic Server 10.3.3.0 Fri Apr 9 00:05:28 PDT 2010 1321401

Title Bar (possible features are):

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

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

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

Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

Domain Select the domain for which you want to view data.

Data Server	The name of the data server.
WLS Server	Select the WebLogic server for which you want to see data.
Cluster	The name of the cluster.

Server Runtime Table

Connection	The name of the connection.
Location	The name of the WLS Server located on the specified connection.
Cluster	The name of the cluster.
Activation Time	The time when the server was started.*
Admin Server	Indicates whether the server is an Administration Server.*
Admin Server Host	The address on which the Administration Server is listening for connections.*
Admin Server Listen Port	The port on which the Administration Server is listening for connections.*
Admin Server Listen Port Secure	Indicates whether the port that the server uses for administrative traffic is configured to use a secure protocol.*
Administration Port	The port on which this server is listening for administrative requests.*
Administration Port Enabled	Indicates whether the administration port is enabled on the server.*
Administration URL	The URL that the server and its clients use for administrative connections.*
Cluster Master	When checked, denotes that the cluster is a cluster master.*
Current Directory	The absolute path of the directory from which the server was started.*
Current Machine	The machine on which the server is running.*
Default URL	The URL that clients use to connect to this server's default network channel.*
Health State	The health state of the server as reported by the server's self-health monitoring.*
Listen Address	The address on which this server is listening for connections through the default network channel.*
Listen Port	The port on which this server is listening for connections.*
Listen Port Enabled	Indicates whether the default listen port is enabled on the server.*
Name	The name of the Java Virtual Machine.*
Open Sockets Current Count	The current number of sockets registered for socket muxing on this server.*

Oracle Home	The Oracle home directory path.*
Parent	The name of the parent of the Java Virtual Machine.*
Pending Restart System Resources	The number of system resources that have not been restarted since the last restart of the WLS Server.
Restart Required	Indicates whether the server must be restarted in order to activate configuration changes.*
Restarts Total Count	The total number of restarts for this server since the cluster was last started.*
SSL Listen Address	The address on which this server is listening for SSL connections.*
SSL Listen Port	The port on which this server is listening for SSL connections.*
SSL Listen Port Enabled	Indicates whether the default SSL listen port is enabled on the server.*
Server Classpath	The class path for this server including domain/lib contents that are automatically picked up and appended to the classpath.*
Server Startup Time	The startup time of the server.*
Shutting Down	Indicates whether the server is shutting down.
Stable State	The current state of the server as an integer.*
State Prev	The state of the server prior to its current state.*
State Val	The current state of the server as an integer.*
WebLogic Home	The WebLogic Home directory path.*
Expired	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
State	The current life cycle state of the server.*

ThreadPool Runtime Table

Connection	The name of the connection.
Completed Request Count	The number of completed requests in the priority queue.*
Delta Completed Request Count	The increase in the amount of completed requests (from the previous polling period to the current polling period).

Execute Thread Idle Count	The number of idle threads in the pool. This count does not include standby threads and stuck threads. The count indicates threads that are ready to pick up new work when it arrives.*
Execute Thread Total Count	The total number of threads in the pool.*
Hogging Thread Count	The threads that are currently being hogged by a request. These threads will either be declared as stuck after the configured timeout or will be returned to the pool. The self-tuning mechanism will backfill if necessary.*
Delta Hogging Thread Count	The increase in the amount of hogging threads (from the previous polling period to the current polling period).
Min Threads Constraints Completed	The number of requests with minimum threads constraint picked up out of order for execution immediately since their minimum threads requirement was not met. This does not include the case where threads are idle during schedule.*
Min Threads Constraints Pending	The number of requests that should be executed now to satisfy the minimum threads requirement.*
Pending User Request Count	The number of pending user requests in the priority queue. The priority queue contains requests from internal subsystems and users. This is just the count of all user requests.*
Queue Length	The number of pending requests, which consist of the total number of internal system requests and user requests, in the priority queue.*
Shared Capacity For Work Managers	The maximum amount of requests that can be accepted in the priority queue.*
Standby Thread Count	The number of threads in the standby pool. Surplus threads that are not needed to handle the present work load are designated as standby and added to the standby pool. These threads are activated when more threads are needed.*
Suspended	Indicates if the RequestManager is suspended. A suspended manager will not dequeue work and dispatch threads until it is resumed.*
Throughput	The mean number of requests completed per second.*
Name	The name of the Java Virtual Machine.
Parent	The name of the parent of the Java Virtual Machine.

JRockit Runtime Table

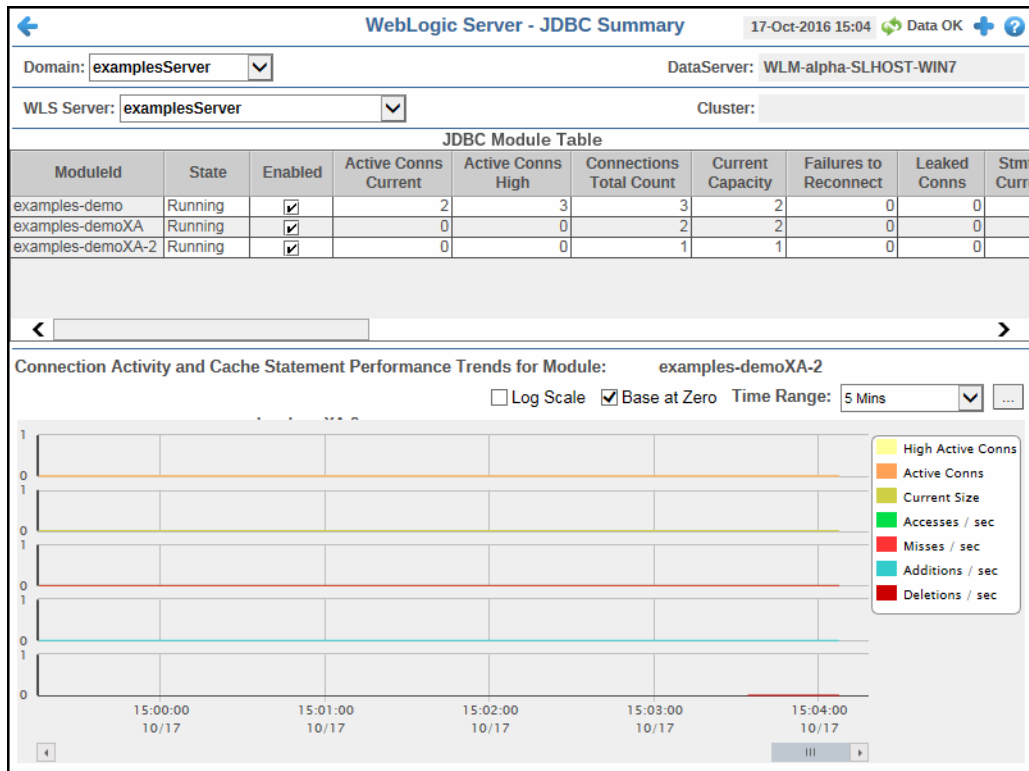
Connection	The name of the connection.
All Processors Average Load	The average load of all processors in the host computer.*
Concurrent	Indicates whether the virtual machine's garbage collector runs in a separate Java thread concurrently with other Java threads.*

Free Heap	The amount, in bytes, of Java heap memory that is currently free in the virtual machine.*
Free Physical Memory	The amount, in bytes, of physical memory that is currently free on the host computer.*
GC Handles Compaction	Indicates whether the virtual machine's garbage collector compacts the Java heap.*
GcAlgorithm	The type of garbage collector (GC) that the virtual machine is using.*
Generational	Indicates whether the virtual machine's garbage collector uses a nursery space. A nursery is the area of the Java heap that the virtual machine allocates to most objects.*
Heap Free Current	The current amount of memory, in bytes, that is available in the JVM heap.*
Heap Free Percent	Percentage of the maximum memory that is free.*
Heap Size Current	The current size, in bytes, of the JVM heap.*
Heap Size Max	The maximum free memory configured for this JVM.*
Incremental	Indicates whether the virtual machine's garbage collector collects (increments) garbage as it scans the memory space and dumps the garbage at the end of its cycle. With a non-incremental garbage collector, garbage is dumped as soon as it is encountered.*
JVM Description	The description of the Java Virtual Machine.*
Java VM Vendor	The vendor of the Java Virtual Machine that the server is running.*
Java Vendor	The vendor of Java that the server is running.*
Java Version	The Java version of the Java Virtual Machine.*
JVM Processor Load	A snapshot of the load that the virtual machine is placing on all processors in the host computer. If the host contains multiple processors, the value represents a snapshot of the average load.*
Jvm Type	The Java Virtual Machine type.*
Last GC End	The time at which the last garbage collection run ended.*
Last GC Start	The time at which the last garbage collection run started.*
Name	The name of the Java Virtual Machine.*
Number Of Daemon Threads	The number of daemon Java threads currently running in the Virtual Machine across all processors.*

Number Of Processors	The number of processors on the virtual machine's host computer. If this is not a Symmetric Multi-Processor (SMP) system, the value will be 1.*
OS Name	The name of the operating system on which the JVM is running.*
OS Version	The version of the operating system on which the JVM is running.*
Parallel	Indicates whether the virtual machine's garbage collector is able to run in parallel on multiple processors if multiple processors are available.*
Parent	The name of the immediate parent.*
Total Garbage Collection Count	The number of garbage collection runs that have occurred since the virtual machine was started.*
Total Garbage Collection Time	The number of milliseconds that the virtual machine has spent on all garbage collection runs since the virtual machine was started.*
Total Heap	The amount, in bytes, of memory currently allocated to the virtual machine's Java heap.*
Total Number Of Threads	The number of Java threads (daemon and non-daemon) that are currently running in the virtual machine across all processors.*
Total Nursery Size	The amount, in bytes, of memory that is currently allocated to the nursery.*
Total Physical Memory	The amount (in bytes) of physical memory on the host computer.*
Uptime	The amount of time, in milliseconds, that the virtual machine has been running.*
Used Heap	The amount of Java heap memory, in bytes, that is currently being used by the virtual machine.*
Used Physical Memory	The amount of physical memory, in bytes, that is currently being used on the host computer.*
Vendor	The name of the JVM vendor.*
Version	The current version of the Java Virtual Machine.*
Heap Used Current	The current amount of JVM heap memory, in bytes, that is being used.*
Memory Used Percent	The percentage of JVM heap memory that is being used.*
Server Version Info Region	Lists the WebLogic server version number and date that it was installed.

WLS JDBC Summary

View JDBC module utilization, performance data, and trend data for a specific WebLogic server.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data.
- Data Server** The name of the data server.

WLS Server Select the WebLogic server for which you want to see data.

Cluster The name of the cluster.

JDBC Module Table


Module Id	The name of the data source.
State	The current state of the data source.
Enabled	Indicates whether the data source is enabled or disabled. When checked, the data source is enabled.*
Active Conns Current	The number of connections currently in use by applications.*
Active Conns High	The highest number of active database connections in the instance of the data source since the data source was instantiated.*
Connections Total Count	The cumulative total number of database connections created in this data source since the data source was deployed.*
Current Capacity	The current count of JDBC connections in the connection pool in the data source.*
Failures to Reconnect	The number of times that the data source attempted to refresh a database connection and failed.*
Leaked Conns	The number of leaked connections. *
Stmt Cache Current Size	The number of prepared and callable statements currently cached in the statement cache.*
Stmt Cache Hits	The cumulative, running count of the number of times statements from the cache were used.*
Stmt Cache Missed	The number of times that a statement request could not be satisfied with a statement from the cache.*
Stmt Cache Accesses	The cumulative, running count of the number of times that the statement cache was accessed.*
Stmt Cache Additions	The cumulative, running count of the number of statements added to the statement cache.*
Stmt Cache Deletions	The cumulative, running count of statements discarded from the cache.*
Reserve Request	The cumulative, running count of requests for a connection from this data source.*
Failed Reserve Requests	The cumulative, running count of requests for a connection from this data source that could not be fulfilled.*
Wait Secs High	The cumulative total number of database connections created in this data source since the data source was deployed.*
Waiting Conn	The number of connection requests waiting for a database connection.*
Waiting Conn Fail	The cumulative, running count of requests for a connection from this data source that had to wait before getting a connection and eventually failed to get a connection.*


	Waiting Conn High	The highest number of application requests concurrently waiting for a connection from this instance of the data source.*
	Waiting Conn Success	The cumulative, running count of requests for a connection from this data source that had to wait before getting a connection and eventually succeeded in getting a connection.*
	Waiting Conn Total	The cumulative, running count of requests for a connection from this data source that had to wait before getting a connection, including those that eventually got a connection and those that did not get a connection.*
	Connection Delay Time	The average amount of time, in milliseconds, that it takes to create a physical connection to the database.*
	Driver Version	The driver class name of the JDBC driver used to create database connections.*
	time_stamp	The data and time that the data in the row was last updated.*
	Connection Activity and Cache Statement Performance Trends for Module: (Module Name)	Shows connection and open cursor data for the connection.
	High Active Conns	-- Traces the highest number of active database connections in the instance of the data source since the data source was instantiated.
	Active Conns	-- Traces the number of connections currently in use by applications.
	Current Size	-- Traces the number of prepared and callable statements currently cached in the statement cache.
	Accesses/sec	-- Traces the cumulative, running count of the number of times that the statement cache was accessed.
	Misses/sec	-- Traces the number of times (per second) that a statement request could not be satisfied with a statement from the cache.
	Additions/sec	-- Traces the cumulative, running count of the number of statements added to the statement cache.
	Deletions/sec	-- Traces the cumulative, running count of statements discarded from the cache.
	Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero


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

Time Range

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




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

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

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

WLS ThreadPool Summary

View threadpool utilization, performance, and trend data for a specific WebLogic server.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data.
- Data Server** The name of the data server.

WLS Server Select the WebLogic server for which you want to see data.

Cluster The name of the cluster.

Thread Pool for Selected Server Table

Completed Requests	The number of completed requests in the priority queue.*
Delta Completed	The increase in the amount of completed requests (from the previous polling period to the current polling period).
Throughput	The mean number of requests completed per second.*
Hogging Threads	The threads that are currently being hogged by a request. These threads will either be declared as stuck after the configured timeout or will be returned to the pool. The self-tuning mechanism will backfill if necessary.*
Queue Length	The number of pending requests, which consist of the total of internal system requests and user requests, in the priority queue.*
Pend User Requests	The number of pending user requests in the priority queue.*
Exec Threads Idle	The number of idle threads in the pool.*
Exec Threads Total	The total number of threads in the pool.*
Standby Threads	The number of threads in the standby pool.*
Suspended	Indicates if the RequestManager is suspended. A suspended manager will not dequeue work and dispatch threads until it is resumed.*
time_stamp	The date and time the data in the row was last updated.

Trend Graph Shows connection and open cursor data for the connection.

Completed -- Traces the number of completed requests in the priority queue.

Throughput -- Traces the mean number of requests completed per second.


Hogging Threads-- Traces the number of threads that are currently being hogged by a request.


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


Base at Zero



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

Time Range

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



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

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

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

Work Manager

View server runtime data for all work managers on a specific WebLogic Server.

The screenshot shows the 'WLS Server - Work Manager' interface. At the top, there's a 'Summary' tab and a title bar with the date '17-Oct-2016 15:20' and a 'Data OK' status. Below the title bar, there are dropdown menus for 'Domain: examplesServer', 'DataServer: WLM-alpha-SLHOST-WIN7', 'WLS Server: examplesServer', and 'Cluster:'. The main area contains a table titled 'Server Runtime' with the following columns: Name, Application, Completed, Delta Compl..., Rate Compl..., Pending, and Stuck. The table lists various work managers and their associated applications, with numerical values in the 'Completed', 'Delta', 'Rate', 'Pending', and 'Stuck' columns. For example, 'DataRetirementWorkManager' has 0 in all these columns. 'JmsDispatcher' has 0 in 'Completed', 70,624 in 'Delta', and 0 in 'Rate', 'Pending', and 'Stuck'. The table is scrollable, and there are navigation arrows at the bottom.

Name	Application	Completed	Delta Compl...	Rate Compl...	Pending	Stuck
DataRetirementWorkManager		0	0	0.00	0	0
default	examples-multiDataSource-d...	0	0	0.00	0	0
default	esphere	0	0	0.00	0	0
default	examples-demo	0	0	0.00	0	0
default	jdbcRowSetsEar	0	0	0.00	0	0
default	SamplesSearchWebApp	0	0	0.00	0	0
default	ejb30	0	0	0.00	0	0
default	slweb	0	0	0.00	0	0
default	mejb	0	0	0.00	0	0
default	bea_wls9_async_response	0	0	0.00	0	0
default	examplesWebApp	0	0	0.00	0	0
default	examples-jms	0	0	0.00	0	0
default	examples-demoXA	0	0	0.00	0	0
default	webServicesJwsSimpleEar	0	0	0.00	0	0
default	mainWebApp	9	0	0.00	0	0
default	extServletAnnotationsEar	0	0	0.00	0	0
default	bea_wls_internal	0	0	0.00	0	0
default	examples-demoXA-2	0	0	0.00	0	0
default	asyncServletEar	0	0	0.00	0	0
default	jspSimpleTagEar	0	0	0.00	0	0
default	xmlBeanEar	0	0	0.00	0	0
default	stockEar	0	0	0.00	0	0
default	ejb20BeanMgedEar	0	0	0.00	0	0
direct		0	0	0.00	0	0
ImageWorkManager		0	0	0.00	0	0
JmsAsyncQueue		70,624	0	0.00	0	0
JmsDispatcher		0	0	0.00	0	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.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data.
- Data Server** The name of the data server.

WLS Server Select the WebLogic server for which you want to see data.

Cluster The name of the cluster.

Server Runtime Table

Name	The name of the work manager.
Application	The name of the application with which the work manager is associated.*
Completed	The number of requests that have been completed.*
Delta Completed	The increase in the amount of completed requests (from the previous polling period to the current polling period).
Rate Completed	The rate of completed requests (per second).
Pending	The number of requests waiting in the queue.*
Stuck	The number of threads that are "stuck."*
Expired	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
time_stamp	The date and time the data in the row was last updated.

Persistent Stores

View available utilization and performance data for all configurations on a specific domain.

← **Persistent Stores - Detail Tables** 18-Nov-2016 14:50 Data OK + ?

Domain: TestDomain

Persistent Store Runtime						
Name	Location	HealthState	CreateCo...	DeleteCount	ObjectCo...	Physic
WseeFileStore_auto_1	AdminServer	HEALTH_OK	1	1	9	
WseeJaxwsFileStore_auto_1	ManagedServer1	HEALTH_OK	2	2	21	
WseeJaxwsFileStore_auto_2	ManagedServer2	HEALTH_OK	2	2	21	

« »

Persistent Store Connection Runtime				
Name	Location	PersistentStoreRuntime	CreateCo...	DeleteCo...
weblogic.messaging.JMServer-1.body	ManagedServer1	FileStore-1	0	0
weblogic.messaging.WseeJmsServer_auto_1.body	AdminServer	WseeFileStore_auto_1	0	0
weblogic.transaction.tlog	ManagedServer2	_WLS_ManagedServer2	0	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.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

Domain Select the domain for which you want to view data.

Persistent Store Runtime Table

Name	The name of the configuration.
Location	The name of the WLS Server located on the specified connection.
Health State	The health state of the store.*
Create Count	The number of create requests issued to the persistent store.*
Delete Count	The number of delete requests issued by this persistent store.*
Object Count	The number of objects contained in the persistent store.*
Physical Write Count	The number of times the persistent store flushes its data to durable storage.*
Read Count	The number of read requests issued to this persistent store.*
Update Count	The number of update requests issued by this persistent store.*
time_stamp	The date and time the data in the row was last updated.

Persistent Store Connection Runtime Table

Name	The name of the configuration.
Location	The name of the WLS Server located on the specified connection.
Persistent Store Runtime	The name of the persistent store.
Create Count	The number of create requests issued by this connection.*
Delete Count	The number of delete requests issued by this connection.*
Object Count	The number of objects contained in this connection.*
Read Count	The number of read requests issued by this connection.*
Update Count	The number of update requests issued by this connection.*
time_stamp	The date and time the data in the row was last updated.

Application Views View

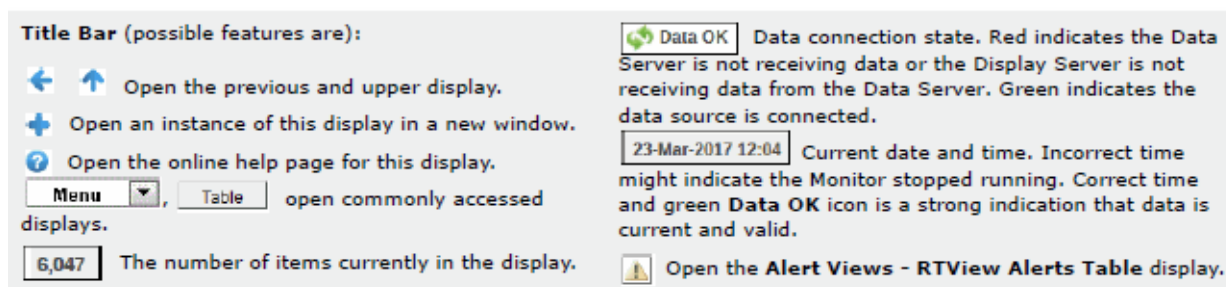
These displays present several views of performance metrics for applications on clusters and a particular WebLogic server.

- **"Cluster Apps Table"**: View performance and utilization metrics for all applications on a particular cluster, or for all applications on all clusters.
- **"Cluster App Summary"**: View session information for a particular application in graph, heatmap, and trend formats.
- **"Server Apps Heatmap"**: Shows a heatmap view of the status and alerts of all applications within a specific WebLogic server.
- **"Server Apps Summary"**: Track performance, utilization, and trend data for all applications on a single WebLogic server.
- **"Server Apps Trends"**: View trend data for a single application on a particular WebLogic server.
- **"App Components Heatmap"**: Provides a heatmap view of the status and alerts of all application components contained within each application on a particular WebLogic server.
- **"App Components Summary"**: View performance, utilization, and trend data for all application components on a single WebLogic Server.

Cluster Apps Table

View performance and utilization metrics for all applications on a particular cluster, or for all applications on all clusters.

Domain	Cluster	Application	Sessions	New	Sessions	Invocations	Exec Time
TestDomain		bea_wls9_async_response	0	0	0.00	0.00	0.00
TestDomain		bea_wls_deployment_internal	0	0	0.00	0.00	0.00
TestDomain		bea_wls_internal	0	0	0.00	0.00	0.00
TestDomain		bea_wls_management_inter...	0	0	0.00	0.00	0.00
TestDomain	Cluster-114-2		0	0			
TestDomain	TestCluster-114	bea_wls9_async_response	0	0	0.00	0.00	0.00
TestDomain	TestCluster-114	bea_wls_cluster_internal	0	0	0.00	0.00	0.00
TestDomain	TestCluster-114	bea_wls_deployment_internal	0	0	0.00	0.28	0.14
TestDomain	TestCluster-114	bea_wls_internal	0	0	0.00	0.00	0.00
TestDomain	TestCluster-114	wlm	0	0	0.00	0.00	0.00
Win7Server		bea_wls9_async_response	0	0	0.00	0.00	0.00
Win7Server		bea_wls_internal	0	0	0.00	0.00	0.00
Win7Server		esphere	0	0	0.00	0.00	0.00
Win7Server		myocm_rtvdata	0	0	0.00	0.00	0.00
Win7Server		ocmon	0	0	0.00	0.00	0.00
Win7Server		ocmon_rtvdata	0	0	0.00	0.00	0.00
Win7Server		ocmon_rtvquery	0	0	0.00	0.00	0.00
Win7Server		wlm	0	0	0.00	0.00	0.00
Win7Server		wlm_rtvdata	0	0	0.00	0.00	0.00
Win7Server		wlm_rtvquery	0	0	0.00	0.00	0.00
examplesServer		SamplesSearch/WebApp	0	0	0.00	0.00	0.00
examplesServer		asyncServletEar	0	0	0.00	0.00	0.00
examplesServer		bea_wls9_async_response	0	0	0.00	0.00	0.00
examplesServer		bea_wls_internal	0	0	0.00	0.00	0.00
examplesServer		ejb20BeanMgedEar	0	0	0.00	0.00	0.00
examplesServer		ejb30	0	0	0.00	0.00	0.00
examplesServer		esphere	0	0	0.00	0.00	0.00
examplesServer		examplesWebApp	0	0	0.00	0.00	0.00
examplesServer		extServletAnnotationsEar	0	0	0.00	0.00	0.00
examplesServer		jdbcRowSetsEar	0	0	0.00	0.00	0.00
examplesServer		jspSimpleTanEar	0	0	0.00	0.00	0.00



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

Fields and Data

This display includes:

Domain	Select the domain for which you want to view data.
Cluster	The name of the cluster.
Cluster Count	Displays the total number of clusters listed in the table.

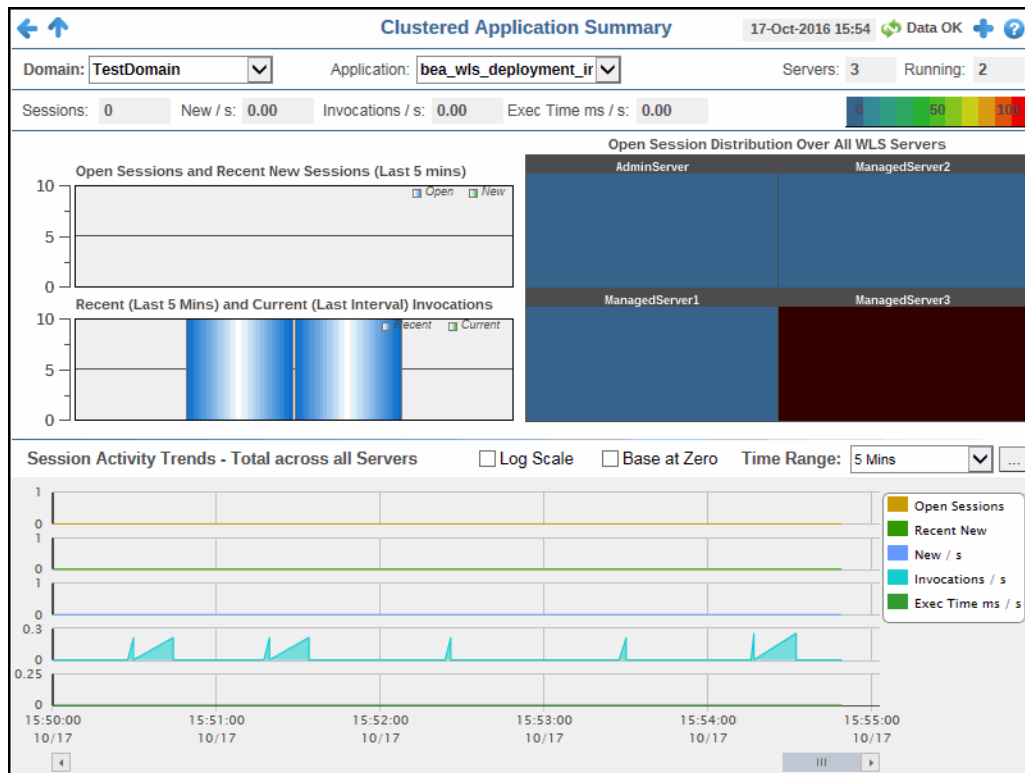
WebLogic Cluster Applications Table

Click a row to view metrics for a single topic in the ["Cluster App Summary"](#) display.

Domain	The name of the domain.
Cluster	The name of the cluster.
Application	The name of the application.
Sessions Open	The number of open sessions on the application.
New Sessions	The number of new sessions since the last polling update.
Session New/sec	The rate of new sessions (per second).
Invocations/sec	The rate of invocations (per second).
Exec Time ms/sec	The rate of execution time in milliseconds (per second).
Server Count	The total number of existing servers on the application.
Running Count	The total number of running servers on the application.
Percent Not Running	The percentage of servers that are not running.
time_stamp	The date and time this row of data was last updated.

Cluster App Summary

View session information for a particular application in graph, heatmap, and trend formats.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.


- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:


Domain Select the domain containing the application for which you want to view data.


Application	Select the application for which you want to view data.
Servers	Displays the total number of servers on the application.*
Running	Displays the total number of servers running on the application.*
Sessions	Displays the number of open sessions.*
New/s	Displays the rate of new sessions being opened (per second).*
Invocations/s	Displays the rate of invocations (per second).*
Exec Time ms/s	Displays the rate of execution time in milliseconds (per second).*
Open Sessions and Recent New Sessions (Last 5 mins) bar graph	Displays the currently open sessions and the sessions created in the last 5 minutes.
Recent (Last 5 Mins) and Current (Last Interval) Invocation bar graph	Displays the number of recent invocations (last 5 minutes) and the number of current invocations (created since the last polling interval).
Open Session Distribution Over All WLS Servers heat map	Displays the number of open sessions for each WLS server in heatmap form based on the color gradient bar  .
Trend Graph	Shows connection and open cursor data for the connection. Open Sessions -- Traces the number of open sessions. Recent New -- Traces the number of newly created sessions. New/s -- Traces the number of sessions created per second. Invocations/s -- Traces the number of invocations per second. Exec Time ms/s -- Traces the execution time in milliseconds per second.
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero



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

Time Range

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




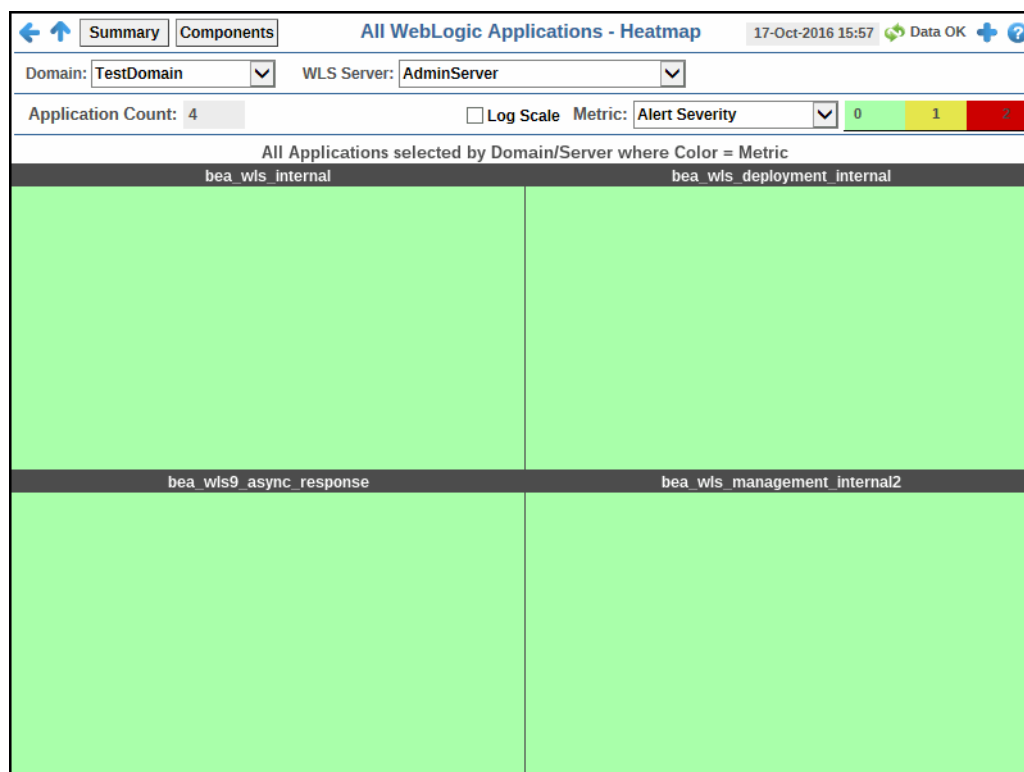
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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








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




Server Apps Heatmap

This display provides a heatmap view of the status and alerts of all applications within a specific WebLogic server. The heatmap is organized so that each rectangle represents a collection contained within a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each collection or click on a rectangle to drill-down to the “[Server Apps Trends](#)” display and view metrics for that particular collection. You can click the table icon  in this display to navigate to the “[Server Apps Summary](#)” display.



Title Bar (possible features are):

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

-  Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
-  23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
-  Open the **Alert Views - RTView Alerts Table** display.

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

Fields and Data

This display includes:

Domain	Select the domain containing the WebLogic Server for which you want to view data.
WLS Server	Select the WebLogic server for which you want to view data.
Application Count	The total number of applications on the server.
Log Scale	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

Metric

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by connection, where each rectangle represents a collection. Mouse-over any rectangle to display the current values of the metrics for the collection. Click on a rectangle to drill-down to the associated ["Server Apps Summary"](#) display for a detailed view of metrics for that particular collection.

Alert Severity

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



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


2 -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

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

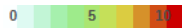
0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

Alert Count


The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

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


Open Sessions

The total number of open sessions in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **WlsOpenSessionsHigh**, which is **10**. The middle value in the gradient bar indicates the middle value of the range (the default is **5**).


**Open Sessions/
sec**

The number of sessions opened per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


Invocations/sec

The number of invocations per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


**Invocation Total
Count**

The total number of invocations in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


Exec Time/sec

The execution time per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Exec Time Total

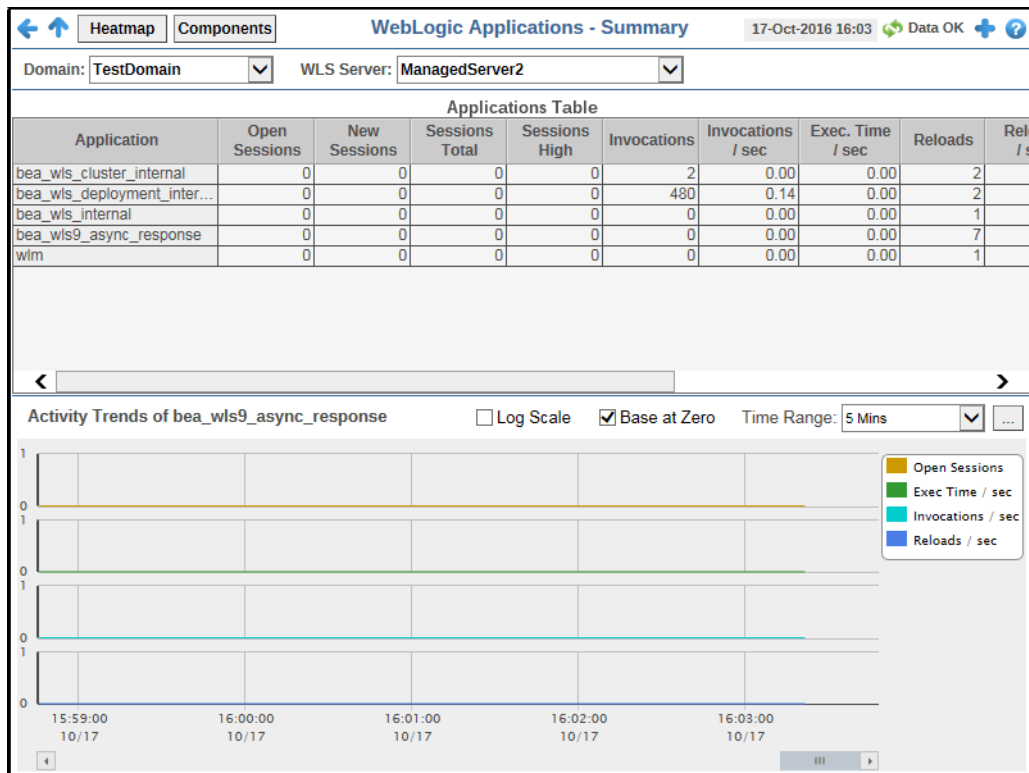
The total amount of execution time in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Reload Total
Count**

The total reload count in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Server Apps Summary

Track performance, utilization, and trend data for all applications on a single WebLogic server.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic Server for which you want to view data.

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

Applications Table

This table describes all topics on the selected server. Click a row to view metrics for a single topic in the [“Server Apps Trends”](#) display.

Application	The name of the application.
Open Sessions	The number of open sessions.
New Sessions	The number of new sessions.
Sessions Total	The total number of sessions.
Invocations	The number of invocations.
Invocations/sec	The number of invocations per second.
Exec. Time/sec	The rate of execution time in milliseconds per second.
Reloads	The number of reloads.
Reloads / sec	The rate of reloads (per second).
Status	The status of the application.
Deployment State	The current status of the application's deployment.
Expired	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
time_stamp	The date and time this row of data was last updated.

Activity Trends of <application>

Shows the following:

- Open Sessions** -- Traces the total number of open sessions in the application.
- Exec Time/sec** -- Traces the execution time per second in the application.
- Invocations/sec** -- Traces the number of invocations per second.
- Reloads/sec** -- Traces the number of reloads per second.


Log Scale


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


Base at Zero

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

Time Range

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




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

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

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

Server Apps Trends

View trend data for a single application on a particular WebLogic server.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.


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


Fields and Data


This display includes:

Domain

Select the domain containing the WebLogic Server for which you want to view data.

- WLS Server** Select the WebLogic server containing the application for which you want to view data.
- Application** Select the application for which you want to view data.
- Metric Trends for:** Shows message data for the selected collection.
- Open Sessions** -- Traces the total number of open sessions in the application.
 - New Sessions** -- Traces the number of new sessions in the application.
 - Invocations/sec** -- Traces the number of invocations per second in the application.
 - Thread Requests** -- Traces the number of thread requests in the application.
 - Hogging Threads** -- Traces the number of hogging threads in the application.
 - JVM CPU %** -- Traces the JVM CPU percentage in the application.
 - Max Heap MB** -- Traces the max heap used, in megabytes, in the application.
 - Used Heap MB** -- Traces the used heap, in megabytes, in the application.
- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



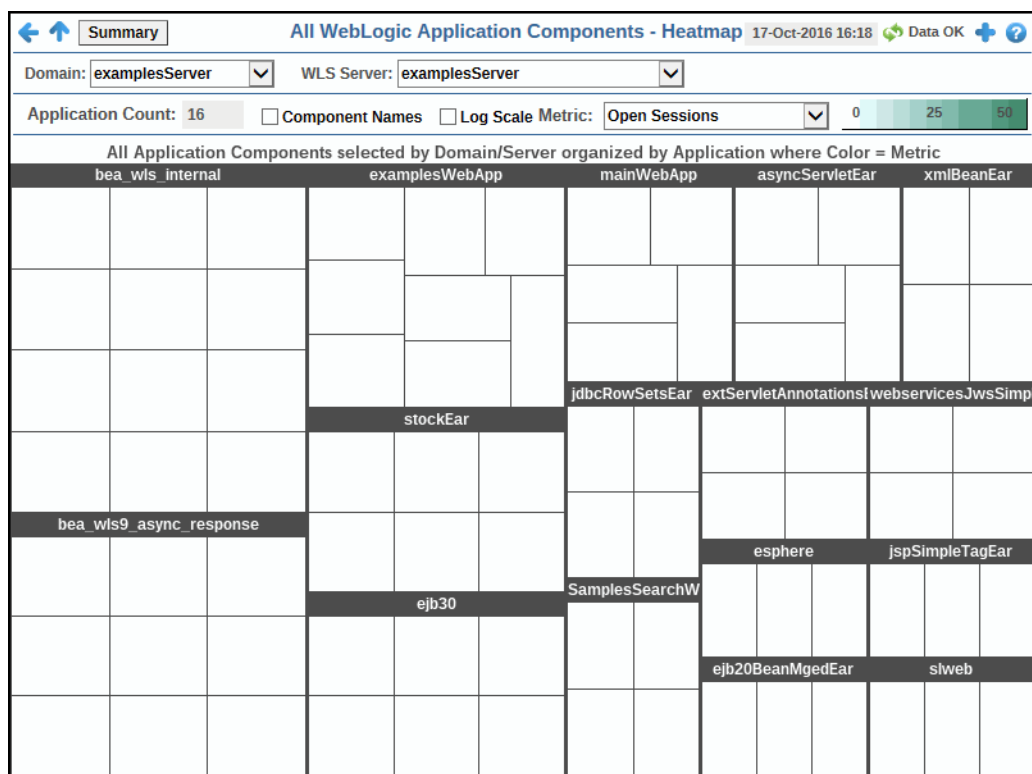
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

App Components Heatmap

This display provides a heatmap view of the status and alerts of all application components contained within each application on a particular WebLogic server. The heatmap is organized so that each rectangle represents a collection contained within a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each collection or click on a rectangle to drill-down to the “[WebLogic Single Application Summary](#)” display and view metrics for that particular collection. You can click the **Summary** button in this display to navigate to the “[Server Apps Summary](#)” display.



Title Bar (possible features are):

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

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

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

Open the Alert Views - RTView Alerts Table display.


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


Fields and Data


This display includes:


- Domain** Select the domain containing the WebLogic Server for which you want to view data.
- WLS Server** Select the WebLogic server for which you want to view data.
- Application Count** The total number of application components in the WebLogic server.
- Component Names** Select this check box to display the names of the application components in the heatmap.
- Log Scale** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

Metric Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by connection, where each rectangle represents a collection. Mouse-over any rectangle to display the current values of the metrics for the collection. Click on a rectangle to drill-down to the associated [“Server Apps Trends”](#) display for a detailed view of metrics for that particular collection.

Open Sessions The number of open sessions in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Open Sessions/sec The number of open sessions per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Invocations/sec The number of invocations per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Invocation Total Count The total number of invocations in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Execution Time Total

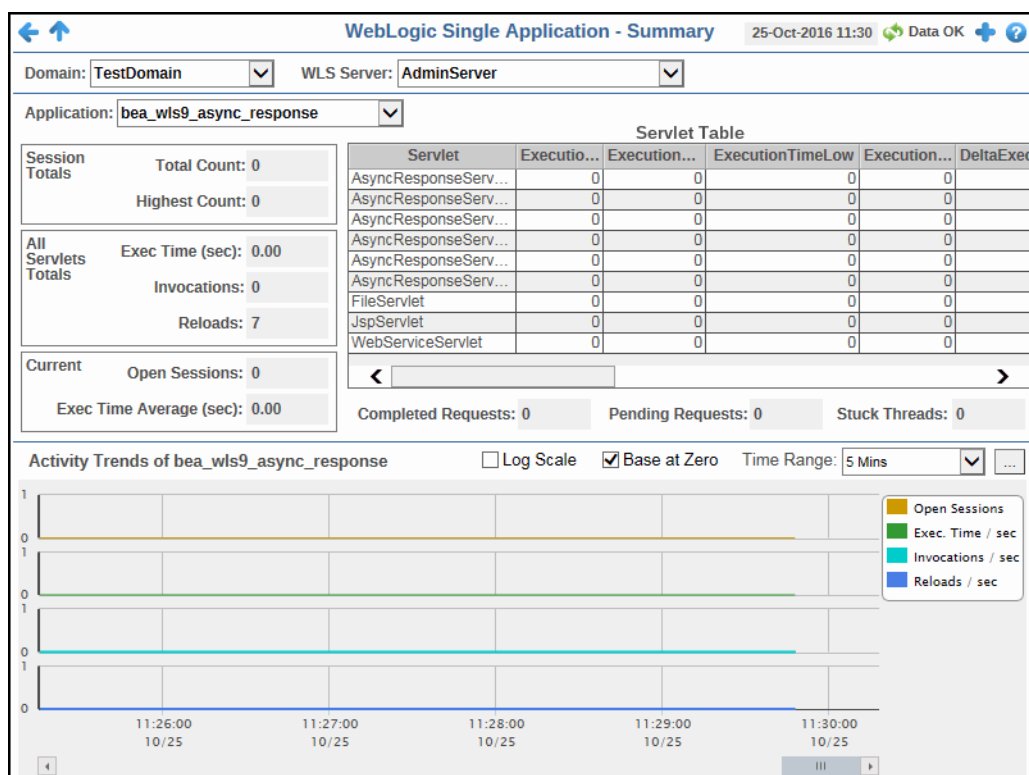
The total execution time in a given item (index) associated with the rectangle. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.


Reload Total Count

The total number of reloads in a given item (index) associated with the rectangle. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

WebLogic Single Application Summary

View performance, utilization, and trend data for a single application component on a single WebLogic server. This display is only accessible by clicking on an application component in the "App Components Heatmap".





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

Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic server for which you want to view data.
- WLS Server** Select the WebLogic server containing the application for which you want to view data.
- Application** Select the application for which you want to view data.

Sessions Totals

- Total Count** The total number of sessions on the application.
- Highest Count** The highest total number of sessions on the application.

All Servlet Totals

- Exec Time (sec)** The amount of time all invocations of all servlets have executed since being created.*
- Invocations** The total count of the times all servlets have been invoked.*
- Reloads** The total count of the number of times all servlets have been reloaded.*

Current

- Open Sessions** The number of currently open sessions.
- Exec Time Average (sec)** The current amount of time invocations of the servlet are being executed.*

Servlet Table

This table describes metrics for each servlet on the application.

Servlet	The name of the servlet.
Execution Time Average	The average amount of time all invocations of the servlet have executed since being created.*
Execution Time High	The amount of time the single longest invocation of the servlet has executed since being created.*
Execution Time Low	The amount of time the single shortest invocation of the servlet has executed since being created.*
Execution Time Total	The total amount of time all invocations of the servlet have executed since being created.*
Delta Execution Time	The increase in the execution time (from the previous polling period to the current polling period).
Rate Execution Time Total	The total time taken to execute requests per second.
Invocation Total Count	The total count of the times this servlet has been invoked.*
Delta Invocation Total Count	The increase in the amount of invocations (from the previous polling period to the current polling period).
Rate Invocation Total Count	The total number of invocations per second.*
Reload Total Count	The total count of the number of times this servlet has been reloaded.*
Delta Reload Total Count	The increase in the amount of reloads (from the previous polling period to the current polling period).
Rate Reload Total Count	The number of times this servlet has been reloaded per second.
Expired	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
time_stamp	The date and time this row of data was last updated.
Completed Requests	The total number of completed requests on the application.
Pending Requests	The total number of pending requests on the application.
Stuck Threads	The total number of stuck threads on the application.
Metric Trends for: <application> Graph	Shows message data for the selected collection. Open Sessions -- Traces the total number of open sessions in the application. Exec. Time/sec -- Traces the number of executions per second in the application. Invocations/sec -- Traces the number of invocations per second in the application. Reloads/sec -- Traces the number of reloads per second in the application.


Log Scale


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


Base at Zero

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

Time Range

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



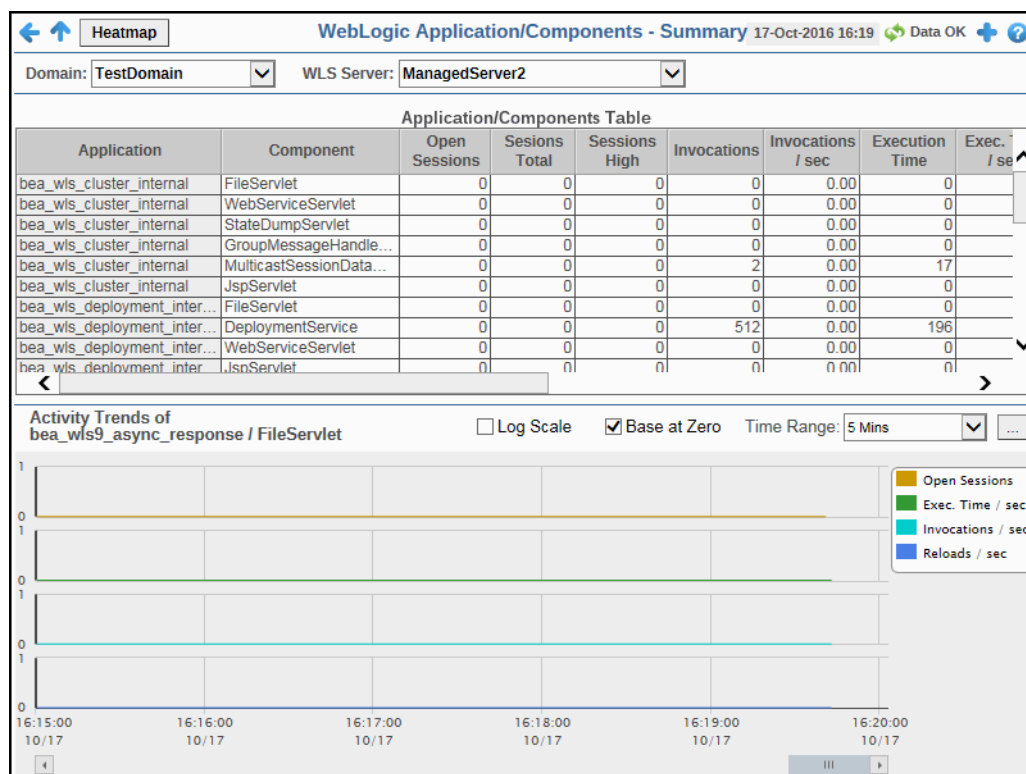
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

App Components Summary

View performance, utilization, and trend data for all application components on a single WebLogic Server.



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

Fields and Data

This display includes:

Domain Select the domain containing the WebLogic server for which you want to view data.

WLS Server Select the WebLogic server containing the application for which you want to view data.

Applications/Components Table

This table describes metrics for each application/component combination.

Application	The name of the application.
Component	The name of the component.
Open Sessions	The current total of open sessions.*
Sessions Total	The total number of sessions that have been opened.*
Sessions High	The highest number of sessions opened at one time.*
Invocations	The total number of invocations of the application or component.*
Invocations/sec	The rate of invocations (per second) of the application or component.*
Execution Time	The amount of time (in milliseconds) it took to execute the last invocation.*
Exec. Time Low	The average amount of time the single shortest invocation of the component has executed since it was most recently deployed.*
Exec. Time High	The average amount of time the single longest invocation of the component has executed since it was most recently deployed.**
Exec. Time Average	The average amount of time it took to execute all invocations of the component since it was most recently deployed.**
Reloads	The total number of times the WebLogic server has reloaded the component since it was last deployed.*
Reloads/sec	The rate of reloads (per second).*
Status	The status of the component.*
Expired	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
time_stamp	The date and time this row of data was last updated.

Metric Trends for: <application> Graph

Shows message data for the selected collection.

Open Sessions -- Traces the total number of open sessions in the application.

Exec. Time/sec -- Traces the number of executions per second in the application.

Invocations/sec -- Traces the number of invocations per second in the application.

Reloads/sec -- Traces the number of reloads per second in the application.


Log Scale


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


Base at Zero

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

Time Range

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



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

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

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

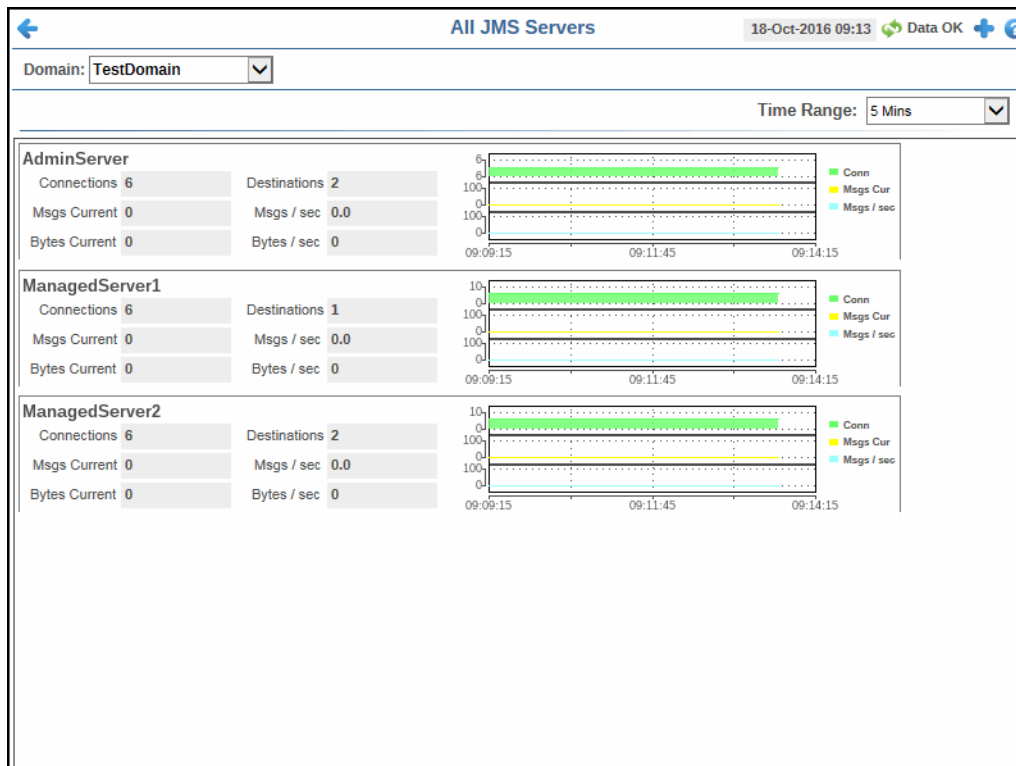
JMS Servers View

These displays present several views of performance metrics for JMS Servers. This view contains the following:

- **"JMS Servers Grid"**: Track performance, utilization, and trend data for all JMS servers on a single domain.
- **"JMS Server Summary"**: View performance, utilization, and trend data for a particular JMS server.
- **"JMS Metric Trends"**: View activity trends for a particular JMS server.
- **"JMS Server Detail"**: Track performance and utilization metrics for all JMS servers on a single WebLogic server.

JMS Servers Grid

Track performance, utilization, and trend data for all JMS servers on a single domain.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic servers for which you want to view data.
- Time Range** Select the period of time from this drop down for which you want to view data.

Server Trend Graphs

Trend graphs and the following metrics displays for servers within the selected domain.

Conn -- Traces the total number of connections on the server.

Msgs Cur -- Traces the current number of incoming/outgoing messages on the server.

Msgs/sec -- Traces the number of incoming/outgoing messages per second on the server.

Connections The number of the connections on the server.

Msgs Current The current number of messages stored on the JMS server. This number does not include the pending messages.*

Bytes Current The current number of bytes stored on the JMS server.*

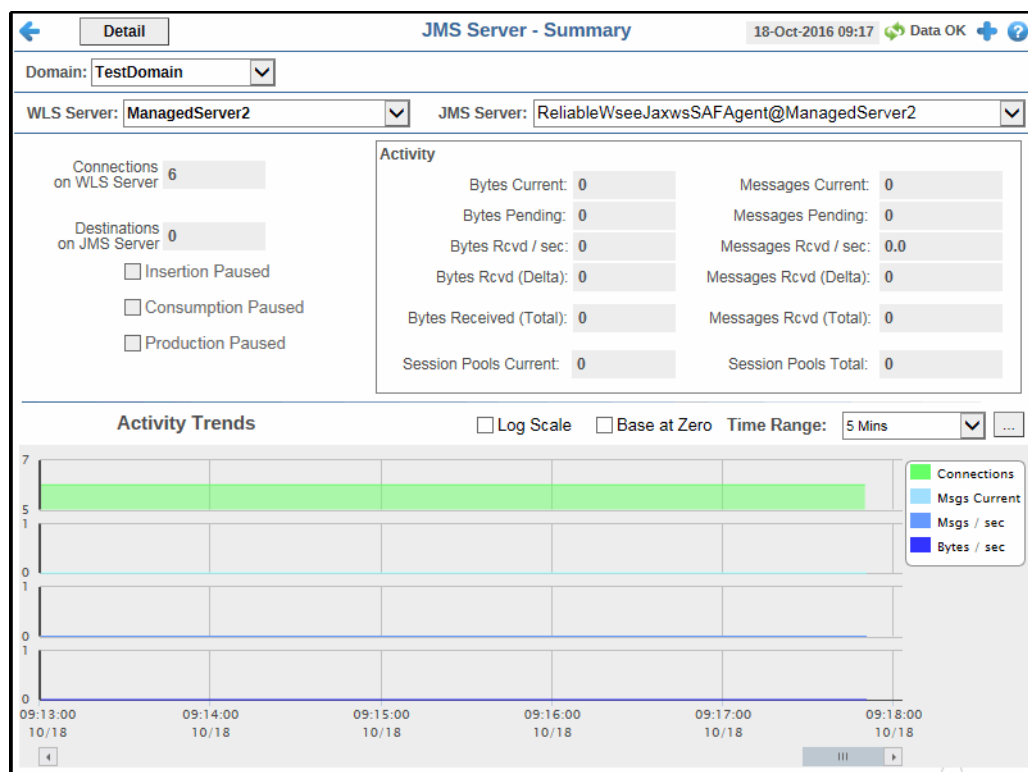
Destinations The current number of destinations stored on the JMS server.*


Msgs/sec The rate of messages received (per second) on the JMS server.*

Bytes/sec The rate of bytes received (per second) on the JMS server.*

JMS Server Summary

View performance, utilization, and trend data for a particular JMS server.





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

Fields and Data

This display includes:

Domain	Select the domain containing the WebLogic server for which you want to view data.
WLS Server	Select the WebLogic server containing the JMS server for which you want to view data.
JMS Server	Select the JMS server for which you want to view data.
Connections on WLS Server	The total number of connections on the WebLogic server.
Destinations on JMS Server	The number of destinations on the selected JMS server.
Insertion Paused	The current insertion paused state of the JMS server.*
Consumption Paused	The current consumption paused state of the JMS server.*
Production Paused	The current production paused state of the JMS server.*

Activity Region


Bytes Current	The current number of bytes stored on the JMS server.*
Bytes Pending	The current number of bytes pending (unacknowledged or uncommitted) on the JMS server.*
Bytes Rcvd/sec	The number of bytes received per second.
Bytes Rcvd (Delta)	The increase in the amount of bytes received (from the previous polling period to the current polling period).*
Bytes Rcvd (Total)	The total number of bytes received since the JMS server was last restarted.*


	Alert Session Pools Current	The current number of session pools instantiated on the server.*
	Messages Current	The current number of messages stored on the JMS server.*
	Messages Pending	The current number of messages pending (unacknowledged or uncommitted) on the JMS server.*
	Messages Rcvd/sec	The number of messages received per second since the server was last restarted.*
	Messages Rcvd (Delta)	The increase in the amount of messages received (from the previous polling period to the current polling period).
	Messages Rcvd (Total)	The total number of messages received since the server was last restarted.*
	Session Pools Total	The total number of session pools on the server.*
Activity Trends	Shows message data for the selected collection.	
	Connections -- Traces the total number of connections on the server.	
	Msgs Current -- Traces the number of current messages.	
	Msgs/sec -- Traces the number of incoming/outgoing messages per second.	
	Bytes/sec -- Traces the number incoming/outgoing bytes per second.	
	Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero


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

Time Range

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




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

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

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

JMS Metric Trends

View activity trends for a particular JMS server.



Title Bar (possible features are):

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

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.


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


Fields and Data


This display includes:

Domain

Select the domain containing the WebLogic server for which you want to view data.

WLS Server	Select the WebLogic server containing the JMS server for which you want to view data.
JMS Server	Select the JMS server for which you want to view data.
Activity Trends	Shows message data for the selected collection. <ul style="list-style-type: none"> Connections -- Traces the total number of connections on the server. Msgs Current -- Traces the number of current messages stored on the server. Msgs/sec -- Traces the number of messages stored per second. Bytes/sec -- Traces the number bytes received on the server per second. Thread Requests -- Traces the number of thread requests. Queue Length -- Traces the length of the queue. CPU % -- Traces the percentage of CPU used. Used Heap Mb -- Traces the amount of heap used, in megabytes.
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
Base at Zero	When this option is checked, zero is set as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click the  button.



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

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

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

WLS Server Select the WebLogic server containing the JMS server for which you want to view data.

Cluster The name of the cluster associated with the WebLogic Server.

JMS Server Runtime Table

This table describes all runtime statistics on the selected server. Click a row to view metrics for a single server in the ["JMS Server Summary"](#) display.

Name	The name of the JMS Server.
Destinations Current	The current number of destinations for the JMS server.*
Destinations Total	The total number of destinations instantiated on the JMS server since the last polling update.*
Messages Received	The total number of messages received by the JMS server since the last polling update.*
Msgs Rcvd Delta	The increase in the amount of messages received (from the previous polling period to the current polling period).
Msgs Rcvd per sec	The rate of messages received (per second) by the server.
Messages Current	The current number of messages stored on the JMS server.*
Messages Pending	The current number of pending messages stored on the JMS server.*
Messages High	The highest number of total messages stored by the JMS server since the last polling update.*
Bytes Received	The number of bytes received by the server since the last polling update.
Bytes Rcvd Delta	The increase in the amount of bytes received (from the previous polling period to the current polling period).
Bytes Rcvd per sec	The rate of bytes received (per second) by the server.*
Bytes Current	The current number of bytes stored on the JMS server.*
Bytes Pending	The current number of bytes pending that are stored on the JMS server.*
Bytes High	The largest number of bytes stored on the JMS server since the last polling update.*
Health State	The current health state of the JMS server.*
Consumption Paused	The current consumption paused state of the JMS server. When checked, consumption is paused.*
Cons Paused Enabled	The current consumption paused state on the JMS server (enabled/disabled).*
Insertion Paused	The current insertion paused state of the JMS server.*
Ins Paused State	The current insertion paused state of the JMS server (enabled/disabled).*

Production Paused	The current production paused state of the JMS server.*
Prod Paused State	The current production paused state of the JMS server (enabled/disabled).*
Session Pools Current	The current number of session pools instantiated on the JMS server.*
Session Pools High	The highest number of session pools instantiated on the server since the last polling update.*
Session Pools Total	The total number of session pools instantiated on the JMS server since the last polling update.*
Destinations	The number of destinations instantiated on the JMS server since the last polling interval.
Transactions	The number of transactions that exist on the JMS server.*
Transactions Pending	The number of pending transactions that exist on the JMS server.*
Server Runtime	The name of the JMS Server.*
time_stamp	The date and time this row of data was last updated.

JMS Runtime Table

This table describes performance metrics about the connectivity of the JMS Server.

Connections Current	The current number of connections to the WebLogic server.*
Connections High	The highest number of connections made to the WebLogic server since the last polling update.*
Connections Total	The total number of connections made to the WebLogic server since the last polling update.*
Connections Delta	The increase in the amount of connections (from the previous polling period to the current polling period).
JMS Servers Current	The number of JMS servers currently deployed on the WebLogic server.*
JMS Servers High	The highest number of JMS servers that were deployed on the WebLogic server since the server was started.*
JMS Servers Total	The total number of JMS servers that were deployed on the WebLogic server since the server was started.*
HealthState	The current state of health of the JMS service.*
JMSServers	The number of JMS servers deployed on the WebLogic server.*
Connections	The number of connections to the WebLogic server.*
JMSPooledConnections	The number of JMS pooled connectons on the server.
ServerRuntime	The name of the server.
time_stamp	The date and time this row of data was last updated.


JMS Destinations View

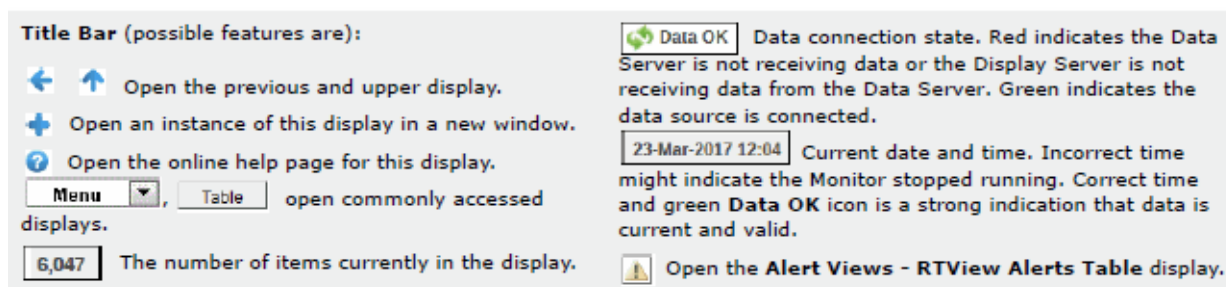
These displays present several views of performance metrics for JMS Destinations. This view contains the following:

- **"Totals by Destination"**: View performance, utilization, and trend data for a particular destination.
- **"Detail by Destination"**: A tabular view that allows you to view performance and utilization metrics for destinations for each JMS Server, as well as destination metrics across all JMS Servers.
- **"Distribution by Server"**: Shows performance and utilization metrics and trends for a destination on a single domain.
- **"Destination by Server"**: Shows metrics and trends for a particular destination on a JMS Server.
- **"Detail by Server"**: Shows metrics for all destinations on one or all modules on a JMS Server.

Totals by Destination

View performance, utilization, and trend data for a particular destination.





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

Fields and Data

This display includes:

Domain	Select the domain containing the WebLogic server for which you want to view data.
Type	Select the type of destination (Queue or Topic) for which you want to view data, or select All Types .
Module	Select the module containing the destination for which you want to view data.
Destination	Select the destination for which you want to view data.
Type	The type of destination (queue or topic) selected.
Consumers	The current number of consumers accessing the destination.
Consumers (High)	The highest number of consumers accessing the destination since the last polling update.
Avg Msg Bytes Rcvd	The average number of bytes received in the destination since the last polling update.
Activity Region	

Bytes

Bytes Current	The current number of bytes stored in the destination.*
Bytes High	The highest number of bytes stored in the destination since the last polling update.*
Bytes Pending	The current number of pending bytes stored in the destination.*
Bytes Rcvd (Delta)	The increase in the amount of bytes received (from the previous polling period to the current polling period).
Bytes Rcvd (Total)	The number of bytes received in this destination since the last polling update.*


Messages


	Messages Current	The current number of messages in the destination.*
	Messages High	The highest number of messages in the destination since the last polling update.*
	Messages Pending	The current number of pending messages in the destination.*
	Messages Rcvd (Delta)	The increase in the amount of messages received (from the previous polling period to the current polling period).
	Messages Rcvd (Total)	The total number of messages in the destination since the last polling update.*
Activity Trends	Shows message data for the selected collection.	
	Consumers -- Traces the total number of consumers on the destination.	
	Msgs Current -- Traces the number of current messages.	
	Msgs/sec -- Traces the number of messages received per second in the destination.	
	Bytes/sec -- Traces the number of bytes received per second in the destination.	
	Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero


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

Time Range

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



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

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

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

Module Select the module containing the destination for which you want to view data.

Destination Select the destination for which you want to view data.

Destination Statistics for each JMS Server Table

Each row in this table describes all destination statistics for a particular JMS Server. Click a row to view additional metrics in the ["Destination by Server"](#) display.

JMS Server	The name of the JMS Server.
Type	The type of destination (queue or topic).
Bytes Current	The current number of bytes stored in the destination.*
Bytes High	The highest number of bytes stored in the destination since the last polling update.*
Bytes Pending	The current number of pending bytes stored in the destination.*
Bytes Rcvd	The number of bytes received in this destination since the last polling update.*
Rcvd Delta	The increase in the amount of bytes received (from the previous polling period to the current polling period).
Rcvd Rate	The rate of bytes received (per second) into the destination.*
Consumers Current	The current number of consumers accessing the destination.*
Consumers High	The highest number of consumers accessing the destination since the last polling update.*
Consumers Total	The total number of consumers accessing the destination since the last polling update.*
Msgs Current	The current number of messages in the destination.*
Msgs Deleted	The number of messages that have been deleted from the destination.*
Msgs High	The highest number of messages in the destination since the last polling update.*
Msgs Moved	The number of moved messages in the destination since the last polling update.*
Msgs Pending	The current number of pending messages in the destination.*
Msgs Rcvd	The number of messages received in the destination since the last polling update.*
Rcvd Delta	The increase in the amount of messages received (from the previous polling period to the current polling period).
Rcvd Rate	The rate of messages received (per second) in the destination since the last polling update.*
Msgs Thresh Time	The amount of time in the threshold condition since the last polling update.

Bytes Thresh Time	The amount of time in the threshold condition since the last polling update.
Paused	Indicates whether or not the destination is paused at the current time.*
Consumption Paused	Indicates the consumption paused state of the destination.
Consumption Paused State	The current consumption paused state of the destination.
Production Paused	Indicates the production paused state of the destination.
Production Paused State	The current production paused state of the destination.
Insertion Paused	Indicates the insertion paused state of the destination.
Insertion Paused state	The current insertion pause state of the destination.
State	The current health state of the destination.
time_stamp	The date and time this row of data was last updated.

Totals across All Servers Table

This table describes destination metric totals across all servers.

Bytes Current	The current number of bytes stored across all servers.
Bytes High	The highest number of bytes on any one of the servers at one time since the last polling update.*
Bytes Pending	The total number of bytes pending on all servers.*
Bytes Rcvd	The number of bytes received across all servers.*
Rcvd Delta	The increase in the amount of bytes received (from the previous polling period to the current polling period).
Rcvd Rate	The rate of messages received (per second) by all servers.
Consumers Current	The number of current consumers across all servers.
Consumers High	The highest number of consumers at one time across all servers.*
Consumers Total	The total number of consumers across all servers.*
Msgs Current	The current number of messages stored across all destinations on the all servers.*
Msgs Deleted	The number of deleted messages across all destinations on all servers.*
Msgs High	The highest number of total messages stored across all destinations for all servers.*
Msgs Moved	The number of messages across all destinations for all servers that were moved.*
Msgs Pending	The current number of pending messages across all servers.

Msgs Rcvd

The total number of messages received across all servers.

Rcvd Delta

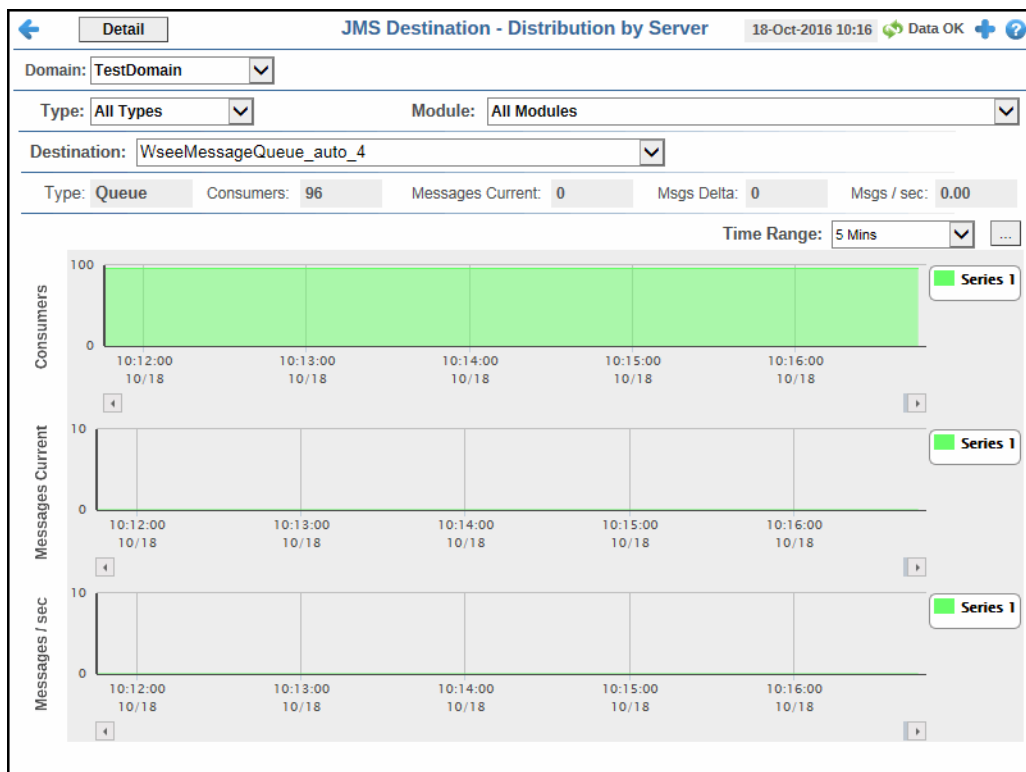
The increase in the amount of messages received (from the previous polling period to the current polling period).

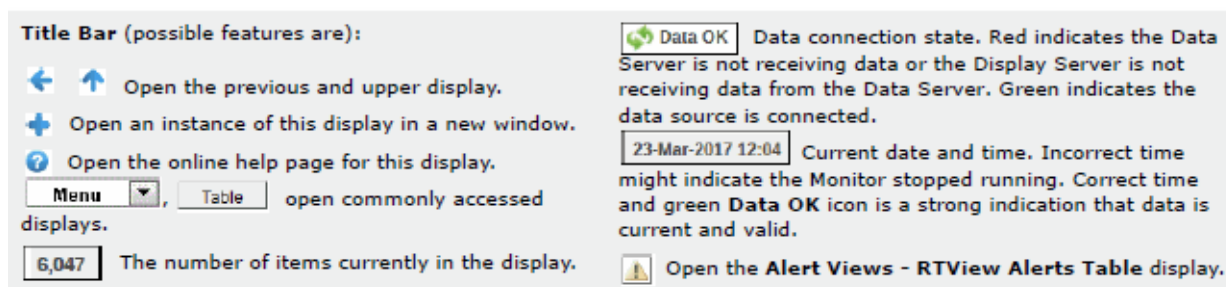
Rcvd Rate

The rate of messages received (per second) across all servers.

Distribution by Server

Track performance and utilization metrics and trends for a destination on a single domain.





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

Fields and Data

This display includes:

Domain	Select the domain containing the destination for which you want to view data.
Type	Select the type of destination (queue or topic) for which you want to view data, or select All Types .
Module	Select the module containing the destination for which you want to view data.
Destination	Select the destination for which you want to view data.
Type	The type of destination (queue or topic).
Consumers	The number of consumers.
Msgs Current	The current number of messages on the destination.*
Msgs Delta	The increase in the amount of messages received (from the previous polling period to the current polling period).
Msgs/sec	The rate of messages received (per second) by the destination.

**Server Trend
Graphs**


Trend graphs and the following metrics displays for servers within the selected domain.


Consumers -- Traces the total number of consumers on the server.


Messages Current -- Traces the current number of incoming/outgoing messages on the server.



Messages/sec -- Traces the number of incoming/outgoing messages per second on the server.

Time Range

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




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

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

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

Destination by Server

Track performance, utilization, and trend data for a particular destination on a single JMS Server.



Title Bar (possible features are):

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

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

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

Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

Domain Select the domain containing the WebLogic Server for which you want to view data.

WLS Server	Select the WebLogic server containing the JMS Server for which you want to view data.
JMS Server	Select the JMS Server containing the destination for which you want to view data
Type	Select the type of destination (queue or topic) for which you want to view data, or select All Types .
Module	Select the module containing the destination for which you want to view data.
Destination	Select the destination for which you want to view data.
Type	The type (queue or topic) of the selected destination.
Consumers	The current number of consumers accessing the destination.*
Consumer (High)	The highest number of consumers accessing the destination since the last polling update.*
Insertion Paused	Indicates the insertion paused state of the destination.*
Consumption Paused	Indicates the consumption paused state of the destination.*
Production Paused	Indicates the production paused state of the destination.*
Activity Region	
Bytes Current	The current number of bytes stored in the destination.*
Bytes High	The highest number of bytes stored in the destination since the last polling update.*
Bytes Pending	The current number of pending bytes stored in the destination.*
Bytes Rcvd (Delta)	The increase in the amount of bytes received (from the previous polling period to the current polling period).
Bytes Rcvd (Total)	The number of bytes received in this destination since the last polling update.*
Messages Current	The current number of messages in the destination.*
Messages High	The highest number of messages in the destination since the last polling update.*
Messages Pending	The current number of pending messages in the destination.*
Messages Rcvd (Delta)	The increase in the amount of messages received (from the previous polling period to the current polling period).
Messages Rcvd (Total)	The number of messages received in the destination since the last polling update.*
Activity Trends	Shows message data for the selected collection. Consumers -- Traces the total number of consumers accessing the destination. Msgs Current -- Traces the number of current messages. Msgs/sec -- Traces the number of messages received per second. Bytes/sec -- Traces the number bytes received per second.


Log Scale


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


Base at Zero

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

Time Range

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



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

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

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

Detail by Server

Track performance and utilization metrics for all destinations on a single JMS Server.

Destination Name	Type	Bytes Current	Bytes High	Bytes Pending	Bytes Rcvd	Rcvd Delta	Rcvd Rate
WseeBufferedRequestErrorQueue_auto_2	Queue	0	0	0	0	0	0.0
WseeBufferedRequestQueue_auto_2	Queue	0	0	0	0	0	0.0
WseeBufferedResponseErrorQueue_auto_2	Queue	0	0	0	0	0	0.0
WseeBufferedResponseQueue_auto_2	Queue	0	0	0	0	0	0.0

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.
- open commonly accessed displays.

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

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

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

Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

- Connection** Select the connection for which you want to view collection data.
- Domain** Select the domain containing the WebLogic Server for which you want to view data.

- WLS Server** Select the WebLogic server containing the JMS Server for which you want to view data.
- JMS Server** Select the JMS Server containing the destination for which you want to view data
- Type** Select the type of destination (queue or topic) for which you want to view data, or select **All Types**.
- Module** Select the module containing the destination for which you want to view data.

JMS Destinations on the selected JMS Server Table

This table describes all destinations on the selected JMS Server.

Destination Name	The name of the destination.
Type	The type of destination (queue or topic).
Bytes Current	The current number of bytes stored in the destination.*
Bytes High	The highest number of bytes stored in the destination since the last polling update.*
Bytes Pending	The current number of pending bytes stored in the destination.*
Bytes Rcvd	The number of bytes received in this destination since the last polling update.*
Rcvd Delta	The increase in the amount of bytes received (from the previous polling period to the current polling period).
Rcvd Rate	The rate of bytes received (per second) into the destination.*
Consumers Current	The current number of consumers accessing the destination.*
Consumers High	The highest number of consumers accessing the destination since the last polling update.*
Consumers Total	The total number of consumers accessing the destination since the last polling update.*
Msgs Current	The current number of messages in the destination.*
Msgs Deleted	The number of messages that have been deleted from the destination.*
Msgs High	The highest number of messages in the destination since the last polling update.*
Msgs Moved	The number of moved messages in the destination since the last polling update.*
Msgs Pending	The current number of pending messages in the destination.*
Msgs Rcvd	The number of messages received in the destination since the last polling update.*
Rcvd Delta	The increase in the amount of messages received (from the previous polling period to the current polling period).
Rcvd Rate	The rate of messages received (per second) in the destination since the last polling update.*

Msgs Thresh Time	The amount of time in the threshold condition since the last polling update.
Bytes Thresh Time	The amount of time in the threshold condition since the last polling update.
Paused	Indicates whether or not the destination is paused at the current time.*
Consumption Paused	Indicates the consumption paused state of the destination.
Consumption Paused State	The current consumption paused state of the destination.
Production Paused	Indicates the production paused state of the destination.
Production Paused State	The current production paused state of the destination.
Insertion Paused	Indicates the insertion paused state of the destination.
Insertion Paused state	The current insertion pause state of the destination.
State	The current health state of the destination.
JMSServer Runtime	The name of the JMS Server.
time_stamp	The date and time this row of data was last updated.

JMS Bridges View

This view contains tabular and summary displays that show performance metrics for JMS bridges. The following displays are available:

- [“All Bridges Table”](#): A tabular view that allows you to view performance and utilization metrics for all JMS Bridge Destinations.
- [“Bridge Summary”](#): A summary view that allows you to view performance and utilization metrics for a bridge associated with a particular JMS Server.

All Bridges Table

This table allows you to view performance and utilization metrics for all JMS Bridge Destinations.

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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

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

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

Open the Alert Views - RTView Alerts Table display.

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

Fields and Data

This display includes:

JMS Bridge Destinations Table

This table describes all JMS Bridge Destinations.

Name	The user-specified name of the instance.
Source Destination Name	The name of the source destination from which the messaging bridge instance reads messages.
Target Destination Name	The name of the target destination where the messaging bridge instance sends the messages it receives from the source destination.
Async Enabled	When checked, denotes that the messaging bridge instance forwards in asynchronous messaging mode.*
Batch Interval	The maximum amount of time, in milliseconds, that a messaging bridge instance waits before sending a batch of messages in one transaction.*
Batch Size	The number of messages that are processed within one transaction.*
Deployment Order	The priority that the server uses to determine when it deploys an item.*
Durability Enabled	Denotes whether or not the messaging bridge allows durable messages.*
Idle Time Maximum	The maximum amount of time, in seconds, that a message bridge instance remains idle.*
Preserve Message Property	Specifies if message properties are preserved when messages are forwarded by a bridge instance.*
QOS Degradation Allowed	Specifies if this messaging bridge instance allows the degradation of its quality of service (QOS) when the configured QOS is not available.*
Quality Of Service	The quality of service for this messaging bridge instance.*
Reconnect Delay Increase	The incremental delay time, in seconds, that a messaging bridge instance increases its waiting time between one failed reconnection attempt and the next retry.*
Reconnect Delay Maximum	The longest time, in seconds, that a messaging bridge instance waits between one failed attempt to connect to the source or target and the next retry.*
Reconnect Delay Minimum	The minimum amount of time, in seconds, that a messaging bridge instance waits before it tries to reconnect to the source or target destination after a failure.*
Selector	The filter for messages that are sent across the messaging bridge instance.*
Started	Specifies the initial operating state of a targeted messaging bridge instance.*
time_stamp	The date and time this row of data was last updated.
Transaction Timeout	The amount of time, in seconds, that the transaction manager waits for each transaction before timing it out.*
Source Destination	The source destination from which the messaging bridge instance reads messages.*
Target Destination	The target destination where the messaging bridge instance sends the messages it receives from the source destination.*

JMS Messaging Bridge Runtime Table

This table describes the messaging bridge metrics.

Name	The name of the messaging bridge.*
Location	The location of the messaging bridge.*
State	The current state of the messaging bridge. If Inactive , the color of this row is set to red.*
Description	If the message bridge is not running, the reason why it is not running is listed here.*
time_stamp	The date and time this row of data was last updated.

Bridge Summary

This summary view allows you to view performance and utilization metrics for a bridge associated with a particular JMS Server.

←

JMS Bridge - Summary by Bridge

18-Oct-2016 10:38 Data OK + ?

Domain: TestDomain

WLS Server: ManagedServer2
JMS Server: WseeJmsServer_auto_3

Bridge Name:

Source Destination:

Target Destination:

☐ AsyncEnabled
Batch Interval:
Reconnect Delay Increase:

☐ DurabilityEnabled
Batch Size:
Reconnect Delay Maximum:


☐ PreserveMsgProperty
Idle Time Maximum:
Reconnect Delay Minimum:

☐ QOSDegradationAllowed
Transaction Timeout:
Quality Of Service:

☐ Started
Time Stamp:
Deployment Order:

Selector:

JMS Messaging Bridge Runtime



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

Fields and Data

This display includes:

Domain	Select the domain containing the WebLogic Server for which you want to view data.
WLS Server	Select the WebLogic server containing the JMS Server for which you want to view data.
JMS Server	Select the JMS Server containing the destination for which you want to view data
Bridge Name	The name of the bridge associated with the JMS Server.
Source Destination	The source destination from which the messaging bridge instance reads messages.*
Target Destination	The target destination where the messaging bridge instance sends the messages it receives from the source destination.*
Async Enabled	When checked, denotes that the messaging bridge instance forwards in asynchronous messaging mode.*
Durability Enabled	Denotes whether or not the messaging bridge allows durable messages.*
Preserve Msg Property	Specifies if message properties are preserved when messages are forwarded by a bridge instance.*
QOS Degradation Allowed	Specifies if this messaging bridge instance allows the degradation of its quality of service (QOS) when the configured QOS is not available.*
Started	Specifies the initial operating state of a targeted messaging bridge instance.*
Batch Interval	The maximum amount of time, in milliseconds, that a messaging bridge instance waits before sending a batch of messages in one transaction.*
Batch Size	The number of messages that are processed within one transaction.*
Idle Time Maximum	The maximum amount of time, in seconds, that a message bridge instance remains idle.*
Transaction Timeout	The amount of time, in seconds, that the transaction manager waits for each transaction before timing it out.*

time_stamp	The date and time the data in this display was last updated.
Reconnect Delay Increase	The incremental delay time, in seconds, that a messaging bridge instance increases its waiting time between one failed reconnection attempt and the next retry.*
Reconnect Delay Maximum	The longest time, in seconds, that a messaging bridge instance waits between one failed attempt to connect to the source or target and the next retry.*
Reconnect Delay Minimum	The minimum amount of time, in seconds, that a messaging bridge instance waits before it tries to reconnect to the source or target destination after a failure.*
Quality of Service	The quality of service for this messaging bridge instance.*
Deployment Order	The priority that the server uses to determine when it deploys an item.*
Selector	The filter for messages that are sent across the messaging bridge instance.*

JMS Messaging Bridge Runtime Table

This table describes the messaging bridge metrics.

Location	The location of the messaging bridge.*
State	The current state of the messaging bridge. If Inactive , the color of this row is set to red.*
Description	If the message bridge is not running, the reason why it is not running is listed here.*

JMS Connections View

This view includes a display that presents a view of JMS performance metrics for a particular server.

- [“JMS Connections - Detail by Server”](#): A series of tables that allow you to view performance and utilization metrics for all JMS connections, JMS consumers, JMS producers, and JMS durable subscribers on a particular JMS server.

WLS Server Select the WebLogic server containing the JMS Server for which you want to view data.

JMS Server Select the JMS Server containing the destination for which you want to view data

JMS Connections Table

This table describes all JMS connections on the selected JMS server.

Name	The name of the connection.
ClientID	The client id for the connection.*
HostAddress	The host address of the client JVM as a string.*
Sessions Current Count	The current number of sessions for the connection.*
Sessions High Count	The highest number of sessions for the connection since the last polling update.
Sessions Total Count	The total number of sessions for the connection since the last polling update.
time_stamp	The date and time this row of data was last updated.

JMS Consumers Table

This table describes all JMS consumers on the selected JMS server.

Name	The name of the consumer.
Active	When checked, denotes that the consumer is active.*
Durable	When checked, denotes that the consumer is durable.*
Bytes Pending Count	The number of bytes pending by the consumer.*
Bytes Received Count	The number of bytes received by the consumer since the last polling update.*
Delta	The increase in the amount of bytes received (from the previous polling period to the current polling period).
Rate	The rate of bytes received (per second) by the consumer.
Messages Pending Count	The number of pending messages by the consumer.*
Messages Received Count	The number of messages received by the consumer since the last polling update.*
Delta	The increase in the amount of messages received (from the previous polling period to the current polling period).
Rate	The number of messages received per second by the consumer.
JMS Connection Runtime	The name of the JMS connection.
JMS Session Runtime	The name of the JMS session.
Destination Name	The name of the destination for the consumer.

Selector	The name of the selector associated with the consumer.
time_stamp	The date and time this row of data was last updated.

JMS Producers Table

This table describes all JMS producers on the selected JMS server.

Name	The name of the producer.
Bytes Pending Count	The current number of bytes that are pending by the producer.*
Bytes Sent Count	The number of bytes sent by the producer since the last polling update.*
Delta Bytes Sent Count	The increase in the amount of bytes sent (from the previous polling period to the current polling period).
Rate Bytes Sent Count	The number of bytes sent per second by the producer.
Messages Pending Count	The number of messages pending by the producer.*
Messages Sent Count	The number of messages sent by the producer since the last polling update.*
Delta Messages Sent Count	The increase in the amount of messages sent (from the previous polling period to the current polling period).
Rate Messages Sent Count	The number of messages sent by the producer per second.
JMS Connection Runtime	The name of the JMS connection.
JMS Session Runtime	The name of the JMS session.
time_stamp	The date and time this row of data was last updated.

JMS Durable Subscribers Table

This table describes all JMS durable subscribers on the selected JMS server.

ClientID	The client ID for the durable subscriber.*
Active	When checked, denotes that the subscription is being used by a durable subscriber.*
No Local	Indicates whether the durable subscriber receives local messages that it has published.*
Bytes Current Count	The number of bytes received by the durable subscriber.*
Bytes Pending Count	The current number of bytes that are pending by the durable subscriber.*
Messages Current Count	The number of messages still available by this durable subscriber.*
Messages Deleted Current Count	The number of messages deleted from the destination.*
Messages High Count	The highest number of messages for the durable subscriber since the last polling update.

Messages Moved Current Count	The number of messages that have been moved from the destination.*
Messages Pending Count	The number of messages pending by this durable subscriber.*
Messages Received Count	The number of messages received by the durable subscriber since the last polling update.
Current Consumer Info Client ID	The client ID for the consumer.*
Current Consumer Info Connection Address	The connection address of the consumer.*
Current Consumer Info Durable	When checked, denotes that the current consumer is durable.*
Current Consumer Info Name	The name of the current consumer.*
Current Consumer Info No Local	Indicates whether the consumer receives local messages it has published itself.*
Current Consumer Info Selector	The message selector defined for this consumer.*
Name	The name of the durable subscriber.*
Subscription Name	The subscription name for the durable subscriber.*
Selector	The message selector defined for the durable subscriber.*
time_stamp	The date and time this row of data was last updated.

CHAPTER 7 RTView DataServer for Solace

The RTView DataServer for Solace provides a way to create connections and modify default configuration settings for the Solace solution package and sends collected data to RTView Central, which contains the displays associated with the RTView DataServer for Solace that help you to monitor the health and performance across your Solace components.

RTView Central contains the following Views and their associated displays that will be populated with data collected via the RTView DataServer for Solace:

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

The RTView *DataCollector* for Solace is also available for use with the RTView DataServer for Solace. RTView DataCollector for Solace is used for collecting solution package data and sending it to one or more RTView DataServers. The RTView DataCollector for Solace is useful if you need to distribute data collection.

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

Message Routers

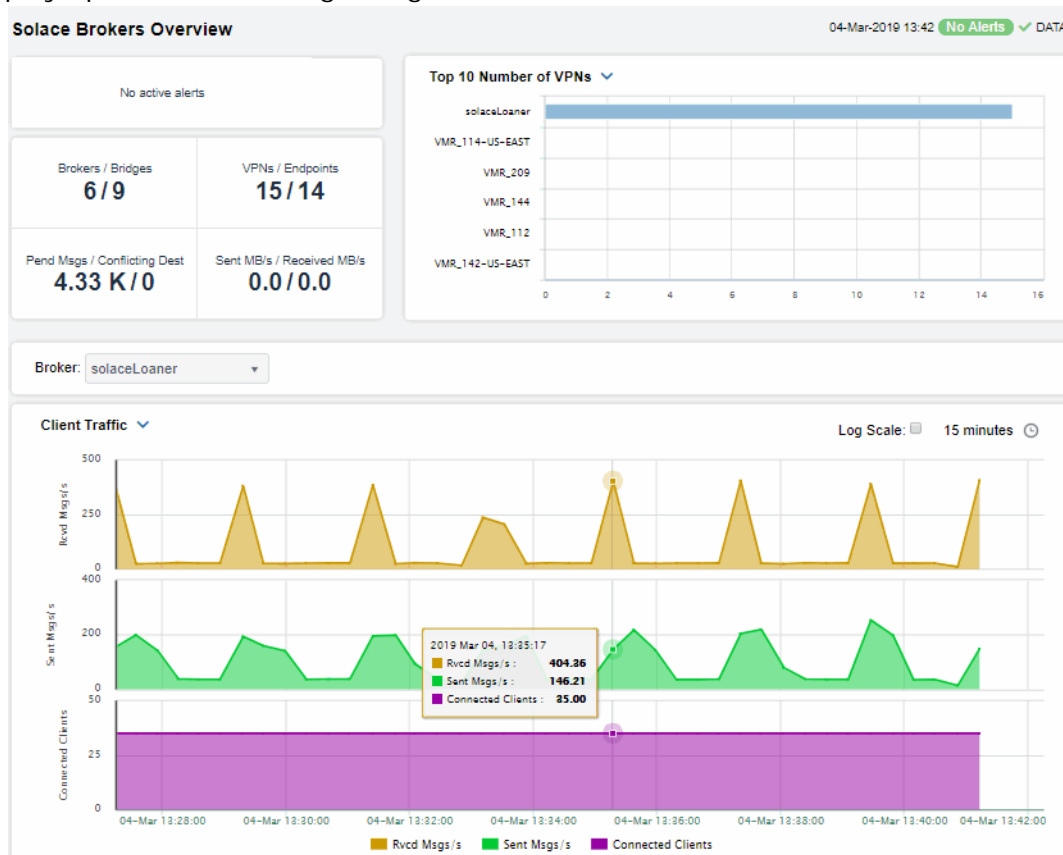
These displays provide detailed metrics for message routers and their connected message routers. Displays in this View are:

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

Message Routers Overview

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

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



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.
- 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.

CRITICAL

Total number of current critical alerts for message routers on the selected data server.

WARNING

Total number of current critical alerts for message routers on the selected data server.

Message Routers/Bridges

Total number of message routers/bridges on the selected data server.

VPNs/Endpoints

Total number of VPNs/endpoints on the selected data server.

Pending Msgs/Conflicting Dest

Total number of pending messages/conflicting destinations on the selected data server.

Sent MBs/Received MBs

Total number of MBs sent/MBs received on the selected data server.

Top 10 Number of VPNs

Ten message routers with the greatest number of connected VPNs.

Msg Router


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



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


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



Time Settings



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

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



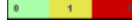
Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

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

Log Scale




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

Message Routers Heatmap

Use the **Show** dropdown menu to view the current status of **All** message routers, or just **Expired** or **Unexpired** message routers, in a heatmap. Each rectangle in the heatmap is a single message router where the rectangle size represents the number of connections. The rectangle color maps where the current value is on its color gradient  bar. Select a message router from the drop-down menu. For example, by default, **Alert Severity** is shown:

Alert Severity

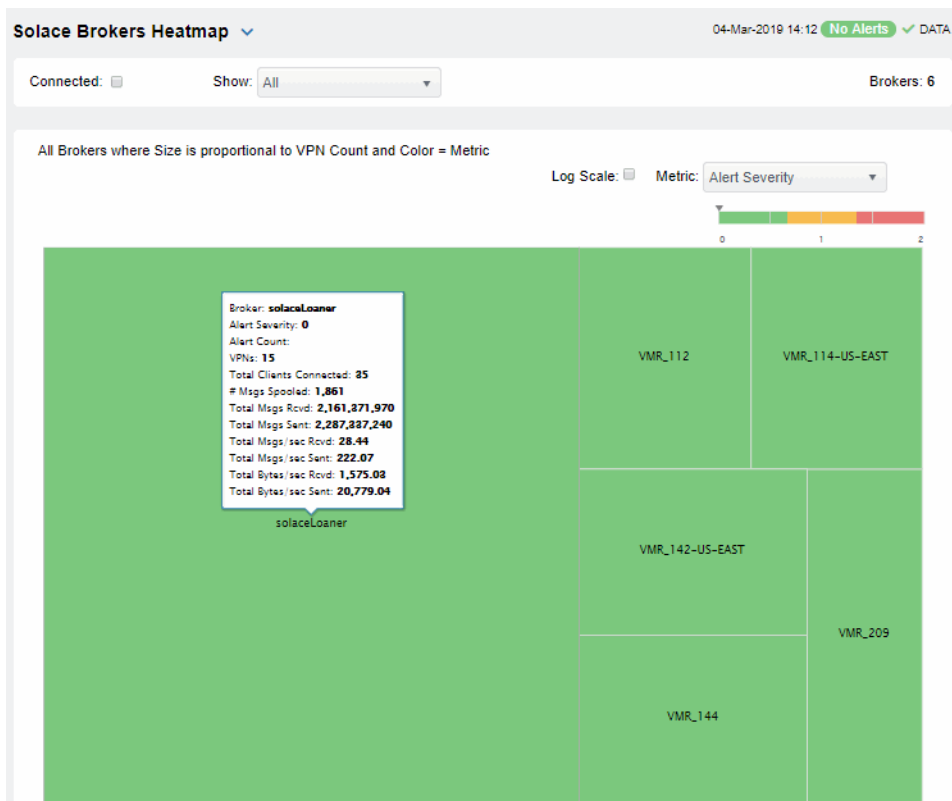
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

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

Each metric has its own color gradient bar (scroll down for more [“Metric Options”](#)).

Mouse over a rectangle to see additional details. Use the check-box ☒ to include / exclude **Connected** message routers and enable **Log Scale** mode. Click a rectangle to drill-down to details about a message router in the [“Message Router Summary”](#) display.

Consider keeping this display open for monitoring your Solace message routers at a glance.



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.

Metric Options

Choose a metric from the drop-down menu:







Alert Severity

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

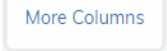
Alert Count

The total number of critical and warning alerts. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

# Msgs Spooled	The total number of spooled messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterPendingMsgsHigh . The middle value in the gradient bar indicates the middle value of the range.
Total Msgs Rcvd	The total number of received messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of total messages received in the heatmap. The middle value in the gradient bar indicates the average count.
Total Msgs Sent	The total number of sent messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of total messages sent in the heatmap. The middle value in the gradient bar indicates the average count.
Total Msgs/ sec Rcvd	The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterInboundMsgRateHigh . The middle value in the gradient bar indicates the middle value of the range.
Total Msgs/ sec Sent	The total number of messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterOutboundMsgRateHigh . The middle value in the gradient bar indicates the middle value of the range.
Total Bytes/ sec Rcvd	The total number of bytes received per second in the message router. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterInboundByteRateHigh . The middle value in the gradient bar indicates the middle value of the range.
Total Bytes/ sec Sent	The total number of bytes sent per second in the message router. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterOutboundByteRateHigh . The middle value in the gradient bar indicates the middle value of the range.

Message Routers Table

Use the **Show:** dropdown menu to view the current status of **All** message routers, or just **Expired** or **Unexpired** message routers, in a tabular format. Each row in the table is a different message router.

By default, a subset of available metrics is shown. Click  (**More Columns**) to toggle to the complete set of metrics available.

Msg Routers: 23 (in the upper right portion) is the number of message routers in the display. Select a message router from the drop-down menu. Sort data in numerical or alphabetical order on column headers. Use the check-boxes ☒ to include / exclude **Connected** and **Expired** message routers.

Data shown in the ["Message Routers Heatmap"](#) is shown as well as additional metrics.

Double-click a row to drill-down and investigate in the “[Message Router Summary](#)” display. See “[Column Values](#)” for details about metrics shown.

Solace Message Routers Table 14-Aug-2018 14:49 2 Alerts DATA ?

Show Connected Only: ☐ Show: All Msg Routers: 27 [More Columns](#)

Message Router	Connected	Alert Level	Host Name	VPNs	Total Clients Connected	Endpoints	Bridge
Sol_Mule_Azure			SolaceVMR0	2	3	3	
Sol_Mule_GCP			solace-vmr-instance-1	2	2	0	
Sol_Mule_SC			vmr-mr8v6yiwahm	1	0	4	
solDemo			solace	14	36	14	
Team-1-138-62			13.56.60.2	0	0	0	
Team-1-159-41			ip-10-0-151-49	6	11	27	
Team-1-demo-vmr-1			104.196.188.129	0	0	0	
Team-1-sgdemo1			sgdemo.solace.com	0	0	0	
Team-1-VMR88			52.187.108.151	0	0	0	
Team-2-Ali-cloud			47.74.235.254	0	0	0	
Team-2-AWS			ec2-35-177-122-45.e	2	0	6	
Team-2-Azure			demothon-team20.sol	0	0	0	
Team-2-Google-cloud			35.187.64.112	0	0	0	
Team-2-lab-appliance			london.solace.com	1	0	5	






Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

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.

Column Values

Message Router	The name of the message router.
Connected	<p>The message router state:</p> <ul style="list-style-type: none">  Red indicates that the message router is NOT connected.  Green indicates that the message router is connected.
Alert Severity	<p>The current alert severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts.
Expired	<p>When checked, performance data about the sensor has not been received within the time specified (in seconds) in the \$solRowExpirationTime field in the conf\rtvapm_solmon.properties file. The \$solRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the sensor. To view/edit the current values, modify the following lines in the .properties file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the Expired check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
Host Name	The name of the host.
Platform	The name of the platform.
OS Version	The version of the operating system.
Up Time	The amount of time that the message router has been up and running.
VPNs	The total number of VPNs configured on the message router.
Total Clients	The total number of clients associated with the message router.
Total Clients Connected	The total number of clients that are currently connected to the message router.
Clients Using Compression	The number of clients who send/receive compressed messages.
Clients Using SSL	The number of clients using SSL for encrypted communications.

Max Client Connections	The maximum number of available client connections.
Endpoints	The total number of Endpoints configured on the message router.
Bridges	The total number of bridges configured on the message router.
Local Bridges	The total number of local bridges configured on the message router.
Remote Bridges	The total number of remote bridges configured on the message router.
Remote Bridge Subscriptions	The total number of remote bridge subscriptions configured on the message router.
Routing Enabled	This check box is checked when the message router is configured to route messages to other message routers.
Routing Interface	The name of the interface configured to support message routing.
Total # Conflicting Destinations	The total number conflicting destinations.
Pending Messages	The number of pending messages on the message router.
Total Client Msgs Rcvd	The total number of client messages received on the message router.
Total Client Msgs Sent	The total number of client messages sent by the message router.
Total Client Msgs Rcvd/sec	The total number of client messages received per second by the message router.
Total Client Msgs Sent/ sec	The total number of client messages sent by the message router.
Total Client Bytes Rcvd	The total number of client bytes received by the message router.
Total Client Bytes Sent	The total number of client bytes sent by the message router.
Total Client Bytes Rcvd/sec	The total number of client bytes received per second by the message router.
Total Client Bytes Sent/sec	The total number of client bytes sent per second by the message router.
Total Client Direct Msgs Rcvd	The total number of direct client messages received by the message router.
Total Client Direct Msgs Sent	The total number of direct client messages sent from the message router.
Total Client Direct Msgs Rcvd/sec	The total number of direct client messages received per second by the message router.
Total Client Direct Msgs Sent/sec	The total number of direct client messages sent per second by the message router.
Total Client Direct Bytes Rcvd	The total number of direct client bytes received by the message router.
Total Client Direct Bytes Sent	The total number of direct client bytes sent by the message router.
Total Client Direct Bytes Rcvd/sec	The total number of direct client bytes received per second by the message router.

Total Client Direct Bytes Sent/sec	The total number of direct client bytes sent per second by the message router.
Total Client Non-Persistent Msgs Rcvd	The total number of non-persistent client messages received by the message router.
Total Client Non-Persistent Msgs Sent	The total number of non-persistent client messages sent by the message router.
Total Client Non-Persistent Msgs Rcvd/sec	The total number of non-persistent client messages received per second by the message router.
Total Client Non-Persistent Msgs Sent/ sec	The total number of non-persistent client messages sent per second by the message router.
Total Client Non-Persistent Bytes Rcvd	The total number of non-persistent client bytes received by the message router.
Total Client Non-Persistent Bytes Sent	The total number of non-persistent client bytes sent by the message router.
Total Client Non-Persistent Bytes Rcvd/sec	The total number of non-persistent client bytes received per second by the message router.
Total Client Non-Persistent Bytes Sent/sec	The total number of non-persistent client bytes sent per second by the message router.
Total Client Persistent Msgs Rcvd	The total number of persistent client messages received by the message router.
Total Client Persistent Msgs Sent	The total number of persistent client messages sent by the message router.
Total Client Persistent Msgs Rcvd/sec	The total number of persistent client messages received per second by the message router.
Total Client Persistent Msgs Sent/ sec	The total number of persistent client messages sent per second by the message router.
Total Client Persistent Bytes Rcvd	The total number of persistent client bytes received by the message router.
Total Client Persistent Bytes Sent	The total number of persistent client bytes sent by the message router.
Total Client Persistent Bytes Rcvd/sec	The total number of persistent client bytes received per second by the message router.
Total Client Persistent Bytes Sent/ sec	The total number of persistent client bytes sent per second by the message router.
Avg Egress Bytes/min	The average number of outgoing bytes per minute.
Avg Egress Compressed Msgs/min	The average number of outgoing compressed messages per minute.
Avg Egress Msgs/min	The average number of outgoing messages per minute.
Avg Egress SSL Msgs/min	The average number of outgoing messages per minute being sent via SSL-encrypted connections.
Avg Egress Uncompressed Msgs/min	The average number of uncompressed outgoing messages per minute.
Avg Ingress Bytes/min	The average number of incoming bytes per minute.

Avg Ingress Compressed Msgs/min	The average number of compressed incoming message per minute.
Avg Ingress Msgs/min	The average number of incoming messages per minute.
Average Ingress SSL Msgs/min	The average number of incoming messages per minute being received via SSL-encrypted connections.
Avg Ingress Uncompressed Msgs/min	The average number of uncompressed messages per minute.
Current Egress Bytes/sec	The current number of outgoing bytes per second.
Current Egress Compressed Msgs/sec	The current number of outgoing compressed messages per second.
Current Egress Msgs/sec	The current number of outgoing messages per second.
Current Egress SSL Msgs/sec	The current number of outgoing messages per second sent via SSL-encrypted connections.
Current Egress Uncompressed Msgs/sec	The current number of outgoing uncompressed messages per second.
Current Ingress Bytes/sec	The current number of incoming bytes per second.
Current Ingress Compressed Msgs/sec	The current number of incoming compressed messages per second.
Current Ingress Msgs/sec	The current number of incoming messages per second.
Current Ingress SSL Msgs/sec	The current number of incoming messages per second received via SSL-encrypted connections.
Current Ingress Uncompressed Msgs/sec	The current number of incoming uncompressed messages per second.
Ingress Comp Ratio	The percentage of incoming messages that are compressed.
Egress Comp Ratio	The percentage of outgoing messages that are compressed.
Egress Compressed Bytes	The number of outgoing compressed bytes.
Egress SSL Bytes	The number of outgoing compressed bytes being sent via SSL-encrypted connections.
Egress Uncompressed Bytes	The number of outgoing uncompressed bytes.
Ingress Compressed Bytes	The number of incoming compressed bytes.
Ingress SSL Bytes	The number of incoming bytes via SSL-encrypted connections.
Ingress Uncompressed Bytes	The number of incoming uncompressed bytes.
Total Egress Discards	The total number of outgoing messages that have been discarded by the message router.
Total Egress Discards/sec	The total number of outgoing messages per second that have been discarded by the message router.
Total Ingress Discards	The total number of incoming messages that have been discarded by the message router.

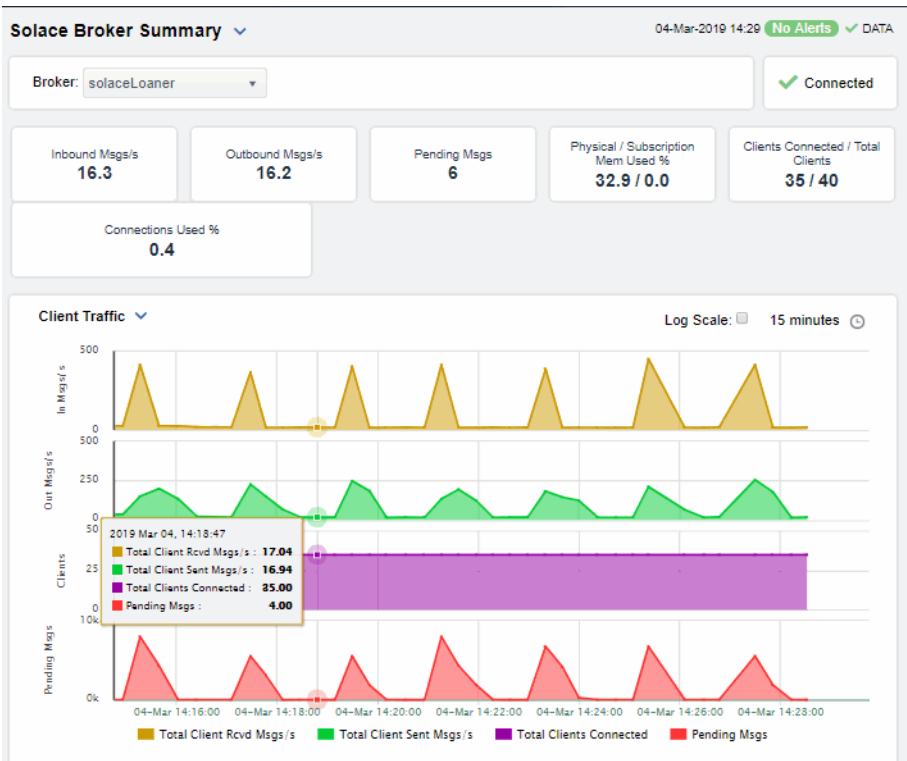
Total Ingress Discards/sec	The total number of incoming messages per second that have been discarded by the message router.
Client Authorization Failures	The number of failed authorization attempts
Client Connect Failures (ACL)	The number of client connection failures caused because the client was not included in the defined access list.
Subscribe Topic Failures	The number of failed attempts at subscribing to topics.
TCP Fast Retrans Sent	The total number of messages that were retransmitted as a result of TCP Fast Retransmission (one or more messages in a sequence of messages that were not received by their intended party that were sent again).
Memory (KB)	The total available memory (in kilobytes) on the message router.
Memory Free (KB)	The total amount of available memory (in kilobytes) on the message router.
Memory Used (KB)	The total amount of memory used (in kilobytes) on the message router.
Memory Used %	The percentage of total available memory that is currently being used.
Swap (KB)	The total available swap (in kilobytes) on the message router.
Swap Free (KB)	The total amount of available swap (in kilobytes) on the message router.
Swap Used (KB)	The total amount of swap used (in kilobytes) on the message router.
Swap Used %	The percentage of total available swap that is currently being used.
Subscription Mem Total (KB)	The total amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Free (KB)	The current amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Used (KB)	The current amount of memory (in kilobytes) being used by queue/topic subscriptions.
Subscription Mem Used %	The percentage of available memory being used by queue/topic subscriptions.
Chassis Product Number	The product number of the chassis in which the message router is contained.
Chassis Revision	The revision number of the chassis.
Chassis Serial	The serial number of the chassis.
BIOS Version	The basic input/output system used by the chassis.
CPU-1	The name of the central processing unit (CPU 1) used by the message router.
CPU-2	The name of the central processing unit (CPU 2) used by the message router.
Operational Power Supplies	The number of available power supplies that are operational on the chassis.

Power Redundancy Config	The configuration used by the backup message router.
Max # Bridges	The maximum number of bridges allowed on the message router.
Max # Local Bridges	The maximum number of local bridges allowed on the message router.
Max # Remote Bridges	The maximum number of remote bridges allowed on the message router.
Max # Remote Bridge Subscriptions	The maximum number of remote bridge subscriptions allowed on the message router.
Redundancy Config Status	The status of the redundancy configuration.
Redundancy Status	The status of the redundant message router.
Redundancy Mode	Refer to Solace documentation for more information.
Auto-revert	Refer to Solace documentation for more information.
Mate Router Name	If redundancy is configured, this field lists the redundant message router name (mate message router name).
ADB Link Up	This check box is checked if a message router is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.
ADB Hello Up	Refer to Solace documentation for more information.
Pair Primary Status	The primary status of the message router and its redundant (failover) mate.
Pair Backup Status	Refer to Solace documentation for more information.
Expired	<p>When checked, performance data about the message router has not been received within the time specified (in seconds) in the \$solRowExpirationTime field in the conf\rtvapm_solmon.properties file. The \$solRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the message router. To view/edit the current values, modify the following lines in the .properties file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the Expired check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
Time Stamp	The date and time the row of data was last updated.

Message Router Summary

View trend graphs of current and historical performance metrics of client traffic on a message router. Check general health details for any message router.

Choose a message router from the drop-down menu to view its total number of connected clients, number of incoming messages, **Up Time**, and additional information. You can also view alert statuses and **Spool Status** data for the message router.



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.

The connection status (connected/disconnected).



Inbound Msgs/s

The number of messages received per second.

Outbound Msgs/s

The number of messages sent per second.

Pending Msgs/s

The number of pending messages.

Physical / Subscription Mem Used %	The total percentage of physical memory used / the total percentage of subscription memory used.
Clients Connected / Total Clients	The current number of clients connected / the total number of clients.
Connections Used %	The percentage of connections used.

Trend Graphs

Traces the sum for the selected message router.

Client Traffic

- **In Msgs/s** - Traces the total number of client messages received per second.
- **Out Msgs/s** - Traces the total number of client messages sent per second.
- **Clients** - Traces the total number of connected clients.
- **Pending Msgs** - Traces the total number of pending messages.

Spool Msgs

- **Pending Msgs** - Traces the total number of pending spool messages.
- **Spool Usage MB** - Traces the total amount of space used by spool messages, in megabytes.

Memory


- **Memory Used %** - Traces the percent of memory used.



Subscription Mem Used % - Traces the percent of memory used by subscriptions.


Log Scale



Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

Message Router Sensors


This tabular display contains environmental sensor metrics for a selected message router. Use this display to find out the type, name, value, and status of the sensors.


Select a message router from the drop-down menu. Note that display does not show data for VMRs as it only applies to message routers.

Solace Broker Environmental Sensors							04-Mar-2019 14:38	Alerts	DATA
Broker: solaceLoaner									
Sensor Readings									
Type	Sensor Name	Value	Units	Status	Expired	Time Stamp			
Voltage	BB +1.5V	1.469	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +1.5V AUX	1.490	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +1.5V ESB	1.482	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +1.8V	1.803	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +12V AUX	12.090	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +3.3V	3.337	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +3.3V STB	3.337	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +5V	5.070	volts	OK		04-Mar-2019 14:37:52			
ThermalMargin	CPU1 Therm Margin	-87.000	degrees C			04-Mar-2019 14:37:52			
ThermalMargin	CPU2 Therm Margin	-59.000	degrees C			04-Mar-2019 14:37:52			
Temperature	Chassis Temp.	23.000	degrees C			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 1	7714	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 2	8057	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 3	7714	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 4	7543	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 5	7371	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 6	7288	RPM			04-Mar-2019 14:37:52			
Power system status	Power Redundancy	yes				04-Mar-2019 14:37:52			

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.

Sensor Readings

Each row in the table is a different sensor on the message router.

Type	See vendor documentation for details.
Sensor Name	The name of the sensor.
Value	Lists the value of the sensor.
Units	Lists the unit of measure for the sensor.
Status	The current status of the sensor.

Expired

When checked, performance data about the sensor has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapi_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the sensor. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number of seconds
#
collector.sl.rtvapi.sub=$solRowExpirationTime:45
collector.sl.rtvapi.sub=$solRowExpirationTimeForDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Time Stamp

The date and time the row of data was last updated.

Message Router Provisioning

This display shows provisioning metrics for a single message router. Use this to see the host, platform, chassis, memory, redundancy and fabric data for a specific message router.

Select a message router from the drop-down menus.

Broker Provisioning
04-Mar-2019 14:39
No Alerts
DATA

Broker: solaceLoaner

Host Name: **solace**
 CPU-1: **Intel(R) Xeon(R) CPU E5450 @ 3.00GHz**

CPU-2: **Intel(R) Xeon(R) CPU E5450 @ 3.00GHz**
 Platform: **Solace 3260**

Chassis Product Number: **CHS-3260AC-01-B**
 BIOS Version: **\$5000.86B.10.00.0094.101320081858**

Chassis Revision: **1.4**
 Chassis Serial: **S009000226**
 Power Redundancy Config: **2+2**

Total Memory (KB): **15,965,652**
 Memory Used (KB): **9,332,304**
 Memory Used %: **32.89**

Swap (KB): **2,007,992**
 Swap Used (KB): **0**
 Swap Used %: **0.0**

Operational Power Supplies: **4**
 Mate Router Name:
 Redundancy Config Status: **Shutdown**

Redundancy Status: **Down**
 Redundancy Mode: **N/A**
 Pair Primary Status: **Local Active**


Pair Backup Status: **Shutdown**
 Auto Revert: **false**
 ADB Link To Mate Up: **false**

ADB Hello To Mate Up: **false**

Last Update: **04-Mar-2019 14:39:28**

Fabric

Product	Fw-Version	Card Type	Slot	Serial #
NAB-0801ET-01-A	6.2.0.466	Network Acceleration Blade	1/1	S003000276
		in use by slot 1/1	1/2	
TRB-000000-02-A		Topic Routing Blade	1/3	P004045787
HBA-0204FC-02-A		Host Bus Adapter Blade	1/4	LFC0848B99469
ADB-000000-01-A		Assured Delivery Blade	1/5	S003000844
		empty	2/1	
		empty	2/2	
		empty	2/3	



Host Name	The name of the host.	
Platform	The platform on which the message router is running.	
Chassis Product #	The product number of the chassis in which the message router is contained.	
Chassis Revision #	The revision number of the chassis.	
Chassis Serial #	The serial number of the chassis.	
Power Configuration	The power configuration used by the chassis.	
Operational Power Supplies	The number of available power supplies that are operational on the chassis.	
CPU 1	The name of the central processing unit (CPU 1) used by the message router.	
CPU 2	The name of the central processing unit (CPU 2) used by the message router.	
BIOS	The basic input/output system used by the chassis.	
Memory (KB)	Physical	Lists the Total amount, the Free amount, the Used amount, and the Used % of physical memory.
	Swap	Lists the Total amount, the Free amount, the Used amount, and the Used % of swap memory.
Redundancy	These fields describe a fault tolerant pair of message routers.	
	Mate Router Name	If redundancy is configured, this field lists the redundant message router name (mate message router name).
	Configuration Status	The status of the configuration for the backup message router.
	Redundancy Status	The status of the redundant message router.
	Redundancy Mode	Refer to Solace documentation for more information.
	Primary Status	The status of the primary message router.
	Backup Status	Refer to Solace documentation for more information.
	Auto-Revert	Refer to Solace documentation for more information.

ADB Link Up

This check box is checked if a message router is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.

ADB Hello Up

Refer to Solace documentation for more information.

Fabric**Slot**

Displays the slot number on the network switch.

Card Type

The type of card connected to the particular slot.

Product

The product associated with the particular slot.

Serial #

The serial number of the product.

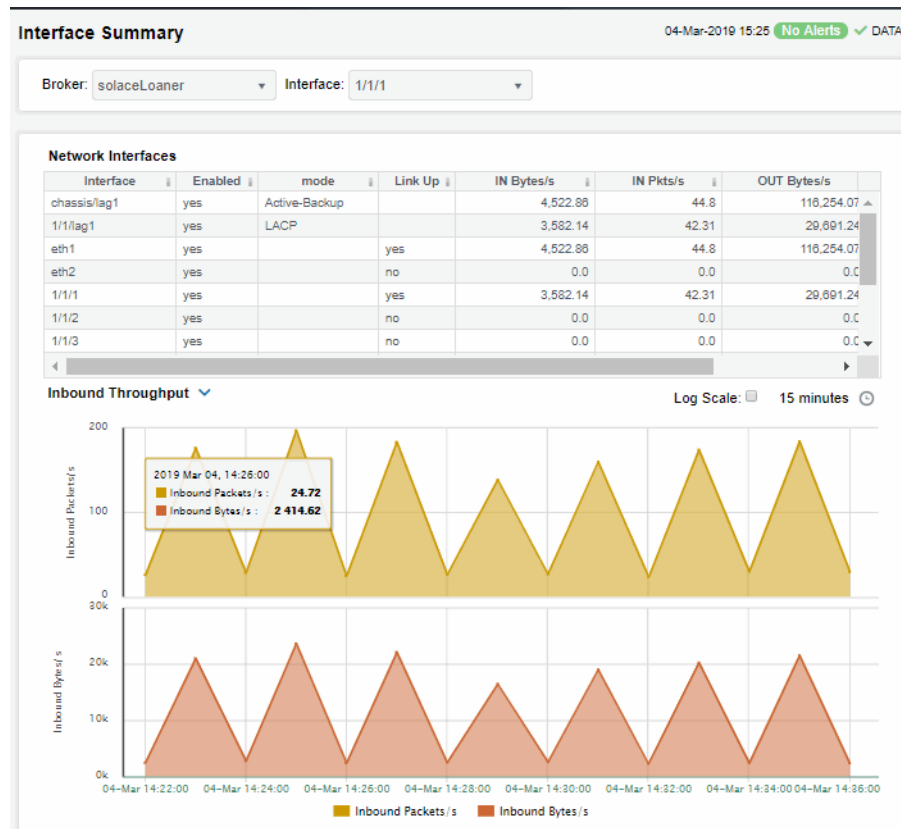
Fw-Version


The firmware version of the product.

Message Router Interface

This display lists all network interfaces on a selected message router, and shows network interface status, in/out throughput per second and additional detailed metrics.

Select a message router and interface from the drop-down menus. Each row in the table is a different network interface. Double-click one to trace its current and historical performance data in the trend graph (bytes in/out and packets in/out per second).






Interface	The name of the network interface.
Enabled	Displays whether or not the network interface is enabled.
mode	Describes how the interface is configured to support networking operations.
Link Up	Indicates whether the interface is electrically signaling on the transmission medium.
IN Bytes/sec	The number of bytes per second contained in incoming messages.
IN Pkts/sec	The number of incoming packets per second.
OUT Bytes/ sec	The number of bytes per second contained in the outgoing messages.
OUT Pkts/sec	The number of outgoing packets per second.

Trend Graphs

Inbound Pkts/ sec	Traces the number of incoming packets per second.
Outbound Bytes/sec	Traces the number of bytes per second contained in the incoming messages.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Time Settings	<p>By default, the time range end point is the current time. To change the time range, click the Time Settings and either:</p> <ul style="list-style-type: none"> choose a Time range from 5 Minutes to 7 Days in the drop-down menu. specify begin/end dates using the calendar . specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .

Restore settings to current time by selecting **now** .

Message Routers Message Spool

Select a message router from the drop-down menu or select **All**. This display shows operational status and spooling performance metrics (if spooling is enabled on the message router) for one or all message routers. Refer to Solace documentation for details about data in this display.





Solace Broker Spool Table 04-Mar-2019 15:35 No Alerts DATA

Broker: - All -

Count: 1

Pool MB	Msg Spool Used By Queue	Msg Spool Used By DTE	Message Count % Usage	Delivered UnAcked Msgs % Usage	Ingress Flow Count	Outgress Flow Count
0.0	13	1	0.0	0.0	18	

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.

Count	The number of message routers that are using spooling in the table.
Connection	The name of the message router.
Config Status	The status of the connection's configuration.
Operational Status	The operational status of the spool on the message router.
Current Spool Usage (MB)	The current amount of spool used in megabytes on the message router (calculated by summing spool used for each endpoint).
Msg Spool Used By Queue	The amount of spool used by the queue.
Msg Spool Used By DTE	The amount of spool used by DTE.
Message Count % Utilization	The percentage of total messages that use the message spool.
Delivered UnAcked Msgs % Utilization	The percentage of messages delivered via the spool that have not been acknowledged.

Ingress Flow Count	The current incoming flow count.
Ingress Flows Allowed	The total number of incoming flows allowed.
Queue/Topic Subscriptions Used	The number of queue/topic subscriptions used.
Max Queue/Topic Subscriptions	The maximum number of queue/topic subscriptions available.
Sequenced Topics Used	The number of sequenced topics used.
Max Sequenced Topics	The maximum number of sequenced topics available.
Spool Files Used	The number of spool files used.
Spool Files Available	The maximum number of spool files available.
Spool Files % Utilization	The percentage of available spool files that have been used.
Active Disk Partition % Usage	The percentage of available active disk partition that has been used.
Standby Disk Partition % Usage	The percentage of available standby disk partition that has been used.
Disk Usage Current (MB)	The current amount of spool disk usage in megabytes.
Disk Usage Max (MB)	The maximum amount of available spool disk usage in megabytes.
Transacted Sessions Used	The current number of transacted sessions.
Transacted Sessions Max	The maximum number of transacted sessions allowed.
Transacted Session Count % Utilization	The percentage of allowable transacted sessions that have been used.
Transacted Session Resource % Utilization	The percentage of allowable transacted session resources that have been used.
Expired	<p>When checked, performance data about the message router has not been received within the time specified (in seconds) in the \$solRowExpirationTime field in the conf\rtvapm_solmon.properties file. The \$solRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the message router. To view/edit the current values, modify the following lines in the .properties file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvapm.sub=\$solRowExpiration Time:45 collector.sl.rtvapm.sub=\$solRowExpiration TimeForDelete:3600</pre> <p>In the example above, the Expired check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>

CSPF Neighbors

These displays provide detailed data and statuses for CSPF neighbor message routers. Check trends on network traffic among CSPF neighbors. Displays in this View are:

- **"Neighbors Table"**: View metrics for Solace neighbor message routers that use the Content Shortest Path First (CSPF) routing protocol to determine the shortest path in which to send messages from one message router to another message router in the Solace network.
- **"Neighbors Diagram"**: Topological view of CSPF Neighbors that shows message router connections and status of servers (Active/Inactive).
- **"Neighbors Summary"**: View detailed performance metrics for a single Solace neighbor message router that uses the CSPF routing protocol.

Neighbors Table

This tabular display shows Content Shortest Path First (CSPF) "neighbor" metrics for a message router. Select a message router from the drop-down menu. View metrics for a Solace neighbor message router that uses the CSPF routing protocol to determine the least cost path in which to send messages from one message router to another message router in the Solace network.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Solace CSPF Neighbors Table 14-Aug-2018 16:16 No Alerts DATA ?

Msg Router: - All - More Columns

Show Ok Only: ☐ Show: All Neighbors: 4

Message Router	Name	Expired	State	Sent Msgs/s	Sent Bytes/s	Connections
VMR-112	ip-172-30-1-144		Ok	0.15	24.38	4
VMR-144	ip-172-30-1-112		Ok	0.2	0.0	4
VMR-144	ip-172-30-1-209		Ok	0.2	0.0	4
VMR-209	ip-172-30-1-144		Ok	0.17	29.41	4

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.

Neighbor Count:	The number of neighbor message routers connected to the selected Msg Router .
Show:	<p>OK Select to <i>only</i> show neighbor message routers that are connected (State is OK). By default, this option is not selected (all neighbor message routers are shown).</p> <p>Expired Select to show <i>both</i> expired and non-expired neighbor message routers. By default, this option is not selected (only non-expired neighbor message routers are shown).</p>

Table:



Each table row is a different neighbor message router.

Message Router	The name of the neighbor message router.
State	The current state of the message router.
Up Time	The amount of time the message router has been up and running.
Connections	The number of connections.
Link Cost Actual	Refer to Solace documentation for more information.
Link Cost Configured	Refer to Solace documentation for more information.
Data Port	Refer to Solace documentation for more information.
Expired	<p>When checked, performance data about the message router has not been received within the time specified (in seconds) in the \$solRowExpirationTime field in the conf\rtvapm_solmon.properties file. The \$solRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the message router. To view/edit the current values, modify the following lines in the .properties file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the Expired check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
Timestamp	The date and time the row of data was last updated.

Neighbors Diagram

Use this topology display to monitor the health of network components: Solace message routers, VMRs and servers. Quickly identify message router neighbors, servers that are inactive and which resources their performance impacts. Drag and drop objects to arrange them on the screen (doing so does not logically impact the Solace message routers, VMRs and servers).

Each object is a Solace message router, VMR or server. Each are labeled with their name and color coded as follows:

-  Red indicates that the object has one or more alerts in a critical state.
-  Yellow indicates that the object has one or more alerts in a warning state.

- Green indicates that there are no alerts on the object.
- Gray indicates that the object is off-line.

Mouse-over objects to see their host IP address.

Right-click on VMR objects and select **Open VMR UI** to open the Solace VMR login web page.

Save: Saves the arrangement of the objects.

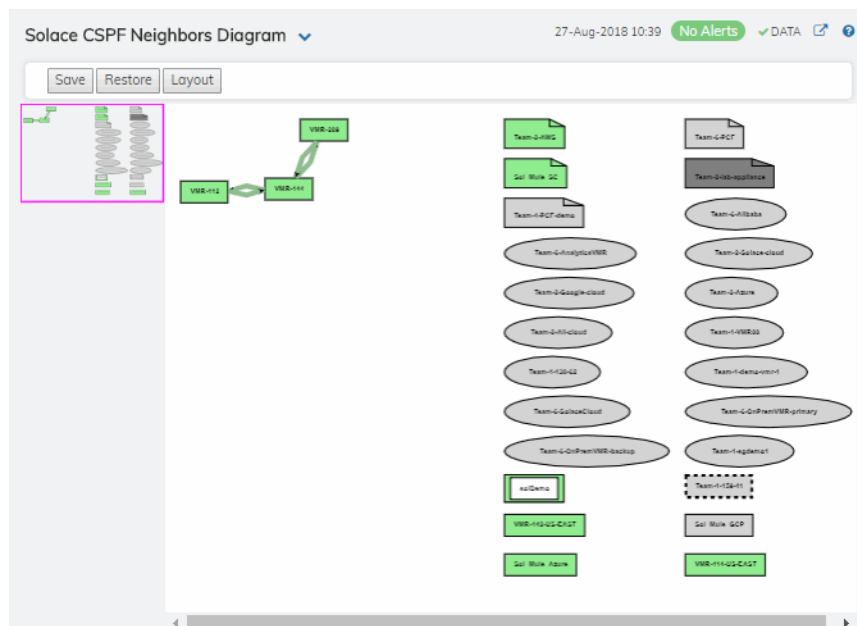
Restore: Returns objects to their previous positions.

Layout: Toggles between two types of layouts. One layout positions objects to the right so you might scroll in that direction to see them. The other layout pulls the objects close together to the left, vertically and in hierarchical order.





Look at the miniature view in (upper left) to see all objects in either layout. Zoom in on an area in the topology by clicking it in the miniature view.




Drill-down to investigate in the "Neighbors Table".


To monitor network bridges and VPNs, see the “[Bridges - Diagram](#)”.




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.

   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.

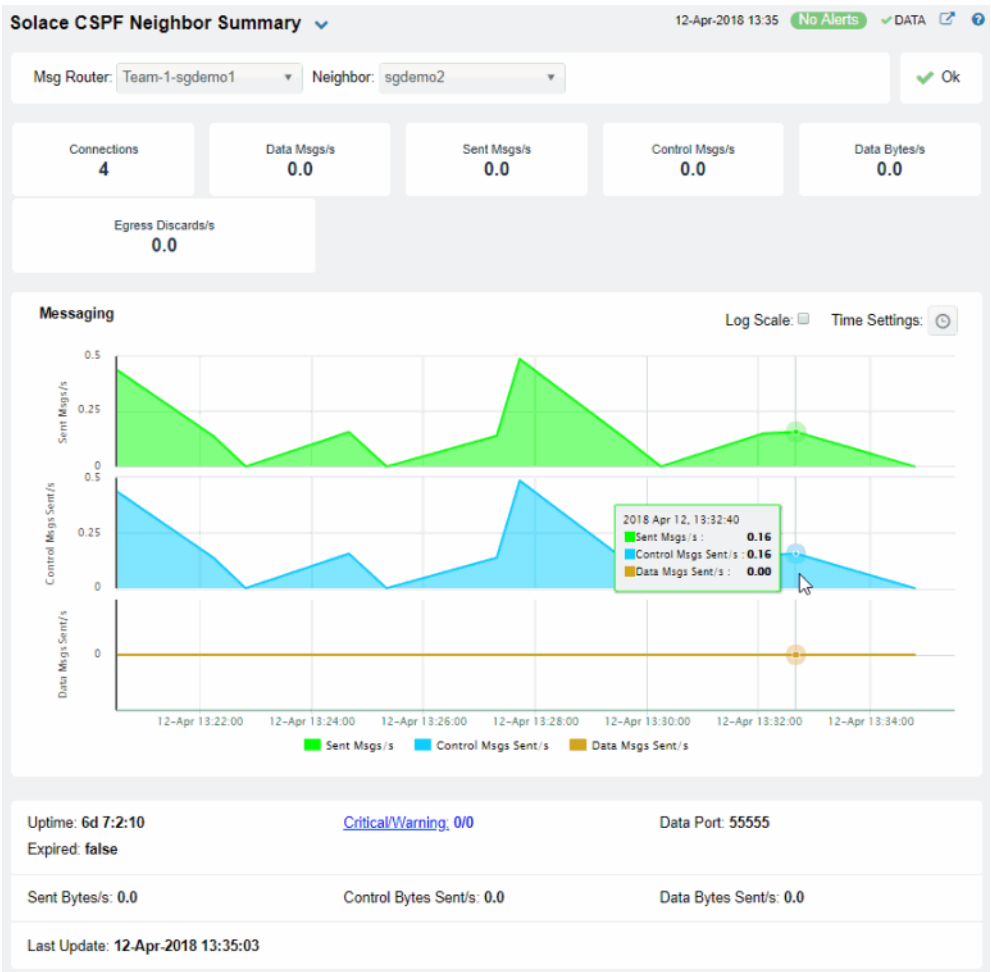
 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Neighbors Summary

View neighbor message router current configuration details and message throughput rates.

Select a message router and a neighbor message router from the drop down menus. Check message throughput rates to the neighbor message router, as well as neighbor **Up Time**, **State**, **Data Port**, number of connections and link costs.

The trend graph traces the current and historical message throughput (**Data**, **Control**, **Discards** and **Total**).



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.

- Neighbor:**




Select the neighbor message router for which you want to show data in the display.
- Connections**


The current number of connections.



Data Msgs/s	Refer to Solace documentation for more information.
Sent Msgs/s	Refer to Solace documentation for more information.
Control Msgs/s	Refer to Solace documentation for more information.
Data Bytes/s	Refer to Solace documentation for more information.
Egress Discards/s	The total number of discarded messages sent from the selected Msg Router to the selected Neighbor message router since the message router was last started.


Trend Graphs

Traces the rates of messages sent from the selected **Msg Router** to the selected **Neighbor** message router.

Sent Msgs/s	Refer to Solace documentation for more information.
Control Msgs/s	Refer to Solace documentation for more information.
Discards/s	Traces the number of discarded messages sent, per second, from the selected Msg Router to the selected Neighbor message router.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Time Settings	By default, the time range end point is the current time. To change the time range, click the Time Settings  and either: <ul style="list-style-type: none"> choose a Time range from 5 Minutes to 7 Days in the drop-down menu. specify begin/end dates using the calendar . specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

VPNs

You can view data for all VPNs configured on a specific message router in heatmap, table, or grid formats, or you can view data for a single VPN. Displays in this View are:

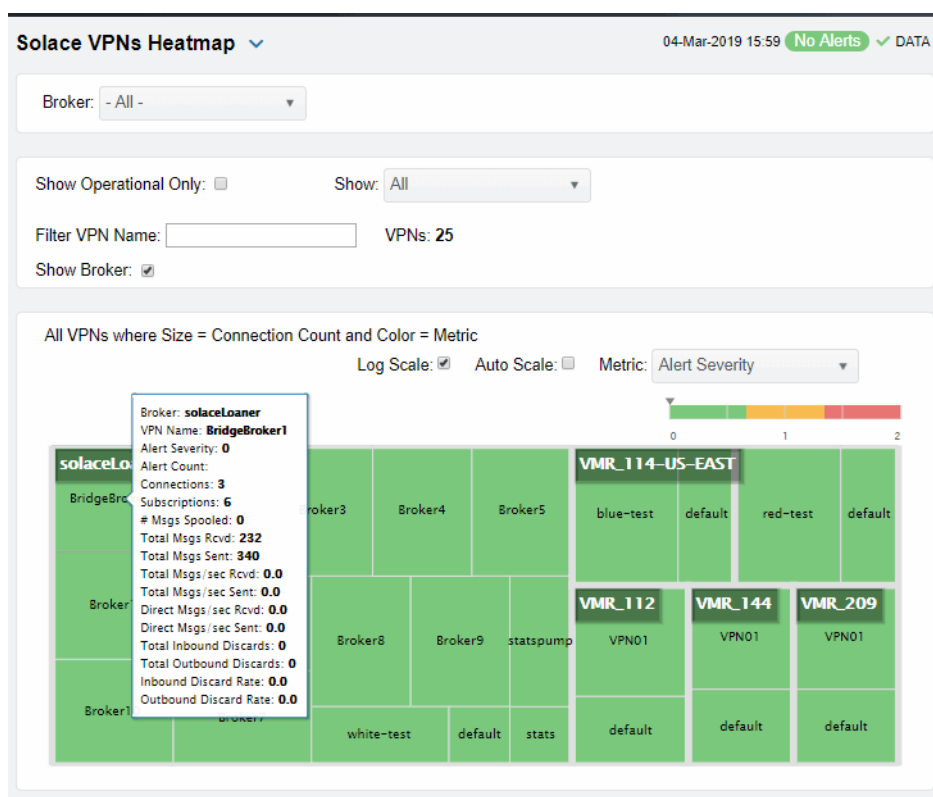
- [“VPNs Heatmap” on page 796](#): A color-coded heatmap view of the current status of all VPNs configured on a specific message router.
- [“VPNs Table” on page 800](#): A tabular view of all available data for all VPNs configured on a specific message router.
- [“VPNs Summary” on page 804](#): Current and historical metrics for a single VPN.

VPNs Heatmap

View the status of all VPNs configured on a specific message router in a heatmap format, which allows you to quickly identify VPNs with critical alerts. Each rectangle in the heatmap represents a VPN. The rectangle color indicates the alert state and rectangle size represents the number of connections.

Select a message router from the **Msg Router** drop-down menu, or enter a search string in the **Filter VPN Name** field, and select a metric from the **Metric** drop-down menu. Use the **Show Operational Only** check-box ☒ to include or exclude non-operational VPNs in the heatmap. Use the **Log Scale** and **Auto Scale** check-boxes ☒ to apply log or auto scale. Use the **Show Message Router** check-box ☒ to include or exclude message router names in the heatmap.

By default, this display shows **Alert Severity**, but you can mouse over a rectangle to see additional metrics. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the “VPNs Summary” display.

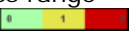









Title Bar (possible features are):


- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.


Operational	When checked, only shows operational message routers.
Filter VPN Name	Enter a string to show only VPNs with this string in their name.
Metric	Choose a metric to view in the display.
Alert Severity	<p>Visually displays the level at which the VPN has or has not exceeded its alarm level threshold. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Connections	<p>The total number of connections. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolVpnConnectionCountHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Subscriptions	<p>The total number of subscriptions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolVpnSubscriptionCountHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
# Msgs Spooled	<p>The total number of spooled messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterPendingMsgsHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

Total Msgs Rcvd

The total number of received messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of messages received in the heatmap. The middle value in the gradient bar indicates the average count.


The **Auto** flag does not impact this metric.

Total Msgs Sent

The total number of sent messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of messages sent in the heatmap. The middle value in the gradient bar indicates the average count.


The **Auto** flag does not impact this metric.

Total Msgs/ sec Rcvd

The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range.


When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Total Msgs/ sec Sent

The number of messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnOutboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range.


When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Total Bytes/ sec Rcvd

The number of bytes contained in messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range.


When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Total Bytes/ sec Sent

The number of bytes contained in direct messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolMsgRouterOutboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range.


When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Direct Msgs/sec Rcvd

The number of direct messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the average number of direct messages received per second in the heatmap. The middle value in the gradient bar indicates the average count.


The **Auto** flag does not impact this metric.

Direct Msgs/sec Sent

The number of direct messages sent per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the average number of direct messages sent per second in the heatmap. The middle value in the gradient bar indicates the average count.


The **Auto** flag does not impact this metric.

Total Inbound Discards

The total number of discarded inbound messages in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of discarded inbound messages in the heatmap. The middle value in the gradient bar indicates the average count.


The **Auto** flag does not impact this metric.

Total Outbound Discards

The total number of discarded outbound messages in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of discarded outbound messages in the heatmap. The middle value in the gradient bar indicates the average count.


The **Auto** flag does not impact this metric.

Inbound Discard Rate

The number of discarded inbound messages per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundDiscardRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Outbound Discard Rate

The number of discarded outbound messages per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnOutboundDiscardRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

VPNs Table

View data shown in the “VPNs Heatmap” display, as well as additional details, in a tabular format. Use this display to view all available data for each VPN associated with a specific message router.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a message router from the **Msg Router** drop-down menu. Each table row is a different VPN associated with the message router. Click a column header to sort column data in numerical or alphabetical order.

Sort data in numerical or alphabetical order on column headers. Use the check-box ☒ to include / exclude non-operational VPNs. Use the **Show** drop-down to see **All** VPNs, **Expired Only** or **Unexpired Only**. Enter a string to show only VPNs with this string in their name.

Double-click a row to drill-down and investigate in the “VPNs Summary” display.

Solace VPNs Table

04-Mar-2019 16:05 No Alerts DATA

Broker: - All - Less Columns

Show Operational Only: Show: All

Filter VPN Name: VPNs: 25

Broker	VPN Name	Alert Level	Alert Count	Connections	Operational
solaceLoaner	BridgeBroker1	✓		3	✓
solaceLoaner	Broker1	✓		3	✓
solaceLoaner	Broker10	✓		3	✓
solaceLoaner	Broker2	✓		3	✓
solaceLoaner	Broker3	✓		3	✓
solaceLoaner	Broker4	✓		3	✓
solaceLoaner	Broker5	✓		3	✓
solaceLoaner	Broker6	✓		3	✓
solaceLoaner	Broker7	✓		3	✓
solaceLoaner	Broker8	✓		3	✓
solaceLoaner	Broker9	✓		3	✓
solaceLoaner	default	✓		0	
solaceLoaner	stats	✓		0	
solaceLoaner	statspump	✓		1	✓
solaceLoaner	white-test	✓		1	✓

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.

Message Router	The name of the message router.
VPN Name	The name of the VPN.
Alert Level	<p>The maximum level of alerts in the row:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the VPN.
Connections	The total number of connections for the VPN.
Operational	When checked, this status indicates that the VPN is enabled and is operating normally.
Total Unique Subscriptions	The total number of unique subscriptions to the VPN.
Total Client Messages Rcvd	The total number of messages received from clients connected to the VPN.
Total Client Messages Sent	The total number of messages sent to clients connected to the VPN.
Total Client Bytes Rcvd	The total number of bytes contained in messages received from clients connected to the VPN.
Total Client Bytes Sent	The total number of bytes contained in messages sent to clients connected to the VPN.
Total Client Msgs/sec Rcvd	The total number of messages received per second from clients connected to the VPN.
Total Client Msgs /sec Sent	The total number of messages sent per second to clients connected to the VPN.
Total Client Bytes/sec Rcvd	The total number of bytes contained in messages received per second from clients connected to the VPN.
Total Client Bytes/sec Sent	The total number of bytes contained in messages sent per second to clients connected to the VPN.
Client Direct Msgs Rcvd	The total number of direct messages received from clients connected to the VPN.
Client Direct Msgs Sent	The total number of direct messages sent to clients connected to the VPN.
Client Direct Bytes Rcvd	The total number of bytes contained in direct messages received from clients connected to the VPN.
Client Direct Bytes Sent	The total number of bytes contained in direct messages sent to clients connected to the VPN.
Client Direct Msgs/sec Rcvd	The total number of direct messages received per second from clients connected to the VPN.
Client Direct Msgs/sec Sent	The total number of direct messages sent per second to clients connected to the VPN.
Client Direct Bytes/sec Rcvd	The total number of bytes contained in the direct messages received per second from clients connected to the VPN.
Client Direct Bytes/sec Sent	The total number of bytes contained in the direct messages sent per second to clients connected to the VPN.

Client NonPersistent Msgs Rcvd	The total number of non-persistent messages received from clients connected to the VPN.
Client NonPersistent Msgs Sent	The total number of non-persistent messages sent to clients connected to the VPN.
Client NonPersistent Bytes Rcvd	The total number of bytes contained in the non-persistent messages received from clients connected to the VPN.
Client NonPersistent Bytes Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client NonPersistent Msgs/sec Rcvd	The total number of non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Msgs/sec Sent	The total number of non-persistent messages sent per second to clients connected to the VPN.
Client NonPersistent Bytes/sec Rcvd	The total number of bytes contained in the non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Bytes/sec Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client Persistent Msgs Rcvd	The total number of persistent messages received from clients connected to the VPN.
Client Persistent Msgs Sent	The total number of persistent messages sent to clients connected to the VPN.
Client Persistent Bytes Rcvd	The total number of bytes contained in persistent messages received from clients connected to the VPN.
Client Persistent Bytes Sent	The total number of bytes contained in persistent messages sent to clients connected to the VPN.
Client Persistent Msgs/sec Rcvd	The total number of persistent messages received per second from clients connected to the VPN.
Client Persistent Msgs/sec Sent	The total number of persistent messages sent per second to clients connected to the VPN.
Client Persistent Bytes/sec Rcvd	The total number of bytes contained in the persistent messages received per second from clients connected to the VPN.
Client Persistent Bytes/sec Sent	The total number of bytes contained in the persistent messages sent per second to clients connected to the VPN.
Total In Discards	The total number of discarded incoming messages.
Total In Discards/sec	The number of discarded incoming messages per second.
Total Out Discards	The total number of discarded outgoing messages.
Total Out Discards/sec	The number of discarded outgoing messages per second.
Max Spool Usage (MB)	The maximum amount of disk storage (in megabytes) that can be consumed by all spooled message on the VPN.
Authentication Type	The defined authentication type on the VPN.

Expired

When checked, performance data about the VPN has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this
number of seconds
#
collector.sl.rtvview.sub=$solRowExpirationTime:4
5
collector.sl.rtvview.sub=$solRowExpirationTimeFo
rDelete:3600
```


In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Time Stamp

The date and time the row data was last updated.

VPNs Summary

Select a message router and a VPN to view details about alerts, connections/destinations, incoming messages and outgoing/pending messages for a single VPN.



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.

Alerts

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Up

Inbound/Outbound Msgs/s	The number of inbound/outbound messages per second.
Pending Msgs	The number of pending messages.
Inbound/Outbound KB/s	The number of inbound/outbound messages in KBs per second.
Current/Max Connections	The total number of current connections / maximum number of supported connections for the VPN.
Current Clients	The number of connected clients.
Current/Max Subscriptions	The total number of current subscribers and maximum number of supported subscribers for the VPN.

Inbound Msgs/s Trend Graphs

Traces the sum of inbound message processing for the selected VPN.



- **Pending Msgs**: The number of pending messages for the VPN.
- **Client Msgs/sec**: The rate of incoming messages (per second) from client.
- **Direct Client Msgs/sec**: The rate of direct incoming messages (per second) from the direct client.


Log Scale



Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

Longest Client Up Time

The number of days, hours and minutes for the longest, currently active, client connection.

Endpoints

The number of endpoints.

Persistent Msgs In	The total number of incoming persistent messages.
Persistent In Msgs/s	The number of incoming persistent messages per second.
Persistent Msgs Out	The total number of outgoing persistent messages.
Persistent Out Msgs/s	The number of outgoing persistent messages per second.
Total In Discards	The total number of incoming messages that were discarded.
Total In Discards/sec	The total number of incoming messages that were discarded, per second.
Critical/Warning Bridges	The number of critical alerts / warning alerts which also opens the Alerts Table .
Non-Persistent Msgs In	The total number of incoming non-persistent messages.
Non-Persistent In Msgs/s	The number of incoming non-persistent messages per second.
Non-Persistent Msgs Out	The total number of outgoing non-persistent messages.
Non-Persistent Out Msgs/s	The number of outgoing non-existent messages per second.
Total Out Discards	The total number of outgoing messages that were discarded.
Total Out Discards/sec	The total number of outgoing messages that were discarded, per second.
Direct Msgs In	The total number of incoming direct messages.
Direct In Msgs/s	The number of incoming direct messages per second.
Direct Msgs Out	The total number of outgoing direct messages.
Direct Out Msgs/s	The number of outgoing direct messages per second.
Expired	<p>When true, performance data about the VPN has not been received within the time specified (in seconds) in the \$solRowExpirationTime field in the conf\rtvapm_solmon.properties file. The \$solRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the .properties file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the Expired check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
Last Update	The date and time of the last data update.

Clients

These displays allow you to view the current and historical metrics for clients configured on a VPN. Displays in this View are:

- **"Clients Table"**: A tabular view of data for all clients configured on a VPN.
- **"Client Summary"**: Current and historical metrics for a single client configured on a VPN.

Clients Table

View VPN clients configured on all message routers, a single message router, all VPNs or a single VPN. Each table row is a different VPN client connection. Use the drop-down menus to show **All**, **Expired** or **Unexpired** clients as well as **All**, **Internal** or **Primary** clients (processes that run on the message router under the Solace OS). Enter a string for **Filter Client Name** to show only clients with this string in their name.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

This display is populated by two caches, SolClientsStats and SolClients. SolClientsStats provides most of the data. SolClients provides the static data. If the SolClients cache encounters an issue the static fields in this display are blank.

Double-click a row to drill-down and investigate in the **"Client Summary"** display.

Solace Broker Clients Table 05-Mar-2019 08:38 No Alerts DATA


Broker: - All - VPN: - All - Less Columns

Show Type: All Show: All Filter Client Name:

Clients: 101

Broker	VPN	Client Name	Alert Level	Alert Count	Sic Subsc
solaceLoaner	BridgeBroker1	#bridge/local/B114toSolDemo/solace/8798/16	✓		
solaceLoaner	BridgeBroker1	#bridge/local/B142toSolDemo/solace/8798/15	✓		
solaceLoaner	BridgeBroker1	#client	✓		
solaceLoaner	Broker1	#bridge/local/testBridgeToNoWhere/solace/8798/9	✓		
solaceLoaner	Broker1	#bridge/remote/B1_to_B2/v:solace/8796/0	✓		
solaceLoaner	Broker1	#client	✓		
solaceLoaner	Broker1	S-HOST10/5236/#00010001	✓		
solaceLoaner	Broker10	#bridge/local/B112toSolDemo/solace/8798/14	✓		
solaceLoaner	Broker10	#bridge/local/B144toSolDemo/solace/8798/12	✓		
solaceLoaner	Broker10	#bridge/local/B209toSolDemo/solace/8798/13	✓		
solaceLoaner	Broker10	#client	✓		
solaceLoaner	Broker10	S-HOST10/5152/#00010001	✓		
solaceLoaner	Broker10	S-HOST10/5448/#00010001	✓		
solaceLoaner	Broker2	#bridge/local/B1_to_B2/solace/8798/10	✓		
solaceLoaner	Broker2	#client	✓		
solaceLoaner	Broker2	S-HOST10/5212/#00010001	✓		
solaceLoaner	Broker3	#bridge/local/Bridge_loanerToVMR144/solace/879	✓		

Page 1 of 3 1 - 40 of 101 items

**Message Router**

Lists the name of the selected message router.

VPN

Lists the name of the selected VPN.

Client Name

The name of the client.

Alert Level

The maximum level of alerts in the row:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Alert Count

Total number of alerts for the client.

Slow Subscriber

This check box will be checked if the client consistently fails to consume their messages at the offered rate (which causes their egress queues to fill up).

Total Egress Flows

The total number of outgoing flows.

Total Ingress Flows

The total number of incoming flows.

Bind Requests

The number of bind requests made by the client.

Subscriptions

The total number of subscriptions.

Subscription Msgs Rcvd

The total number of messages received from subscriptions.

Subscription Msgs Sent

The total number of messages sent from subscriptions.

Type

Lists the type of alert.

Uptime

Lists the amount of time the client has been up and running.

Client ID

Lists the client ID.

Client UserName

Lists the user name for the client.

Client Address

The IP Address of the client.

Profile

The client profile that is assigned to the client.

ACL Profile

The access control list profile to which the client is assigned.

Description

Lists a description of the client.

Platform

Lists the platform of the client.

Software Version	The version of the platform.
Total Flows Out	The total number of outbound message flows for the client.
Total Flows In	The total number of inbound message flows for the client.
# Subscriptions	The number of subscribers connected to the client.
Add Sub Msgs Rcvd	The number of Add Subscription messages received.
Add Sub Msgs Sent	The number of Add Subscription Messages sent.
Already Exists Msgs Sent	Refer to Solace documentation for more information.
Assured Ctrl Msgs Rcvd	Refer to Solace documentation for more information.
Assured Ctrl Msgs Sent	Refer to Solace documentation for more information.
Total Client Msgs Rcvd	The total number of messages received by the client.
Total Client Msgs Sent	The total number of messages sent by the client.
Total Client Bytes Rcvd	The total number of bytes contained within the messages received by the client.
Total Client Bytes Sent	The total number of bytes contained within the messages sent by the client.
Total Client Msgs Rcvd/sec	The total number of messages received per second by the client.
Total Client Msgs Sent/sec	The total number of messages sent per second by the client.
Total Client Bytes Rcvd/sec	The total number of bytes contained within the messages received per second by the client.
Total Client Bytes Sent/sec	The total number of bytes contained within the messages sent per second by the client.
Ctl Bytes Rcvd	The number of control data bytes received by the client.
CTL Bytes Sent	The number of control data bytes sent by the client.
Ctl Msgs Rcvd	The number of control data messages received by the client.
Ctl Msgs Sent	The number of control data messages sent by the client.
Client Data Bytes Rcvd	The number of bytes contained within the data messages received by the client.
Client Data Bytes Sent	The number of bytes contained within the data messages sent by the client.
Client Data Msgs Rcvd	The number of data messages received by the client.
Client Data Msgs Sent	The number of data messages sent by the client.
Client Direct Msgs Rcvd	The number of direct messages received by the client.
Client Direct Msgs Sent	The number of direct messages sent by the client.
Client Direct Bytes Rcvd	The number of bytes contained within direct messages received by the client.
Client Direct Bytes Sent	The number of bytes contained within direct messages sent by the client.

Client Direct Msgs Rcvd/sec	The number of direct messages received per second by the client.
Client Direct Msgs Sent/sec	The number of direct messages sent per second by the client.
Client Direct Bytes Rcvd/sec	The number of bytes contained within the messages received per second by the client.
Client Direct Bytes Sent/sec	The number of bytes contained within the messages sent per second by the client.
Client NonPersistent Msgs Rcvd	The number of non-persistent messages received by the client.
Client NonPersistent Msgs Sent	The number of non-persistent messages sent by the client.
Client NonPersistent Bytes Rcvd	The number of bytes contained within the non-persistent messages received by the client.
Client NonPersistent Bytes Sent	The number of bytes contained within the non-persistent messages sent by the client.
Client NonPersistent Msgs Rcvd/sec	The number of non-persistent messages received per second by the client.
Client NonPersistent Msgs Sent/sec	The number of non-persistent messages sent per second by the client.
Client NonPersistent Bytes Rcvd/sec	The number of bytes contained within the non-persistent messages received per second by the client
Client NonPersistent Bytes Sent/sec	The number of bytes contained within the non-persistent messages sent per second by the client
Client Persistent Msgs Rcvd	The number of persistent messages received by the client.
Client Persistent Msgs Sent	The number of persistent messages sent by the client.
Client Persistent Bytes Rcvd	The number of bytes contained within the persistent messages received by the client.
Client Persistent Bytes Sent	The number of bytes contained within the persistent messages sent by the client.
Client Persistent Msgs Rcvd/sec	The number of persistent messages received per second by the client.
Client Persistent Msgs Sent/sec	The number of persistent messages sent per second by the client.
Client Persistent Bytes Rcvd/sec	The number of bytes contained within the persistent messages received per second by the client.
Client Persistent Bytes Sent/sec	The number of bytes contained within the persistent messages sent per second by the client.
Denied Dup Clients	Refer to Solace documentation for more information.
Denied Subscribe Permission	The number of denied subscription requests due to improper permissions.
Denied Subscribe Topic-ACL	The number of denied subscriptions to topics due to the fact that the client requesting was not on the Access Control List.
Denied Unsubscribe Permission	The number of denied unsubscribe requests due to improper permissions.
Denied Unsubscribe Topic-ACL	The number of denied unsubscribe requests to topics due to the fact that the client requesting was not on the Access Control List.

DTO Msgs Rcvd	The number of Deliver-To-One messages received by the client.
Egress Compressed Bytes	The number of compressed bytes contained within outgoing messages.
Ingress Compressed Bytes	The number of compressed bytes contained within incoming messages.
Total Ingress Discards	The total number of discarded incoming messages.
Total Egress Discards	The total number of discarded outgoing messages.
Total Ingress Discards/sec	The total number of discarded incoming messages per second.
Total Egress Discards/sec	The total number of discarded outgoing messages per second.
Keepalive Msgs Rcvd	The number of Keepalive messages received by the client.
Keepalive Msgs Sent	The number of Keepalive messages sent by the client.
Large Msgs Rcvd	The number of large messages received by the client.
Login Msgs Rcvd	The number of login message received by the client.
Max Exceeded Msgs Sent	The number of responses sent by the client informing the connected message router(s) that the number of the message(s) sent exceeded the maximum allowed.
Not Enough Space Msgs Sent	The number of responses sent by the client informing the connected message router(s) that the size of the message(s) sent exceeded the maximum allowable size, or that the message caused the client's Local Spool Quota to exceed the maximum amount of space.
Not Found Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Add Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Remove Msgs Sent	Refer to Solace documentation for more information.
Remove Subscription Msgs Rcvd	The number of remove subscription requests received by the client.
Remove Subscription Msgs Sent	The number of remove subscription requests sent by the client.
Subscribe Client Not Found	The number of subscription requests for clients that were not found.
Unsubscribe Client Not Found	The number of unsubscribe requests for clients that were not found.
Update Msgs Rcvd	Refer to Solace documentation for more information.
Update Msgs Sent	Refer to Solace documentation for more information.

Expired

When checked, performance data about the client has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the client. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after
this number of seconds
#
collector.sl.rtvapm.sub=$solRowExpirationTime:
45
collector.sl.rtvapm.sub=$solRowExpirationTimeF
orDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Timestamp

The date and time the row of data was last updated.

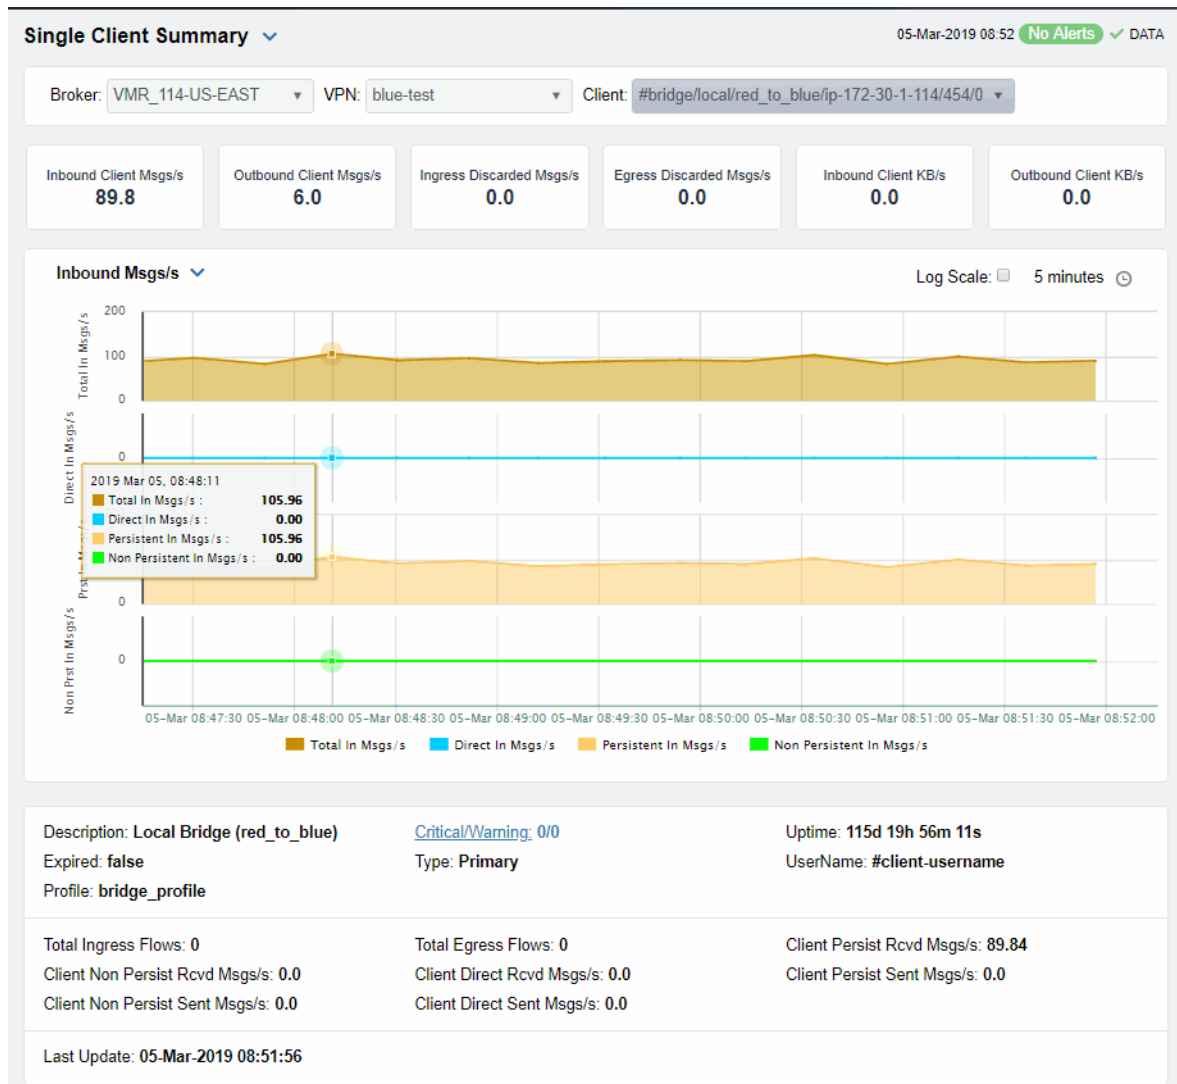
Client Summary

View the current and historical metrics for a single VPN client.

Select a message router, VPN and client from the drop-down menus. You can view the **Client Type**, the **User Name**, the **Client ID**, the associated **Platform**, the current **Up Time**, and additional information specific to the client. You can also view the total number of incoming and outgoing messages, as well as the number of incoming and outgoing persistent, non-persistent, direct, and discarded messages.

This display is populated by two caches, SolClientsStats and SolClients. SolClientsStats provides most of the data. SolClients provides the static data. If the SolClients cache encounters an issue the graphic elements that have no data are replaced with **N/A**.

This display also includes a trend graph containing the current and historical incoming messages per second, outgoing messages per second, incoming direct messages per second, and outgoing direct messages per second.



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.

Inbound Client Msgs /sec The number of incoming client messages per second.

Outbound Client Msgs /sec The number of outgoing client messages per second.

Ingress Discarded Msgs /sec The number of discarded ingress messages per second.

Egress Discarded Msgs /sec	The number of discarded egress messages per second.
Inbound Client KB/sec	The amount of incoming data from the client in KBs per second.
Outbound Client KB/sec	The amount of outgoing data for the client in KBs per second.

Trend Graphs

Traces the sum of message processing for the selected client.



- **Total In Msgs/sec**: The number of incoming messages (per second) for the client.
- **Dir-In Msgs/sec**: The number of incoming direct messages (per second) for the client.
- **Persistent In Msgs/sec**: The number of incoming persistent messages (per second) for the client.
- **Non Persistent In Msgs/sec**: The number of incoming non-persistent messages (per second) for the client.


Log Scale



Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

Description

The description of the client.

Expired

When checked, performance data about the client has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the client. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number of
seconds
#
collector.sl.rtvview.sub=$solRowExpirationTime:45
collector.sl.rtvview.sub=$solRowExpirationTimeForDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Profile

The client's profile.

Total Ingress Flows

The number of inflows coming to the client.

Persistent Msgs In/sec

The number of persistent incoming messages per second.

Persistent Msgs Out/sec

The number of persistent outgoing messages per second.

Last Update

The date and time of the last data update.

Critical/Warning

The number of critical alerts / warning alerts which also opens the **Alerts Table**.

Non Persistent Msgs In/sec

The number of non-persistent incoming messages per second.

NonPersistent Msgs Out/sec

The number of non-persistent outgoing messages per second.

Uptime

If the VPN's **Local Status** is **Up**, this field displays the length of time that the VPN has been up and running.

Username

The client's user name.

Bind Requests

The number of bind requests received by the client.

Direct In Msgs /sec

The number of non-persistent incoming messages per second.

Direct Out Msgs /sec

The number of non-persistent outgoing messages per second.

Bridges

These displays provide process data for bridges configured on a VPN. Displays in this View are:

- **"Bridges Table"**: A tabular view of all available process performance data for all bridges configured on a VPN.
- **"Bridges - Diagram"**: Topological view of Solace network bridges that shows bridge message router connections and status.
- **"Bridge Summary"**: Current and historical metrics for a single bridge.

Bridges Table

This display allows you to view data for all bridges configured for a VPN.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a message router and VPN from the drop-down menus. Use the check-boxes ☒ to include / exclude **Enabled** and **Expired** bridges. Each table row is a different bridge.

Rows listing bridges that are disabled or expired display with a shaded background. Double-click a row to drill-down and investigate in the [“Bridge Summary”](#) display.

Solace Bridges Table

05-Mar-2019 09:01No AlertsDATA

Broker: - All -VPN: blue-testLess Columns

Show Enabled Only: Show: AllFilter Bridge Name: Bridges: 2

Broker	Local VPN	Bridge Name	Remote VPN	Remote R
VMR_114-US-EAST	blue-test	#bridge/v:solace/BridgeBroker1/15	BridgeBroker1	v:solace
VMR_114-US-EAST	blue-test	red_to_blue	red-test	v:ip-172-30-1-1

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.

Message Router	Displays the name of the message router
Local VPN	The name of the local VPN.
Bridge Name	The name of the bridge.
Alert Level	<div>The current level of alerts in the row.<div><div>Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</div><div>Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</div><div>Green indicates that no metrics have exceeded their alert thresholds.</div></div></div>
Alert Count	The total number of active alerts for the process.

Remote VPN	The name of the remote VPN that is connected to the local VPN via the bridge.
Remote Router	The name of the remote message router.
Admin State	Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).
Inbound Operational State	The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)
Outbound Operational State	The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)
Queue Operational State	The current operational status of the queue.
Connection Establisher	Indicates whether the administrator created and configured the bridge directly on the message router using SolAdmin or the command line interface, or indirectly from another message router.
Redundancy	Displays whether the bridge is the primary bridge, the backup bridge, the static bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (none).
Uptime	The current amount of time in which the bridge has been up and running.
Client Name	The name of the client.
Connected Via Addr	The local IP address and port used for the bridge.
Connected Via Interface	The name of the network interface used for the bridge.
Client Direct Bytes Rcvd	The number of bytes contained within direct messages received by the client via the bridge.
Client Direct Bytes/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Bytes Sent	The number of bytes contained within direct messages sent by the client via the bridge.
Client Direct Bytes/sec Sent	The number of bytes contained within direct messages sent per second by the client via the bridge.
Client Direct Msgs/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Msgs Sent	The number of direct messages sent by the client via the bridge.
Client Direct Msgs/sec Sent	The number of direct messages sent per second by the client via the bridge.
Client NonPersistent Bytes Rcvd	The number of bytes contained within non-persistent messages received by the client via the bridge.
Client NonPersistent Bytes/sec Rcvd	The number of bytes contained within non-persistent messages received per second by the client via the bridge.
Client NonPersistent Bytes Sent	The number of bytes contained within non-persistent messages sent by the client via the bridge.
Client NonPersistent Bytes/sec Sent	The number of bytes contained within non-persistent messages sent per second by the client via the bridge.
Client NonPersistent Msgs Rcvd	The number of non-persistent messages received by the client via the bridge.

Client NonPersistent Msgs/sec Rcvd	The number of non-persistent messages received per second by the client via the bridge.
Client NonPersistent Msgs Sent	The number of non-persistent messages sent by the client via the bridge.
Client NonPersistent Msgs/sec Sent	The number of non-persistent messages sent per second by the client via the bridge.
Client Persistent Bytes Rcvd	The number of bytes contained within persistent messages received by the client via the bridge.
Client Persistent Bytes/sec Rcvd	The number of bytes contained within persistent messages received per second by the client via the bridge.
Client Persistent Bytes Sent	The number of bytes contained within persistent messages sent by the client via the bridge.
Client Persistent Bytes/sec Sent	The number of bytes contained within persistent messages sent per second by the client via the bridge.
Client Persistent Msgs Rcvd	The number of persistent messages received by the client via the bridge.
Client Persistent Msgs /sec Rcvd	The number of persistent messages received per second by the client via the bridge.
Client Persistent Msgs Sent	The number of persistent messages sent by the client via the bridge.
Client Persistent Msgs/sec Sent	The number of persistent messages sent per second by the client via the bridge.
Total Client Bytes Rcvd	The number of bytes contained within all messages received by the client via the bridge.
Total Client Bytes/sec Rcvd	The number of bytes contained within all messages received per second by the client via the bridge.
Total Client Bytes Sent	The number of bytes contained within all messages sent by the client via the bridge.
Total Client Bytes/sec Sent	The number of bytes contained within all messages sent per second by the client via the bridge.
Total Client Msgs Rcvd	The total number of all messages received by the client via the bridge.
Total Client Msgs/sec Rcvd	The total number of all messages received per second by the client via the bridge.
Total Client Msgs Sent	The total number of all messages sent by the client via the bridge.
Total Client Msgs/sec Sent	The total number of all messages sent per second by the client via the bridge.
Total Out Discards	The total number of discarded outgoing messages sent by the client via the bridge.
Total Out Discards/sec	The total number of discarded outgoing messages sent per second by the client via the bridge.
Total In Discards	The total number of discarded incoming messages received by the client via the bridge.
Total In Discards/sec	The total number of discarded incoming messages received per second by the client via the bridge.

Expired

When checked, performance data about the bridge has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapi_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the bridge. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this
number of seconds
#
collector.sl.rtvapi.sub=$solRowExpirationTime:4
5
collector.sl.rtvapi.sub=$solRowExpirationTimeFo
rDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.





Timestamp

The date and time the row of data was last updated.

Bridges - Diagram

Use this topology view to monitor the health of your network bridges and VPNs. Quickly identify bridge and VPN connections, their health status and which resources their performance impacts. Drag and drop objects to arrange them on the screen (doing so does not logically impact the network bridges and VPNs). Arrows show the connections between VPNs and bridges.

Each object is a network bridge or VPN. Each is labeled with their name and color coded as follows:

-  Red indicates that the object has one or more alerts in a critical state.
-  Yellow indicates that the object has one or more alerts in a warning state.
-  Green indicates that there are no alerts on the object.
-  Gray indicates that the object is off-line.

Save: Saves the arrangement of the objects.

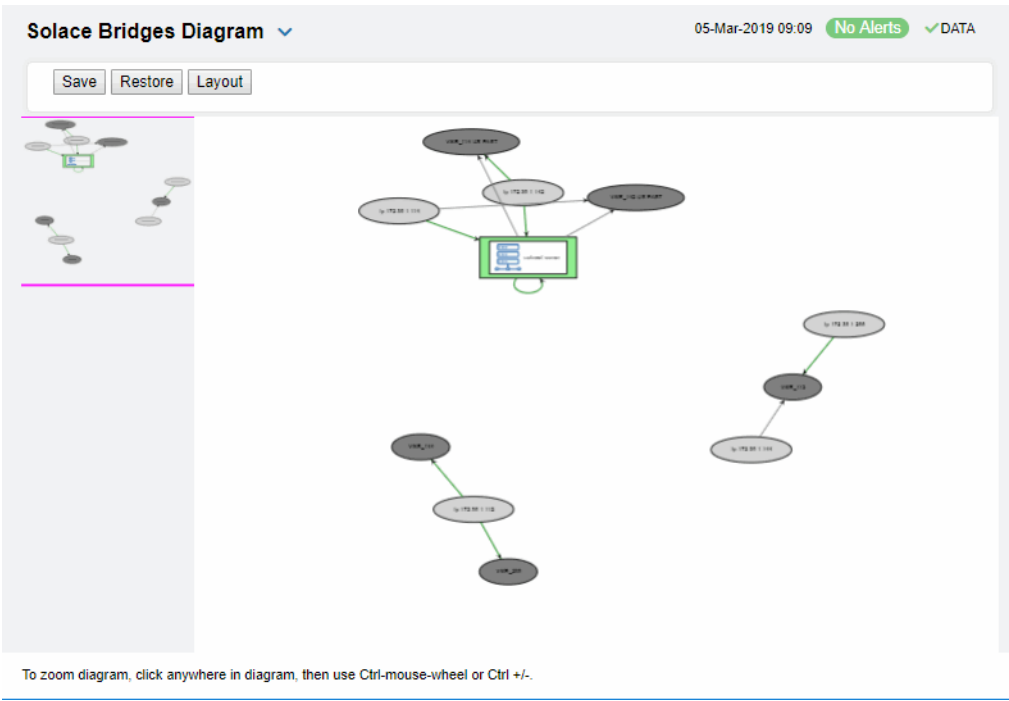
Restore: Returns objects to their previous positions.

Layout: Toggles between two types of layouts. One layout positions objects to the right so you might scroll in that direction to see them. The other layout pulls all the objects close together to the left, vertically, in hierarchical order.

Look at the miniature view in (upper left) to see all objects in either layout. Or zoom into the display using **Ctrl+/-** or **Ctrl+** mouse wheel.


Drill-down to investigate in the [“Bridges Table”](#).


To monitor network message routers, VMRs and servers, see the “[Neighbors Diagram](#)”.



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 [dropdown], Table open commonly accessed displays.

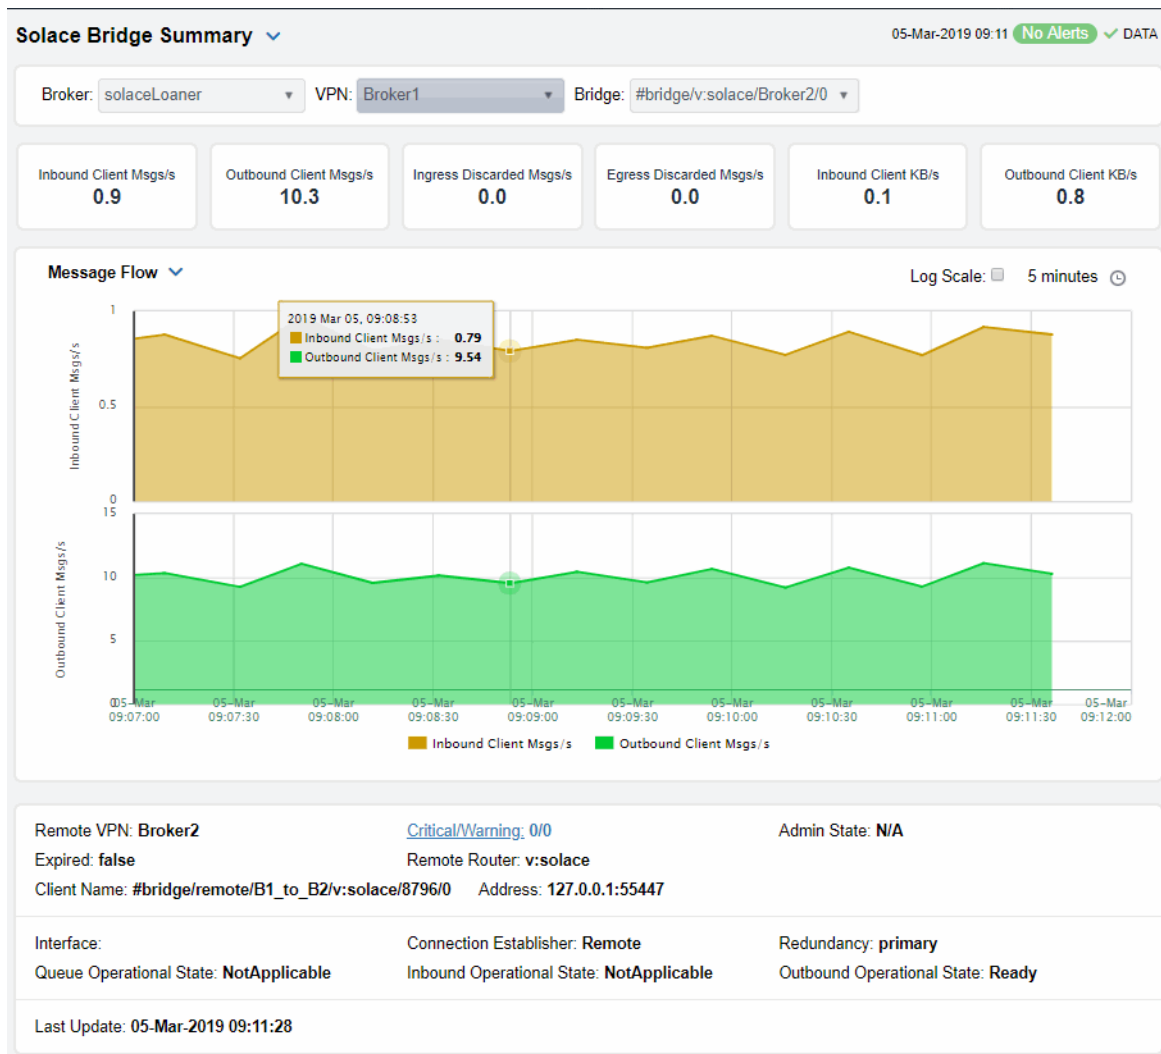
 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **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.

Bridge Summary

This display allows you to view performance details for a specific bridge configured on a VPN.

Select a message router, VPN, and a bridge from the drop-down menus, and use the **Time-Range** to “zoom-in” or “zoom-out” on a specific time frame in the trend graph.



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.

Inbound Client Msgs/s

The number of client messages received per second.

Outbound Client Msgs/s

The number of client messages sent per second.

Ingress Discarded Client Msgs/s

The number of discarded ingress messages per second.

Egress Discarded Msgs/s

The number of discarded egress messages per second.

Inbound Client KB/s

The amount of incoming client data, in KB per second.

Outbound Client KB/s

The amount of outgoing client data, in KB per second.

Messages Flow Trend Graphs

Traces the sum for the selected client.



- **Inbound Client Msgs/s**: The number of client messages received per second.
- **Outbound Client Msgs/s**: The number of client messages sent per second.


Log Scale



Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

Remote VPN

The name of the remote VPN that is connected to the local VPN via the bridge.

Expired

When true, performance data about the bridge has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapi_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the bridge. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number
# of seconds
#
collector.sl.rtvapi.sub=$solRowExpirationTime:45
collector.sl.rtvapi.sub=$solRowExpirationTimeForDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Address

The IP address.

Interface	The interface ID.
Queue Operational State	Refer to Solace documentation for more information.
Last Update	The date and time of the last data update.
Critical/Warning	The number of critical alerts / warning alerts which also opens the Alerts Table .
Remote Router	The remote message router.
Conn Establisher	Refer to Solace documentation for more information.
Inbound Operational State	The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)
Admin State	Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).
Client Name	The name of the client.
Redundancy	Indicates whether the bridge is the primary bridge, the backup bridge, the static bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (none).
Outbound Op State	The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

Endpoints

These displays list data for one or more endpoints configured on a VPN. Displays in this View are:

- ["Endpoints Table"](#)
- ["Endpoint Summary"](#)

Endpoints Table

View all endpoints configured on a VPN. Each row in the table lists the details for a specific endpoint.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a message router and VPN from the drop-down menus. Filter the table using the **Show Ingress Config Status Down Only** check-box ☒ and use the **Show** drop-down menus to include **All**, **Expired** or **Unexpired**.

You can click a column header to sort column data in numerical or alphabetical order, or double-click a row to drill-down and investigate in the “Endpoint Summary” display.

Solace Endpoints Table

05-Mar-2019 09:13 No Alerts DATA

Broker: solaceLoaner VPN: Broker1 Less Columns

Show Ingress Config Status Down Only: Show: All Filter Endpoint Name : Endpoints: 2

Broker	VPN	Endpoint Name	Alert Level	Alert Count	Bind Count	Spooled Messages	Cur Spool Usage MB	High-Wa Mark M
solaceLoaner	Broker1	bridgeq	✓		1	2	0.0	
solaceLoaner	Broker1	q1	✓		0	0	0.0	

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.


Message Router	Displays the name of the message router
VPN	The name of the VPN.
Endpoint Name	The name of the endpoint.
Alert Level	<div>The current alert severity in the row.<div><div>●</div> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</div><div><div>●</div> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</div><div><div>●</div> Green indicates that no metrics have exceeded their alert thresholds.</div></div>
Alert Count	The total number of active alerts for the endpoint.
Bind Count	The total number of binds connected to the endpoint.
Endpoint Type	The type of endpoint (either queue or topic).

Durable	Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an message router restart and are automatically restored as part of an message router's backup and restoration process.
In Config Status	Refer to Solace documentation for more information.
Out Config Status	Refer to Solace documentation for more information.
Type	Refer to Solace documentation for more information.
Access Type	Refer to Solace documentation for more information.
Pending Messages	The total number of pending messages on the endpoint.
Spool Usage (MB)	The total spool usage consumed on the endpoint (in megabytes).
High Water Mark (MB)	The highest level of spool usage on the endpoint (in megabytes).
In Selector	Refer to Solace documentation for more information.
Out Selector	Refer to Solace documentation for more information.
Expired	<p>When checked, performance data about the endpoint has not been received within the time specified (in seconds) in the \$solRowExpirationTime field in the conf\rtvapm_solmon.properties file. The \$solRowExpirationTimeForDelete field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the endpoint. To view/edit the current values, modify the following lines in the .properties file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the Expired check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
Time Stamp	The date and time the row of data was last updated.

Endpoint Summary

This display allows you to view endpoint information, message data, and a trend graph for pending and spool messages for a specific endpoint configured on a VPN. Choose a message router, a VPN, and an endpoint from the drop-down menus, and use the **Time Settings** to "zoom-in" or "zoom-out" on a specific time frame in the trend graph.

This display is provided by default and should be used if you do not want to collect message spool data for specific VPNs. However, if you do want to configure message spool monitoring for specific VPNs, then you should use the **Single Endpoint Summary Rates** display instead, which is not included in the navigation tree by default.



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.

- Pending Messages**
Spool Usage (MB)
Spool Memory HWM MB
- The total number of pending messages on the endpoint.
The current spool usage consumed on the endpoint (in megabytes).
Refer to Solace documentation for more information

Expired

When **true**, performance data about the endpoint has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapi_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the endpoint. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number
of seconds
#
collector.sl.rtvapi.sub=$solRowExpirationTime:45
collector.sl.rtvapi.sub=$solRowExpirationTimeForDelete
:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Durable

Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an message router restart and are automatically restored as part of an message router's backup and restoration process.

Bind Count

The total number of binds connected to the endpoint.

Trend Graphs

Traces the sum of metrics for the endpoint.

- **Spooled Msgs**: The amount of spooled messages, in megabytes.
- **Cur Spool Usage**: The amount of space used by spooled messages, in megabytes.

Log Scale



Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Base at Zero



Select to use zero (**0**) as the Y axis minimum for all graph traces.


Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

Endpoint Type	The type of endpoint.
Egress Config Status	Refer to Solace documentation for more information.
Egress Selector Present	Refer to Solace documentation for more information.
Last Update	The date and time of the last data update.
Critical/Warning	The number of critical alerts / warning alerts which also opens the Alerts Table .
Access Type	Refer to Solace documentation for more information.
Ingress Config Status	Refer to Solace documentation for more information.
Ingress Selector Present	Refer to Solace documentation for more information.

Capacity

These displays provide current message router capacity metrics, alert count and severity at the message router level. Displays in this View are:

- **"Capacity Table"**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for all message routers.
- **"Capacity - Summary"**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for a specific message router.
- **"Capacity Trends"**: View the message router capacity data for a specific message router in a trend graph format.

Capacity Table

View current and HWM (high water mark for the last 30 days) capacity utilization data for all message routers.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

You can view client, spool usage, incoming message, outgoing message, incoming bytes, and outgoing bytes data for the message router. Each table row is a different message router.

Double-click a row to drill-down and investigate in the “Capacity - Summary” display.

Broker Capacity Table 05-Mar-2019 09:19 No Alerts ✓ DATA

Show: All Brokers: 1 [Less Columns](#)

Broker	Alert Level	Alert Count	Current Client Connections	Connections HWM	Connections Max	Connections Reserved	Connections Used %
solaceLoaner	✓		35	35	9,000	135,000	

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.

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.

Message Router

The name of the message router.

Alert Level

The maximum level of alerts in the row:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Alert Count

The total number of active alerts.

Current Client Connections

The current number of clients connected.

Connections HWM

The greatest number of connections in the last 30 days.

Connections Max

The greatest number of connections since the message router last started.

Connections Reserved

The current number of reserved connections.

Connections Used %

The current amount of connections used, in percent.

Connections Used HWM %

The greatest amount of connections used, in percent, in the last 30 days.

Cur Spool Usage MB

The current amount of used spool disk, in megabytes.

Cur Spool Usage HWM	The greatest amount of spool disk used in the last 30 days.
Spool Disk Allocated	The amount of allocated spool disk.
Spool Reserved	The amount of reserved spool disk.
Current Spool Usage %	The current amount of used spool disk, in percent.
Current Spool Usage % HWM	The greatest amount of used spool disk in the last 30 days, in percent.
Delivered Unacked Msgs Util %	Refer to Solace documentation for more information.
Ingress Flow Count	The number of ingress flows.
Ingress Flow HWM	The greatest number of ingress flows in the last 30 days.
Ingress Flows Allowed	The maximum number of ingress flows allowed.
Ingress Flow Count %	The amount of ingress flows in percent.
Ingress Flow Count HWM %	The greatest amount of ingress flows in the last 30 days, in percent.
Ingress Msgs/s	The number of ingress messages per second.
Ingress Msgs/s HWM	The greatest number of ingress messages per second in the last 30 days.
Max Ingress Msgs/s	The maximum number of ingress flows per second allowed.
Ingress Msgs %	The amount of ingress messages in percent.
Ingress Msgs/s HWM %	The greatest amount of ingress messages in the last 30 days, in percent.
Cur Egress Msgs/s	The number of egress messages per second.
Egress Msgs/s HWM	The greatest number of egress messages per second in the last 30 days.
Max Egress Msgs/s	The maximum number of egress flows per second allowed.
Egress Msgs %	The amount of egress messages in percent.
Egress Msgs/s HWM %	The greatest amount of ingress messages in the last 30 days, in percent.
Cur Egress Bytes/s	The amount of egress in bytes per second.
Egress Bytes/s HWM	The greatest amount of egress, in bytes per second, in the last 30 days, in percent.

Expired

When checked, performance data about the VPN has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapi_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number of
seconds
#
collector.sl.rtvapi.sub=$solRowExpirationTime:45
collector.sl.rtvapi.sub=$solRowExpirationTimeForDelete:3600
```

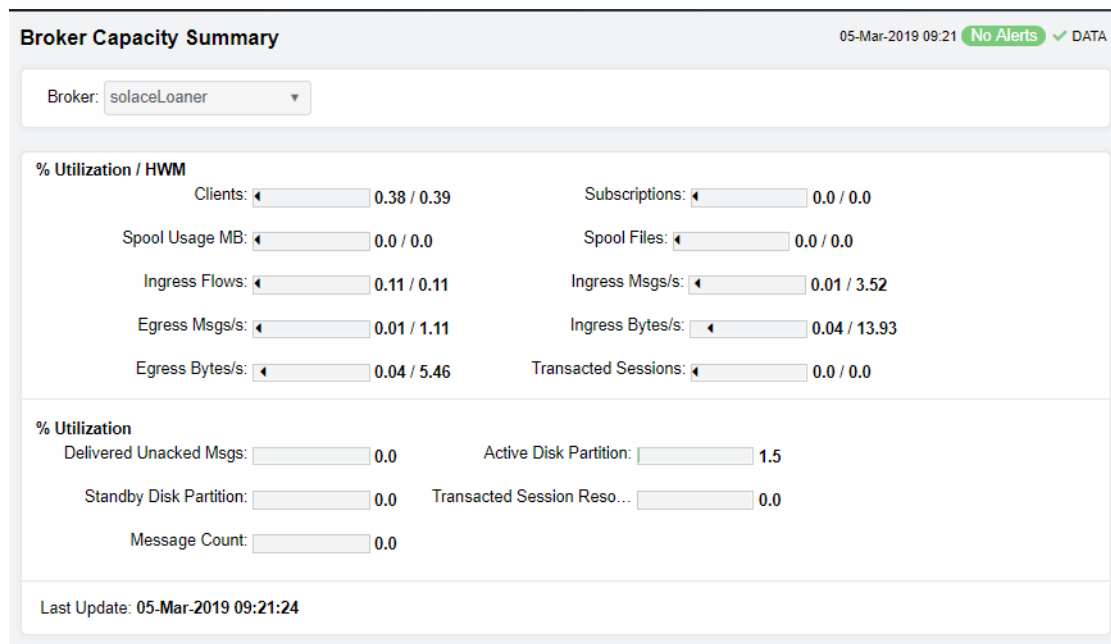
In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.


Time Stamp

The date and time the row of data was last updated.

Capacity - Summary

This display, a pivoted view of the “[Capacity Table](#)”, allows you to view current and HWM (high water mark for the last 30 days) capacity utilization data for a single message router. Select a message router from the drop-down menu to view client, spool usage, incoming message, outgoing message, incoming bytes, and outgoing bytes data for the message router.

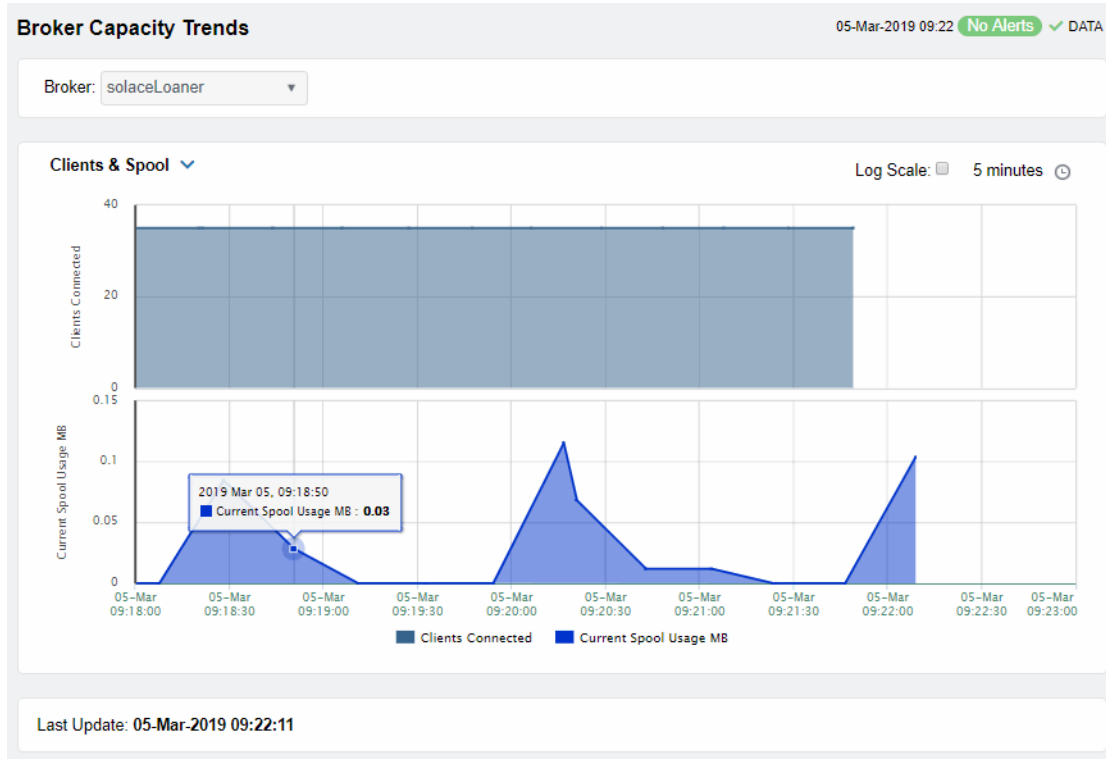




% Utilization/HWM	These values show high water marks (peak capacity utilization) for the last 30 days.
Clients	The current number of clients connected to the message router.
Spool Files	The highest number of spool files on the message router in the past 30 days.
Egress Msgs/s	The highest number of outgoing messages per second on the message router in the past 30 days.
Transacted Sessions	The highest number of transacted sessions on the message router in the last 30 days.
Subscriptions	The highest number of subscriptions on the message router in the last 30 days.
Ingress Flows	The highest number of inflows on the message router in the last 30 days.
Ingress Bytes/s	The highest amount of inflows, in bytes per second, on the message router in the past 30 days.
Spool Usage MB	The highest amount of spool utilization, in megabytes per second, on the message router in the past 30 days.
Ingress Msgs/s	The highest number of incoming messages per second on the message router in the past 30 days.
Egress Bytes/s	The highest number of outgoing messages per second on the message router in the past 30 days.
% Utilization	These values show current capacity utilization.
Delivered Unacked Msgs	The current number of delivered messages that were not acknowledged divided by the maximum number of delivered messages that were not acknowledged allowed on the message router.
Transacted Sessions Reso...	The current number of transacted sessions that were resolved on the message router.
Active Disk Partition	The percentage of available active disk partition that is used.
Message Count	The current number of messages on the message router.
Standby Disk Partition	The percentage of available standby disk partition that has been used.
Last Update	The date and time of the last data update.

Capacity Trends

This display allows you to view a trend graph that traces message router performance data for clients & spool data, message flow and throughput. Select a message router and a performance metric from the drop-down menus.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.



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.

Clients & Spool

The trend graph traces the following performance metrics:

Clients Connected: The current number of clients connected to the message router.




Current Spool Usage: The current spool usage, in megabytes, on the message router.


Message Flow



The trend graph traces the following:


Ingress Msgs/sec: The number of incoming messages per second on the message router.

Egress Msgs/sec: The number of outgoing messages per second on the message router.

- Throughput** The trend graph traces the following:
- Ingress KB/sec:** The amount of incoming per second, in KB, on the message router.
 - Egress KB/sec:** The number of outgoing data per second, in KB, on the message router.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Settings** By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:
- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
 - specify begin/end dates using the calendar .
 - specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

Syslog

The display in this View provides a tabular list of all Syslog events:

- [“All Syslog Events Table”](#): View all Syslog events for all your Solace message routers.

All Syslog Events Table

This table lists all Syslog events collected from one or all Solace message routers. Each row in the table is a different message. Filter messages per single Solace message router or all message routers (choose **All Hosts** from the **Source** drop-down menu), a single tag or **All Tags**, a single severity level or all levels (choose **All Levels** from the **Severity** drop-down menu), and specify a **Time Settings**.

Click a column header to sort column data in numerical, alphabetical or chronological order.

All Syslog Events - Table

Source: All Hosts Tag: All Tags

Count: 4,000 Severity: All Levels Time Range: 5 Mins

Timestamp	Message Timestamp	Host Address	Facility	Severity	Tag	Message Text
15-Feb-2016 07:27:07.175	15-Feb-2016 07:27:07.000	192.168.220.5	local3	NOTICE	solace	solace:NOTICE:SYSTEM:SYSTEM:AUTHENTICATION_SESSION_CLOSE
15-Feb-2016 07:27:07.111	15-Feb-2016 07:27:07.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:07.021	15-Feb-2016 07:27:07.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:06.465	15-Feb-2016 07:27:06.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:06.332	15-Feb-2016 07:27:06.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:06.117	15-Feb-2016 07:27:06.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:05.034	15-Feb-2016 07:27:05.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:04.325	15-Feb-2016 07:27:04.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:04.300	15-Feb-2016 07:27:04.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:04.204	15-Feb-2016 07:27:04.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:03.563	15-Feb-2016 07:27:03.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:03.102	15-Feb-2016 07:27:03.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:02.319	15-Feb-2016 07:27:02.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:01.451	15-Feb-2016 07:27:01.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:00.723	15-Feb-2016 07:27:00.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:27:00.155	15-Feb-2016 07:27:00.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:59.974	15-Feb-2016 07:26:59.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:59.949	15-Feb-2016 07:26:59.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:59.868	15-Feb-2016 06:47:47.000	192.168.220.5	local3	NOTICE	solace	solace:NOTICE:SYSTEM:SYSTEM:AUTHENTICATION_SESSION_OPE
15-Feb-2016 07:26:59.014	15-Feb-2016 07:26:59.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:58.601	15-Feb-2016 07:26:58.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:57.662	15-Feb-2016 07:26:57.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:57.174	15-Feb-2016 06:47:45.000	192.168.220.5	local3	NOTICE	solace	solace:NOTICE:SYSTEM:SYSTEM:AUTHENTICATION_SESSION_CLO
15-Feb-2016 07:26:56.869	15-Feb-2016 07:26:56.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:56.641	15-Feb-2016 07:26:56.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:56.496	15-Feb-2016 07:26:56.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:56.214	15-Feb-2016 07:26:56.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:55.507	15-Feb-2016 07:26:55.000	192.168.220.110	local3	INFO	S-HOST10	event CLIENT CLIENT CLIENT CONNECT vpcndi numConnHighCler
15-Feb-2016 07:26:54.926	15-Feb-2016 07:26:54.000	192.168.220.110	local3	INFO	S-HOST10	logger AFWlab-128-17_1 Start of action: Testing event CONNECTIONS
15-Feb-2016 07:26:54.854	15-Feb-2016 07:26:54.000	192.168.220.110	local3	INFO	S-HOST10	logger AFWlab-128-17_1 End of action
15-Feb-2016 07:26:54.830	15-Feb-2016 07:26:54.000	192.168.220.110	local3	INFO	S-HOST10	event SYSTEM:SYSTEM:CHASSIS:DISK_UTILIZATION_HIGH_CLEAR
15-Feb-2016 07:26:54.586	15-Feb-2016 07:26:54.000	192.168.220.110	local3	INFO	S-HOST10	logger AFWlab-128-17_1 Start of action: Testing event DISK_UTILIZAT
15-Feb-2016 07:26:54.115	15-Feb-2016 07:26:54.000	192.168.220.110	local3	INFO	S-HOST10	logger AFWlab-128-17_1 End of action
15-Feb-2016 07:26:54.069	15-Feb-2016 07:26:54.000	192.168.220.110	local3	WARN	S-HOST10	event SYSTEM:SYSTEM:CHASSIS:DISK_UTILIZATION_HIGH ... Disk
15-Feb-2016 07:26:53.953	15-Feb-2016 07:26:53.000	192.168.220.110	local3	INFO	S-HOST10	logger AFWlab-128-17_1 Start of action: Testing event DISK_UTILIZAT

Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.



Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.




Open the Alert Views - RTView Alerts Table display.



Source: Select the host for which you want to view data, or **All Hosts**.


Tag: Select the message tag for which you want to view data, or **All Tags**.

Severity: Select the message severity level for which you want to view data, or **All Levels**.

Time Settings:

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

Timestamp

The date and time the row of data was last updated.

Message Timestamp

The date and time the message was sent.

Host Address

The host IP address. Refer to Solace documentation for more information.

Facility

The message facility code. Refer to Solace documentation for more information.

Severity

The message severity level. Refer to Solace documentation for more information.

- **INFO**
- **NOTICE**
- **NOTICE or higher**
- **WARN**
- **WARN or higher**
- **ERROR**
- **ERROR or higher**
- **CRITICAL**
- **ALERT**
- **EMERGENCY**

Tag

The host name. Refer to Solace documentation for more information.

Message Text

The content of the message.

CHAPTER 8 RTView DataServer for TIBCO

The RTView DataServer for TIBCO provides a way to create connections and modify default configuration settings for the various TIBCO solution packages and sends collected data to RTView Central, which contains the displays associated with the RTView DataServer for TIBCO that help you to monitor the health and performance across your TIBCO components.

RTView Central contains the following Views and their associated displays that will be populated with data collected via the RTView DataServer for TIBCO:

- ["TIBCO ActiveMatrix"](#)
- ["TIBCO ActiveSpaces \(2.x\)"](#)
- ["TIBCO Adapters"](#)
- ["TIBCO BusinessEvents"](#)
- ["TIBCO BusinessWorks"](#)
- ["TIBCO Enterprise Message Service"](#)
- ["TIBCO FTL"](#)
- ["TIBCO Hawk"](#)
- RTView Manager is also included. For details, see ["RTView Manager"](#).

The RTView *DataCollector* for TIBCO is available for use with the RTView DataServer for TIBCO. RTView DataCollector for TIBCO is used for collecting data and sending it to one or more RTView DataServers. The RTView DataCollector for TIBCO is also useful if you need to distribute data collection.

Note: This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

TIBCO ActiveMatrix

The TIBCO ActiveMatrix Views can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO ActiveMatrix is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central.

TIBCO ActiveSpaces (2.x)

The following TIBCO ActiveSpaces Views can be found under **Components** tab > **Middleware** > **TIBCO ActiveSpaces**:

- **"Spaces View"**: The displays in this View allow you to view the current and historical metrics for all metaspaces and spaces in a heatmap, tabular, or summary format.
- **"Members View"**: The displays in this View allow you to view the current and historical metrics for all members in a particular metaspace, view data for members within a particular space, and view data for all spaces for a particular member.

Spaces View

These displays provide detailed data for all metaspaces and spaces in a heatmap, tabular, or summary format. Displays in this View are:

- **"All Metaspaces Table"**: A tabular view of your metaspaces and their associated metrics.
- **"Metaspace Summary"**: This display allows you to view metrics and trend data for a particular metaspace.
- **"All Spaces Table"**: A tabular view of all spaces contained within a particular metaspace.
- **"All Spaces Heatmap"**: A heatmap view of all spaces contained within a particular metaspace.
- **"Space Summary"**: This display allows you to view metrics and trend data for a particular space.
- **"All Queries Table"**: This display allows you to view queries by domain, metaspace, and space and view the performance metrics for the queries.
- **"Query Summary"**: This display allows you to view performance metrics for a particular query, as well as to view any related queries.

All Metaspaces Table

The table in this display provides a view of all of your metaspaces and their associated metric data including domain, metaspace, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected metaspace in the **"Metaspace Summary"** display

← **ActiveSpaces All Metaspaces - Table** 21-Apr-2017 11:48 Data OK + ?

Domain: Production ▼

Metaspace Count: 3

Domain	Metaspace	Alert Level	Alert Count	Spaces	Members	AS Version	Entries	Replicas
Production	retailV212	●	0	3	4	2.1.2.005	450	
Production	retailV213	●	0	3	5	2.1.3.005	469	
Production	retailV214	●	0	3	4	2.1.4.011	441	

◀ ▶

Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▼, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

- Domain** Select the domain for which you want to view data.
- Metaspace Count** The total number of metaspaces found for the domain selected in the **Domain** dropdown, which are displayed in the **Metaspaces** table.

Metaspaces Table

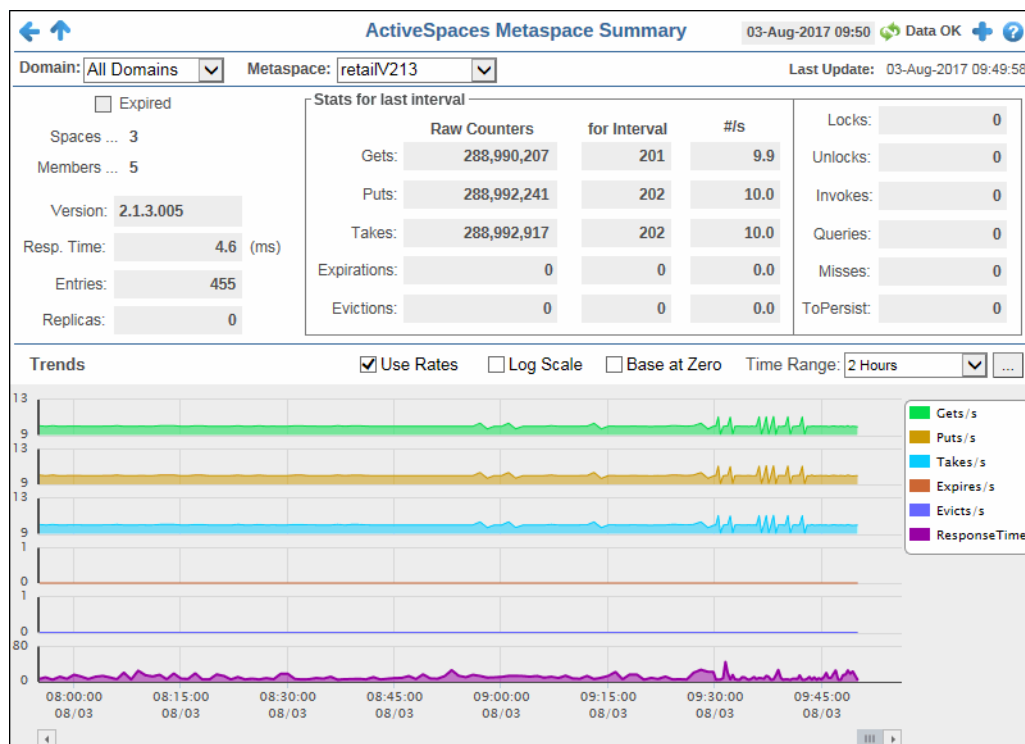
- Domain** The name of the domain.


Metaspace	The name of the metaspace.
Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the host.
Spaces	The number of user spaces defined in the metaspace.*
Members	The number of members (clients and servers) associated with the metaspace.*
AS Version	The metaspace's current version of TIBCO ActiveSpaces.*
Entries	The total number of entries stored in the metaspace.*
Replicas	The total number of replicas stored in the metaspace.*
Response Time	The average response time for the metaspace.*
Gets	The total number of "get" operations performed on the user-spaces defined on the metaspace.*
Gets/interval	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Gets/sec	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
Puts	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
Puts/interval	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Puts/sec	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
Takes	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
Takes/interval	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Takes/sec	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
Expires	The total number of entries in the user-spaces defined on the metaspace that have expired.*
Expires/interval	The number of entries in the user-spaces defined for the metaspace that expired during the current polling interval.*
Expires/sec	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
Evicts	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
Evicts/interval	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
Evicts/sec	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*

Locks	The total number of locks in the user-spaces defined for the metaspace.*
Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Invokes	The remote invocation count.*
Queries	The browser queries count in the user-spaces defined for the metaspace.*
Misses	The total number of misses on the user-spaces defined for the metaspace.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

Metaspace Summary

This display provides a view of the current and historical metrics for a single metaspace. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions, and also traces the average response time.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.

Fields and Data:

- Last Update** The date and time in which the data in the display was last updated.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Spaces** The number of user spaces defined in the metaspace.*
- Members** The number of members (clients and servers) associated with the metaspace.*
- Version** The metaspace's current version of TIBCO ActiveSpaces.
- Resp. Time** The average response time for the metaspace.*
- Entries** The total number of entries stored in the metaspace.*
- Replicas** The total number of replicas stored in the metaspace.*

Stats for last interval

- Gets** **Raw Counters--** The total number of gets for the metaspace.
for interval-- The number of gets for the current interval.
#/s -- The number of gets per second.
- Puts** **Raw Counters--** The total number of puts for the metaspace.
for interval-- The number of puts for the current interval.
#/s -- The number of puts per second.

Takes	Raw Counters -- The total number of takes for the metaspace. for interval -- The number of takes for the current interval. #/s -- The number of takes per second.
Expirations	Raw Counters -- The total number of expirations for the metaspace. for interval -- The number of expirations for the current interval. #/s -- The number of expirations per second.
Evictions	Raw Counters -- The total number of evictions for the metaspace. for interval -- The number of evictions for the current interval. #/s -- The number of evictions per second.
Locks	The total number of locks in the user-spaces defined for the metaspace.*
Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Invokes	The remote invocation count.*
Queries	The browser queries count in the user-spaces defined for the metaspace.*
Misses	The total number of misses in the user-spaces defined for the metaspace.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

Trends

Traces the following:

Gets(/s) -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

Puts(/s)-- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

Takes(/s) -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

Expires(/s) -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

Evicts(/s) -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

Response Time -- traces the average response time.


Use Rates Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).


Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Base at Zero



Select to use zero (**0**) as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

All Spaces Table

The table in this display provides a view of all of your spaces and their associated metric data including domain, metaspace, space, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the [“Space Summary”](#) display.

ActiveSpaces All Spaces - Table 21-Apr-2017 15:55 Data OK

Domain: Production Metaspace: retailV214

Space Count: 3

Domain	Metaspace	Space	Alert Level	Alert Count	Space State	Members	Seeders
Production	retailV214	customers		0	READY	4	
Production	retailV214	inventory		0	READY	4	
Production	retailV214	stores		0	READY	4	

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.


Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

- Domain** Select the domain for which you want to view data.
- Metaspace** Select the metaspace for which you want to view data.

Spaces Table:

- Domain** The name of the domain.


Metaspace	The name of the metaspace.
Space	The name of the space.
Alert Level	<p>The current alert severity.</p> <p> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</p> <p> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</p> <p> Green indicates that no metrics have exceeded their alert thresholds.</p>
Alert Count	The total number of alerts for the host.
Space State	The current state of the space.*
Members	The total number of members in the space.*
Seeders	The number of seeders in the space.*
Min Seeder Count	The defined minimum seeder count (minimum number of seeders that need to be joined to the space before the space becomes ready).*
CapacityPerSeeder	The capacity value for the space in number of entries per seeder.*
Entries	The total number of entries stored in the space.*
Replicas	The total number of replicas stored in the space.*
Gets	The total number of "get" operations performed on the user-spaces defined on the space.*
Gets/interval	The number of "get" operations performed on the user-spaces defined for the space during the current polling interval.*
Gets/sec	The rate of "get" operations (per second) performed on the user-spaces defined for the space.*
Puts	The total number of "put" operations performed on the user-spaces defined on the space.*
Puts/interval	The number of "put" operations performed on the user-spaces defined for the space during the current polling interval.*
Puts/sec	The rate of "put" operations (per second) performed on the user-spaces defined for the space.*
Takes	The total number of "take" operations performed on the user-spaces defined on the space.*
Takes/interval	The number of "take" operations performed on the user-spaces defined for the space during the current polling interval.*
Takes/sec	The rate of "take" operations (per second) performed on the user-spaces defined for the space.*
Expires	The total number of entries in the user-spaces defined on the space that have expired.*
Expires/interval	The number of entries in the user-spaces defined for the space that expired during the current polling interval.*
Expires/sec	The rate of entries in the user-spaces defined for the space that expired (per second).*
Evicts	The total number of entries in the user-spaces defined on the space that have been evicted.*

Evicts/interval	The number of entries performed on the user-spaces defined for the space that were evicted during the current polling interval.*
Evicts/sec	The rate of entries in the user-spaces defined for the space that were evicted (per second).*
Locks	The total number of locks in the user-spaces defined for the space.*
Unlocks	The total number of unlocks in the user-spaces defined for the space.*
Invokes	The remote invocation count.*
Queries	The total number of queries in the user-spaces defined for the space.*
Misses	The total number of misses in the user-spaces defined for the space.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

All Spaces Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your spaces for each available metric. You can view the spaces in the heatmap based on the following metrics: current alert severity, entries, gets per second, puts per second, takes per second, expires per second, and evicts per second. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Labels** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an space. Clicking one of the rectangles in the heatmap opens the "[Space Summary](#)" display, which allows you to see additional details for the selected space.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

Domain Select the domain for which you want to see data.

Metaspace Select the metaspace for which you want to see data.

Fields and Data:

Last Update The date and time in which the data in the display was last updated.

Space Count The total number of spaces found for the selected Domain/Metaspace combination.

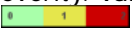
Labels Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.


Log Color Scale Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.


Log Size Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the size for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.


Color Metric Choose a metric to view in the display.

Alert Severity


The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

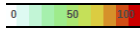
 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

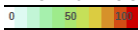
Entries

The total number of entries in the space. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceEntriesHigh**. The middle value in the gradient bar indicates the middle value of the range.

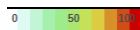
Gets/sec

The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceGetRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

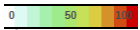
Puts/sec

The number of message sent per second. The color gradient bar  shows the range of the value/color mapping. The current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpacePutRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

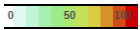
Takes/sec

The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceTakeRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Expires/sec


The number of expires per second. The color gradient bar  shows the range of the value/color mapping. The current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceExpireRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Evicts/sec

The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceEvictsRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Space Summary

This display provides a view of the current and historical metrics for a single space. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions.



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 connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.



Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.
- Space** Select the space for which you want to show data in the display.

Fields and Data:

Last Update	The date and time in which the data in the display was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Members	The total number of members associated with the space.* Note: You can click on this field to open the "Members by Space Table" .
Seeders	The number of seeders in the space.* Note: You can click on this field to open the "Members by Space Table" .
State	The current state of the space.*
Entries	The total number of entries stored in the space.*
Replicas	The total number of replicated entries in the space.*

Stats for last interval

Gets	Raw Counters-- The total number of gets for the space. for interval-- The number of gets for the current interval. #/s -- The number of gets received per second.
Puts	Raw Counters-- The total number of puts for the space. for interval-- The number of puts for the current interval. #/s -- The number of puts received per second.
Takes	Raw Counters-- The total number of takes for the space. for interval-- The number of takes for the current interval. #/s -- The number of takes received per second.
Expirations	Raw Counters-- The total number of expirations for the space. for interval-- The number of expirations for the current interval. #/s -- The number of expirations received per second.
Evictions	Raw Counters-- The total number of evictions for the space. for interval-- The number of evictions for the current interval. #/s -- The number of evictions received per second.
Locks	The total number of locks in the user-spaces defined for the space.*
Unlocks	The total number of unlocks in the user-spaces defined for the space.*
Invokes	The remote invocation count.*
Queries	The total number of queries in the user-spaces defined for the space.*
Misses	The total number of misses on the user-spaces defined for the space.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

Trends

Traces the following:

Gets(/s) -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

Puts(/s)-- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

Takes(/s) -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

Expires(/s) -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

Evicts(/s) -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

Response Time -- traces the average response time.

Use Rates

Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).


Log Scale

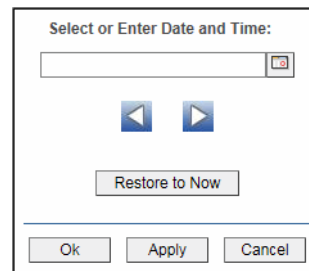
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Base at Zero


Select to use zero (**0**) as the Y axis minimum for all graph traces.



Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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 Queries Table

This display allows you to view queries by domain, metaspace, and space and view the performance metrics for the queries. Clicking on a query in the table opens the “[Query Summary](#)” display.

ActiveSpaces All Queries - Table 16-Mar-2018 13:14 Data OK

Domain: All Domains Metaspace: All Metaspaces Space: All Spaces

Query Count: 4,254 ☐ Show Expired

Domain	Timestamp	Metaspace	space_name	QueryDuration	query_status	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	52	2 key/10 = 10	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	20	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	20	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	20	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	19	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	18	2 key/10 = 10	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	17	2 key/10 = 10	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	16	2 key/10 = 10	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2 value like "valu	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2 year(time) = 20	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2 year(time) = 20	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2 (key%10) in (0,	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2 value not like "\	
tas_domain	16-Mar-2018 13:14:21	ms1	\$spacedef_space	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_state	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$members	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_members	0	2 space_name="	
tas_domain	16-Mar-2018 13:14:21	ms1	\$members	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$members	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$spacedef_space	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_members	0	2 space_name="	
tas_domain	16-Mar-2018 13:14:21	ms1	\$spacedef_space	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_state	0	2	

Page 1 of 5 1 - 200 of 1000 items

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

Domain Select the domain for which you want to show data in the display.

Metaspace Select the metaspace for which you want to show data in the display.

Space Select the space for which you want to show data in the display.

Fields and Data:

Query Count The total number of queries listed in the table.

Show Expired Select this toggle to display expired queries in the table.

Queries Table

Domain The name of the domain containing the query.

Timestamp The date and time that the row in the table was last updated.

Metaspace The name of the metaspace containing the query.

space_name The name of the space containing the query.

Query Duration The duration, in seconds, of the query.*

query_status The status of the query.*

0 - Failed

1 - In progress

2 - Completed

filter The filter used in the query.*

query_type The type of query.*

scan_type Lists whether the query used a table scan or an index scan.*

index_name The name of the index being used in the query.*

limit Lists the maximum number of entries that can be returned when executing a query.*

estimated_cost The estimated execution time of the query.*

actual_cost The actual execution time of the query.*

abort When checked, denotes that the query was aborted.*

StartTime Start time of the query.

EndTime End time of the query.

start_time Internal start time of the query.*

end_time Internal end time of the query.*

request_id The request id of the query.*

parent_request_id The request id of the query's parent.*

member_name The name of the member node.*

member_id The id of the member node.*

process_id The process ID of the member node processing the query.*

Expired When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.


Query Summary

This display allows you to view performance metrics for a particular query, as well as to view any related queries. Data only appears in this display when you select a query from the “[All Queries Table](#)”.

The screenshot shows the 'ActiveSpaces Query Summary' window. At the top, there are dropdown menus for 'Domain' (tas_domain), 'Metaspace' (ms1), and 'Space' (test_space_1). Below these is a table for the 'Selected Query' with columns: Domain, Timestamp, Metaspace, space_name, QueryDuration, query_status, and Filter. The data row shows: tas_domain, 16-Mar-2018 13:16:25, ms1, test_space_1, 52, 2, and key/10 = 10.

Below the table are three sections: 'Space Info' (Domain: tas_domain, Metaspace: ms1, Space: test_space_1), 'Member Info ...' (Name: ms1_as-agent-1, Id: c009c8d1-1b58-59d27f48-124, Process Id: 17664), and 'Query Info' (Duration: 52, Status: 2, Start Time: Oct 4, 2017 5:33:58 PM, End Time: Oct 4, 2017 5:34:50 PM, Limit: 10000, Request Id: c009c8d1-1b58-59d27f48-124-37a, Query Type: Query, Scan Type: TableScan, Index name: KeyIndex, Estimated Cost: 10000000, Actual Cost: 10000000, Filter: key/10 = 10, and an 'Abort' checkbox).

At the bottom, there is a section for 'Related Query Count: 1' and 'Parent Request Id: c009c8d1-1b58-59d57df5-132-48'. Below this is another table for 'Related Queries' with the same columns as the 'Selected Query' table, showing one related query with the same data as the selected query.



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.
- Space** Select the space for which you want to show data in the display.

Fields and Data:

- Selected Query** Lists the details of the query selected from the ["All Queries Table"](#).

Space Info

- Domain** The name of the domain in which the query resides.
- Metaspace** The name of the metaspace in which the query resides.
- Space** The name of the space in which the query resides.

Member Info **Note:** You can click this region to open the ["Member Summary"](#) display.

- Name** The name of the member node.*
- Id** The id of the member node.*
- Process ID** The process ID of the member node processing the query.*

Query Info

- Duration** The duration, in seconds, of the query.*
- Status** The status of the query.*
 - 0** - Failed
 - 1** - In progress
 - 2** - Completed
- Start Time** Start time of the query.

End Time	End time of the query.
Limit	Lists the maximum number of entries that can be returned when executing a query.*
Abort	When checked, denotes that the query was aborted.*
Request Id	The request id of the query.*
Query Type	The type of query.*
Scan Type	Lists whether the query used a table scan or an index scan.*
Index Name	The name of the index being used in the query.*
Estimated Cost	The estimated execution time of the query.*
Actual Cost	The actual execution time of the query.*
Filter	The filter used in the query.*
Related Query Count	The number of queries related to the selected query.
Parent Request Id	The request ID of the query's parent.
Related Queries	Lists the details of any related ("sibling") queries.

Members View

The displays in this view allow you to view the current and historical metrics for all members in a particular metaspace, view data for members within a particular space, and view data for all spaces for a particular member. The available displays in this View are:

- ["All Members Table"](#): A tabular view of all members in a particular metaspace.
- ["All Members Heatmap"](#): A heatmap view of all members in a particular metaspace.
- ["Member Summary"](#): This display allows you to view current and trending data for a single member for a particular metaspace.
- ["Member Summary - Process"](#): This display allows you to view current and trending process statistics for a single member for a particular metaspace.
- ["Member Summary -JVM"](#): This display allows you to view current and trending JVM statistics for a single member for a particular metaspace.
- ["Members by Space Table"](#): A tabular view of all members in a particular space.
- ["Members by Space Heatmap"](#): A heatmap view of all members in a particular space.
- ["Spaces by Member Table"](#): A tabular view of all spaces for a particular member.
- ["Member by Space Summary"](#): This display allows you to view data for a selected member for a particular space.

Filter By:**Domain**

Select the domain for which you want to view data.

Metaspace

Select the metaspace for which you want to view data.

Member Count

The resulting total number of members found in the filtered query, and listed in the **Members** table.

Members Table**Domain**

The name of the domain.

Metaspace


The name of the metaspace.


Member Name


The name of the member.

Alert Level

The current alert severity.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

Alert Count

The total number of alerts for the host.

Management Role

The member's role within the metaspace.

Host Address

The IP address of the host.*

Host Port

The port of the host.*

ProcessID

The process ID of the process being monitored.*

Process Name

The name of the process.*

NumSpaces

The number of spaces in which the metaspace member is a member.*

Entries

The number of entries associated with the member.*

Replicas

The number of replicas.*

Gets

The total number of "get" operations performed on the user-spaces defined on the metaspace.*

Gets/interval

The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*

Gets/sec

The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*

Puts

The total number of "put" operations performed on the user-spaces defined on the metaspace.*

Puts/interval

The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*

Puts/sec

The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*

Takes

The total number of "take" operations performed on the user-spaces defined on the metaspace.*

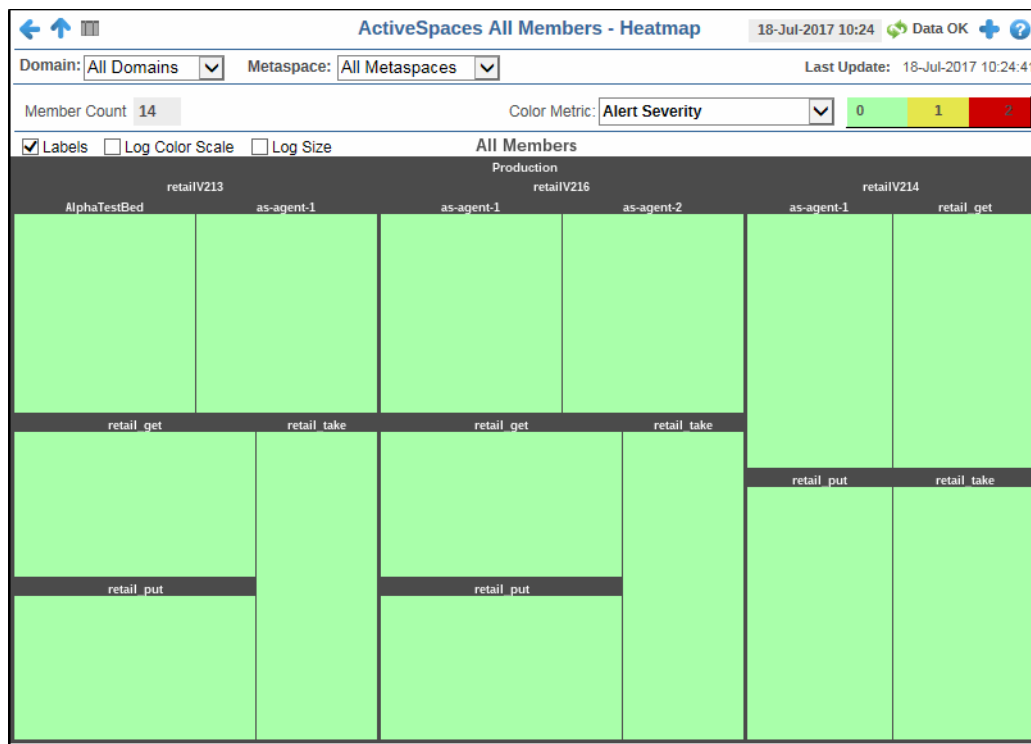
Takes/interval	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Takes/sec	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
Expires	The total number of entries in the user-spaces defined on the metaspace that have expired.*
Expires/interval	The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*
Expires/sec	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
Evicts	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
Evicts/interval	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
Evicts/sec	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
Locks	The total number of locks in the user-spaces defined for the metaspace.*
Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Invokes	The remote invocation count.*
Queries	The total number of queries in the user-spaces defined for the metaspace.*
Misses	The total number of misses in the user-spaces defined for the metaspace.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Join Time	The time that the member joined the space.
Timestamp	The date and time the row data was last updated.
sys_name	The operating system on which the member is running.*
cmd_name	Indicates the command used to start as-admin .*
user_name	The name of the user running the process.*
thread_count	The number of threads running for the process.*
res_mem_size	Indicates the amount of physical memory currently allocated to the member.*
mem_load	The percentage of memory being used.*
peak_res_mem_size	Indicates the peak size of the system resident memory allocated by the system.*
page_size	Indicates the current size of the system pagefiles allocated by the system.*

peak_page_size	Indicates the maximum size of the system pagefiles allowed by the system.*
process_cpu_load	Indicates the load on the CPU (CPU percentage).*
cpu_count	The number of CPUs running on the system.*
jvm_comm_heap_size	The committed JVM heap usage, in megabytes.*
jvm_max_heap_size	The maximum JVM heap usage, in megabytes.*
jvm_used_heap_size	The used JVM heap, in megabytes.*
jvm_comm_nonheap_size	The committed JVM non-heap memory usage, in megabytes.*
jvm_max_nonheap_size	The maximum JVM non-heap memory usage, in megabytes.*
jvm_used_nonheap_size	The used JVM non-heap memory, in megabytes.*
jvm_finalizing_count	The amount of memory freed by the finalize operation on the JVM.*
as_version	The current ActiveSpaces version running.*
JVMMemoryUsedPercent	The percentage of memory used by the JVM.*

All Members Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your members for each available metric. You can view the members in the heatmap based on the following metrics: the current alert severity, the number of entries, the number of gets per second, the number of puts per second, the number of takes per second, the number of expires per second, the number of evictions per second, the percentage of CPU used, the percentage of memory used, and the percentage of JVM memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Labels** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a particular member. Clicking one of the rectangles in the heatmap opens the ["Member Summary"](#) display, which allows you to see additional details for the selected member.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

- Domain** Select the Domain for which you want to view data.
- Metaspace** Select the metaspace for which you want to view data.

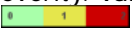
Fields and Data:


- Last Update** The date and time in which the data in the display was last updated.
- Member Count** The number of members found for the selected **Domain/Metaspace** combination.
- Labels** Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.
- Log Color Scale** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.


Log Size Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the size for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.


Color Metric Choose a metric to view in the display.

Alert Severity


The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

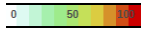
 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

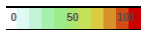
Entries

The total number of entries in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEntriesHigh**. The middle value in the gradient bar indicates the middle value of the range.

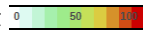
Gets/sec

The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberGetRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

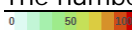
Puts/sec

The number of puts per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberPutRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

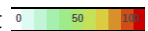
Takes/sec

The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberTakeRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

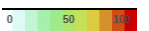
Expires/sec

The number of expires per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberExpireRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

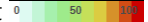
Evicts/sec

The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEvictsRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

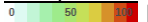
CPU %

The percentage of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberCpuHigh**. The middle value in the gradient bar indicates the middle value of the range.

Memory %


The percentage of memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberMemoryUsedHigh**. The middle value in the gradient bar indicates the middle value of the range.

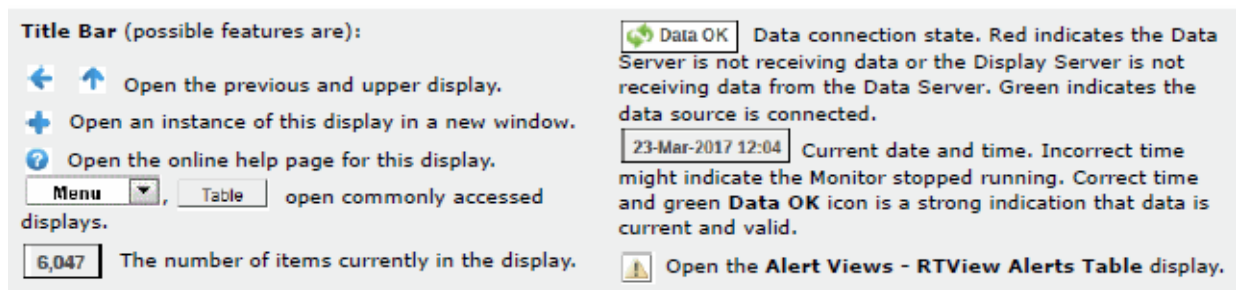
JVM Memory %

The percentage of JVM memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberJvmMemoryUsedHigh**. The middle value in the gradient bar indicates the middle value of the range.

Member Summary

This display provides a view of the current and historical metrics for a single member. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions.





Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.
- Member** Select the space for which you want to show data in the display.

Fields and Data:

- Last Update** The date and time in which the data in the display was last updated.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Spaces** The total number of spaces in which the member is a member.*
Note: Clicking on this field opens the "[Spaces by Member Table](#)" display.
- Name** The name of the member.
- Host** The IP address of the host.
- Process** The process ID and process name (ProcessID:ProcessName).*
- Role** The role of the member.
- Entries** The number of entries for the member.*
- Replicas** The number of replicas for the member.*
- Stats for last interval**

Gets	<p>Raw Counters-- The total number of "get" operations performed on the user-spaces defined on the metaspace.*</p> <p>for Interval-- The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*</p> <p>#/s -- The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*</p>
Puts	<p>Raw Counters-- The total number of "put" operations performed on the user-spaces defined on the metaspace.*</p> <p>for Interval-- The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*</p> <p>#/s -- The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*</p>
Takes	<p>Raw Counters-- The total number of "take" operations performed on the user-spaces defined on the metaspace.*</p> <p>for Interval-- The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*</p> <p>#/s -- The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*</p>
Expirations	<p>Raw Counters-- The total number of entries in the user-spaces defined on the metaspace that have expired.*</p> <p>for Interval-- The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*</p> <p>#/s -- The rate of entries in the user-spaces defined for the metaspace that expired (per second).*</p>
Evictions	<p>Raw Counters-- The total number of entries in the user-spaces defined on the metaspace that have been evicted.*</p> <p>for Interval-- The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*</p> <p>#/s -- The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*</p>
Locks	The total number of locks in the user-spaces defined for the metaspace.*
Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Invokes	The remote invocation count.*
Queries	The total number of queries in the user-spaces defined for the metaspace.*
Misses	The total number of misses in the user-spaces defined for the metaspace.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

Trends

Traces the following:

Gets(/s) -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

Puts(/s) -- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

Takes(/s) -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

Expires(/s) -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

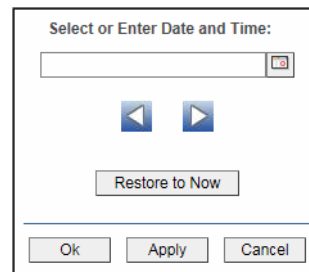
Evicts(/s) -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

Use Rates Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).


Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.



Base at Zero Select to use zero (**0**) as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

Member Summary - Process

This display provides a view of the current and historical process metrics for a single member. The trend graph in the bottom half of the display traces the current and historical process statistics for the selected member.

Note: These metrics are only available for members where system monitoring is enabled.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

Domain	Select the domain for which you want to show data in the display.
Metaspace	Select the metaspace for which you want to show data in the display.
Member	Select the space for which you want to show data in the display.

Fields and Data:

Last Update	The date and time in which the data in the display was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Spaces	The total number of spaces in which the member is a member. Note: Clicking on this field opens the "Spaces by Member Table" display.
Name	The name of the member.
Host	The IP address of the host.*
Process	The process ID and process name (ProcessID:ProcessName).*
Role	The role of the member.
Entries	The number of entries for the member.*
Replicas	The number of replicas for the member.*
Process Stats	
CPU %	Indicates the load on the CPU (CPU percentage).*
CPU Count	The number of CPUs running on the system.*
Thread Count	The number of threads running for the process.*
Memory %	The percentage of memory being used.*
Resident Memory	Peak-- Indicates the peak size of the system resident memory allocated by the system.* Size-- Indicates the amount of physical memory currently allocated to the member.*
Paged Memory	Peak-- Indicates the maximum size of the system pagefiles allowed by the system.* Size-- Indicates the current size of the system pagefiles allocated by the system.*
AS Version	The current ActiveSpaces version running.*
User	The name of the user running the process.*
System	The operating system on which the member is running.*

Command Indicates the command used to start the member process.*

Trends

Traces the following:

CPU %-- traces the percentage of CPU being used.

Memory %-- traces the percentage of memory being used.*

Peak Res Mem-- traces the peak size of the system resident memory allocated by the system.*

Res Mem-- traces the amount of physical memory currently allocated to the member.*

Peak Page Size-- traces the maximum size of the system pagefiles allowed by the system.*

Page Size-- traces the current size of the system pagefiles allocated by the system.*

Thread Count-- traces the number of threads running for the process.*


Log Scale


Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Base at Zero



Select to use zero (0) as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

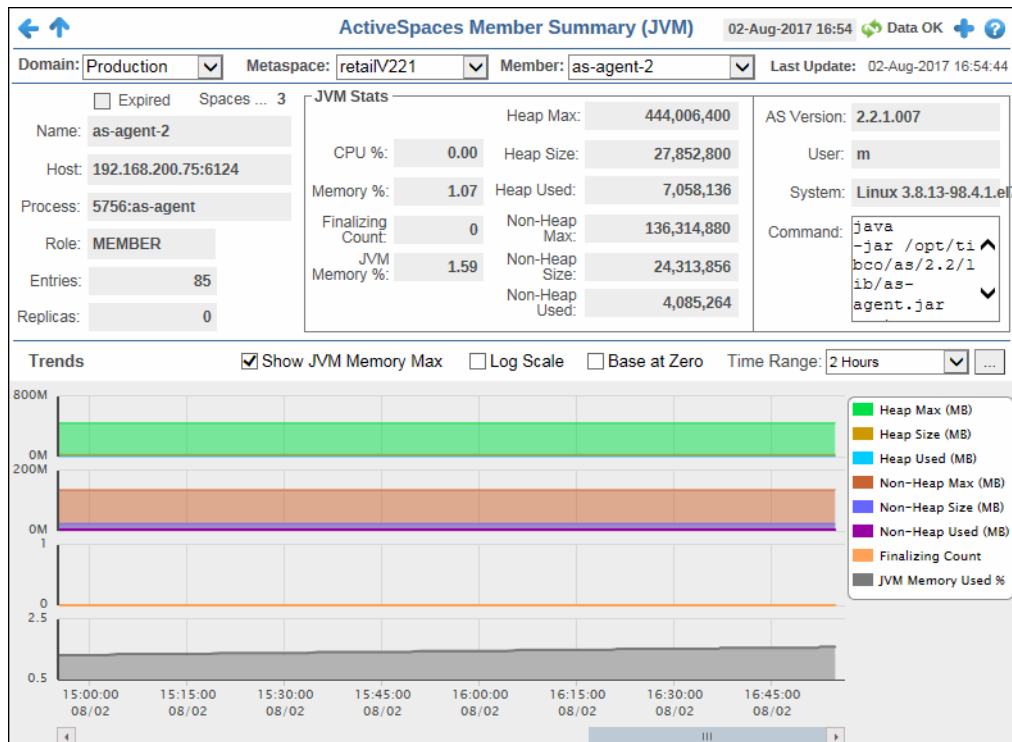
Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Member Summary -JVM

This display provides a view of the current and historical JVM statistics for a single member. The trend graph in the bottom half of the display traces the current and historical JVM metrics for the selected member.

Note: These metrics are only available for Java members where system monitoring is enabled.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

Domain	Select the domain for which you want to show data in the display.
Metaspace	Select the metaspace for which you want to show data in the display.
Member	Select the space for which you want to show data in the display.

Fields and Data:

Last Update	The date and time in which the data in the display was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Spaces	The total number of spaces in which the member is a member.* Note: Clicking on this field opens the " Spaces by Member Table " display.
Name	The name of the member.
Host	The IP address of the host.*
Process	The process ID and process name (ProcessID:ProcessName).*
Role	The role of the member.
Entries	The number of entries for the member.*
Replicas	The number of replicas for the member.*
JVM Stats	
CPU %	The load on the CPU (CPU percentage).*
Memory %	The percentage of memory being used.*
Finalizing Count	The amount of memory freed by the finalize operation on the JVM.*
JVM Memory %	The percentage of Java memory used by the JVM.*
Heap Max	The maximum JVM heap memory that can be used, in megabytes.*
Heap Size	The committed JVM heap size, in megabytes.*
Heap Used	The JVM heap memory currently being used, in megabytes.*
Non-Heap Max	The maximum JVM non-heap memory that can be used, in megabytes.*
Non-Heap Size	The committed JVM non-heap size, in megabytes.*
Non-Heap Used	The JVM non-heap memory currently being used, in megabytes.*
AS Version	The current ActiveSpaces version running.*

User	The name of the user running the process.*
System	The operating system on which the member is running.*
Command	Indicates the commands used to start the member process.*

Trends

Traces the following:

Heap Max (MB)-- traces the maximum JVM heap memory that can be used, in megabytes.*

Heap Size (MB)-- traces the maximum JVM heap usage, in megabytes.*

Heap Used (MB)-- traces the committed JVM heap size, in megabytes.*

Non-Heap Max (MB)-- traces the maximum JVM non-heap memory that can be used, in megabytes.*

Non-Heap Size (MB)-- traces the committed JVM non-heap size, in megabytes.*

Non-Heap Used (MB)-- traces the JVM non-heap memory currently being used, in megabytes.*

Finalizing Count-- traces the amount of memory freed by the finalize operation on the JVM.*

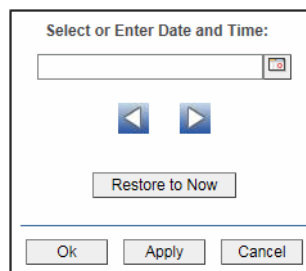
JVM Memory Used %-- traces the percentage of Java memory used by the JVM.*


Show JVM Memory Max When selected, enables the Heap Max (MB) and Non-Heap Max (MB) metrics in the trend graph, which might be useful for removing the maximum metrics from the plot when they differ significantly from the used and committed values.



Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



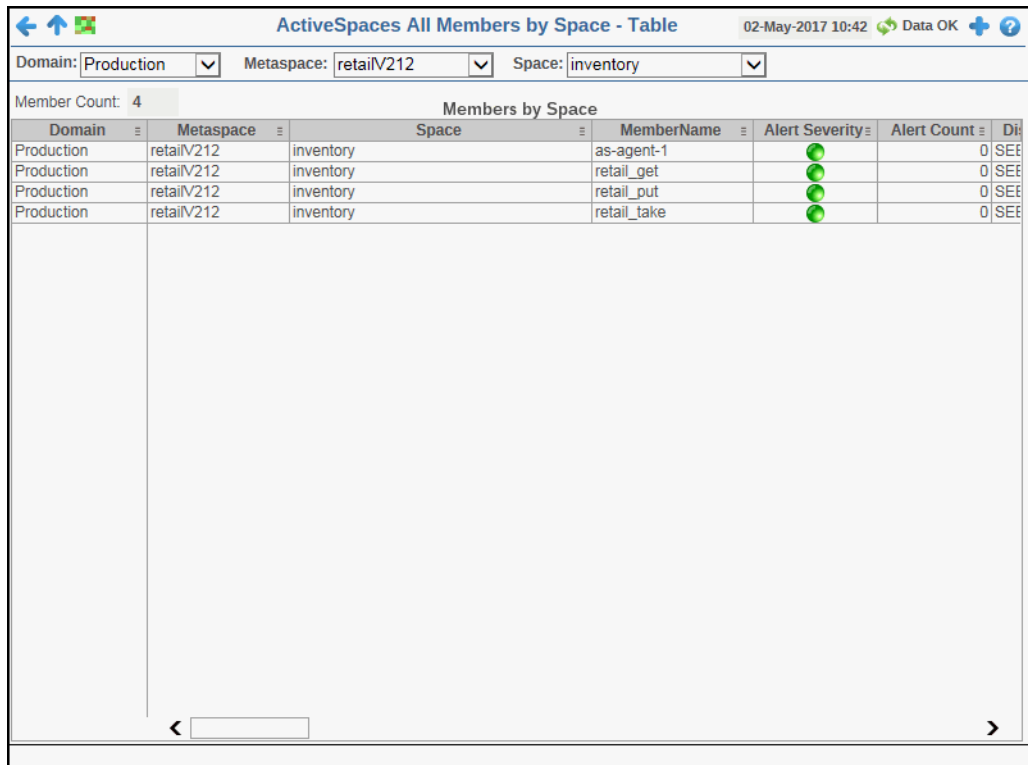
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Members by Space Table

The table in this display provides a view of all of your members and their associated metric data including domain, metaspace, space, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected member in the [“Member by Space Summary”](#) display.



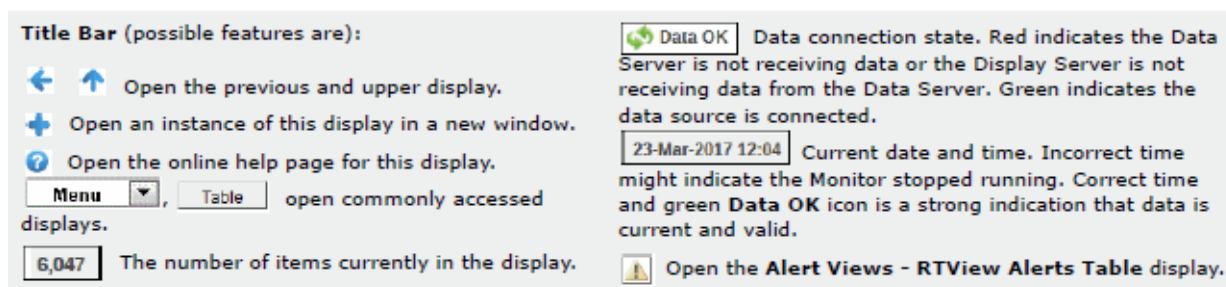
ActiveSpaces All Members by Space - Table 02-May-2017 10:42 Data OK

Domain: Production Metaspace: retailV212 Space: inventory

Member Count: 4

Members by Space

Domain	Metaspace	Space	MemberName	Alert Severity	Alert Count	Drill Down
Production	retailV212	inventory	as-agent-1		0	SEE
Production	retailV212	inventory	retail_get		0	SEE
Production	retailV212	inventory	retail_put		0	SEE
Production	retailV212	inventory	retail_take		0	SEE



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

Domain Select the domain for which you want to view data.

Metaspace Select the metaspace for which you want to view data.

Space Select the space for which you want to view data.

Member Count The resulting total number of members found in the filtered query, which are listed in the **Members by Space** table.

Members by Space Table

Domain The name of the domain.

Metaspace The name of the metaspace.

Space The name of the space.

Member Name The name of the member.

Alert Severity The current alert severity.

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Alert Count The total number of alerts for the host.

Distribution Role The member's role within the space.*

Entries The number of entries.*

% Capacity The percentage of available entries used for the space.

Replicas The number of replicas.*

Gets The total number of "get" operations performed on the user-spaces defined on the metaspace.*

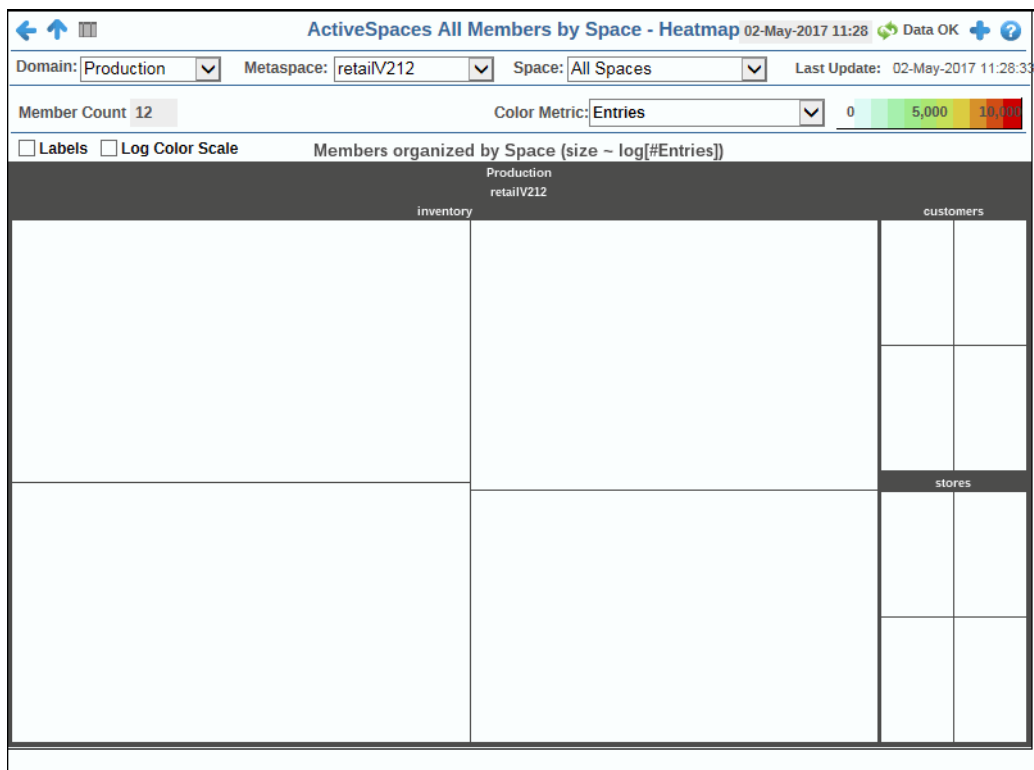
Gets/interval	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Gets/sec	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
Puts	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
Puts/interval	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Puts/sec	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
Takes	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
Takes/interval	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Takes/sec	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
Expires	The total number of entries in the user-spaces defined on the metaspace that have expired.*
Expires/interval	The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*
Expires/sec	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
Evicts	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
Evicts/interval	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
Evicts/sec	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
Locks	The total number of locks in the user-spaces defined for the metaspace.*
Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Invokes	The remote invocation count.*
Queries	The total number of queries in the user-spaces defined for the metaspace.*
Misses	The total number of misses in the user-spaces defined for the metaspace.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
ClientAvgGetMicros	The client's average "get" latency in microseconds.*
ClientAvgPutMicros	The client's average "put" latency in microseconds.*
ClientAvgTakeMicros	The client's average "take" latency in microseconds.*
ClientMaxGetMicros	The client's highest "get" latency in microseconds.*
ClientMaxPutMicros	The client's highest "put" latency in microseconds.*
ClientMaxTakeMicros	The client's highest "take" latency in microseconds.*

ClientTotalGetMillis	The client's cumulative total "get" latency in milliseconds for the current polling period.*
ClientTotalPutMillis	The client's cumulative total "put" latency in milliseconds for the current polling period.*
ClientTotalTakeMillis	The client's cumulative total "take" latency in milliseconds for the current polling period.*
ClientTotalMissMillis	The client's cumulative total "miss" latency in milliseconds for the current polling period.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

Members by Space Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your members for each available metric. You can view the members in the heatmap based on the following metrics: the number of entries, the number of gets per second, the number of puts per second, the number of takes per second, and the number of expires per second, and the number of evictions per second. By default, this display shows the heatmap based on the **Entries** metric.

You can use the **Labels** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a member. Clicking one of the rectangles in the heatmap opens the "[Member Summary](#)" display, which allows you to see additional details for the selected member.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

- Domain** Select the Domain for which you want to view data.
- Metaspace** Select the metaspace for which you want to view data.
- Space** Select the space for which you want to view data.

Fields and Data:

- Last Update** The date and time in which the data in the display was last updated.
- Member Count** The number of members found for the selected **Domain/Metaspace** combination.
- Labels** Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.


Log Color Scale

Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

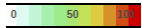
Color Metric

Choose a metric to view in the display.

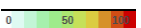
Entries

The total number of entries in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEntriesHigh**. The middle value in the gradient bar indicates the middle value of the range.

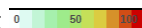
Gets/sec

The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberGetRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

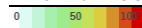
Puts/sec

The number of puts per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberPutRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

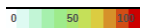
Takes/sec

The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberTakeRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Expires/sec

The number of expires per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberExpireRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Evicts/sec

The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEvictsRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Spaces by Member Table

The table in this display provides a view of all of your spaces (by member name) and the their associated metric data including domain, metaspace, space, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected member in the “[Member by Space Summary](#)” display.

ActiveSpaces All Spaces by Member - Table02-May-2017 13:25Data OK

Domain:All DomainsMetaspace:retailV212Member:as-agent-1

Member Count: 3

Domain	Metaspace	MemberName	Space	Alert Severity	Alert Count	Di
Production	retailV212	as-agent-1	customers		0	SEI
Production	retailV212	as-agent-1	inventory		0	SEI
Production	retailV212	as-agent-1	stores		0	SEI

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

Domain Select the domain for which you want to view data.

Metaspace Select the metaspace for which you want to view data.

Member Select the space for which you want to view data.

Member Count The resulting total number of members found in the filtered query, which are listed in the **Spaces by Members** table.

Spaces by Member Table




Domain The name of the domain.

Metaspace The name of the metaspace.

Member Name The name of the member.

Space The name of the space.

Alert Severity The current alert severity.

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

Alert Count The total number of alerts for the host.

Distribution Role The member's role within the space.

Entries The number of entries.*

% Capacity The percentage of available entries used for the space.

Replicas The number of replicas.*

Gets The total number of "get" operations performed on the user-spaces defined on the metaspace.*

Gets/interval The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*

Gets/sec The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*

Puts The total number of "put" operations performed on the user-spaces defined on the metaspace.*

Puts/interval The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*

Puts/sec The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*

Takes The total number of "take" operations performed on the user-spaces defined on the metaspace.*

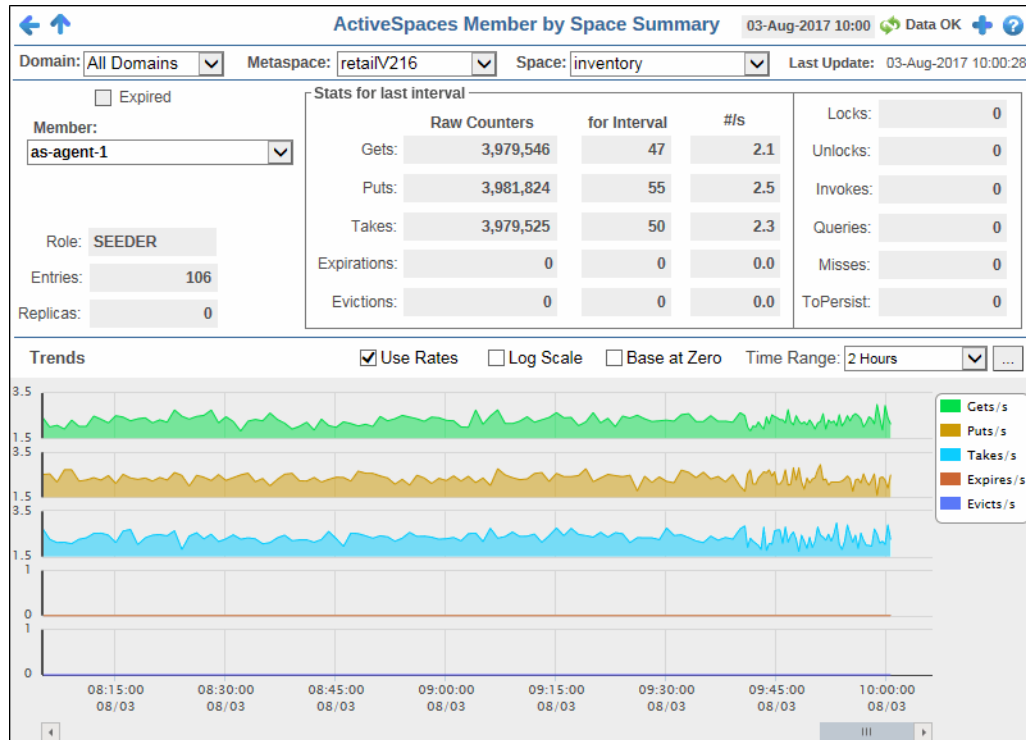
Takes/interval The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*

Takes/sec The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*

Expires	The total number of entries in the user-spaces defined on the metaspace that have expired.*
Expires/interval	The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*
Expires/sec	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
Evicts	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
Evicts/interval	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
Evicts/sec	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
Locks	The total number of locks in the user-spaces defined for the metaspace.*
Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Invokes	The remote invocation count.*
Queries	The total number of queries in the user-spaces defined for the metaspace.*
Misses	The total number of misses in the user-spaces defined for the metaspace.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
ClientAvgGetMicros	The client's average "get" latency in microseconds.*
ClientAvgPutMicros	The client's average "put" latency in microseconds.*
ClientAvgTakeMicros	The client's average "take" latency in microseconds.*
ClientMaxGetMicros	The client's highest "get" latency in microseconds.*
ClientMaxPutMicros	The client's highest "put" latency in microseconds.*
ClientMaxTakeMicros	The client's highest "take" latency in microseconds.*
ClientTotalGetMillis	The client's cumulative total "get" latency in milliseconds for the current polling period.*
ClientTotalPutMillis	The client's cumulative total "put" latency in milliseconds for the current polling period.*
ClientTotalTakeMillis	The client's cumulative total "take" latency in milliseconds for the current polling period.*
ClientTotalMissMillis	The client's cumulative total "miss" latency in milliseconds for the current polling period.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Active Spaces > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

Member by Space Summary

This display provides a view of the current and historical metrics for a single member in a particular space. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

Domain Select the domain for which you want to show data in the display.

Metaspace Select the metaspace for which you want to show data in the display.

Space Select the space for which you want to show data in the display.

Fields and Data:

Last Update The date and time in which the data in the display was last updated.

Expired When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

Member Select the member for which you want to see data.*

Role The member's role within the space.*

Entries The number of entries.*

Replicas The number of replicas.*

Stats for last interval

Gets	Raw Counters -- The total number of gets for the space. for interval -- The number of gets for the current interval. #/s -- The number of gets received per second.
Puts	Raw Counters -- The total number of puts for the space. for interval -- The number of puts for the current interval. #/s -- The number of puts received per second.
Takes	Raw Counters -- The total number of takes for the space. for interval -- The number of takes for the current interval. #/s -- The number of takes received per second.
Expirations	Raw Counters -- The total number of expirations for the space. for interval -- The number of expirations for the current interval. #/s -- The number of expirations received per second.
Evictions	Raw Counters -- The total number of evictions for the space. for interval -- The number of evictions for the current interval. #/s -- The number of evictions received per second.
Locks	The total number of locks in the user-spaces defined for the space.*
Unlocks	The total number of unlocks in the user-spaces defined for the space.*
Invokes	The remote invocation count.*
Queries	The total number of queries in the user-spaces defined for the space.*
Misses	The total number of misses in the user-spaces defined for the space.*
ToPersist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

Trends

Traces the following:

Gets(/s) -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

Puts(/s) -- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

Takes(/s) -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

Expires(/s) -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

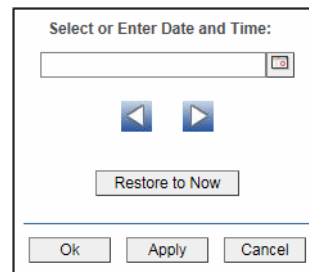
Evicts(/s) -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

Use Rates Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).


Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.



Base at Zero Select to use zero (**0**) as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

TIBCO Adapters

The following TIBCO Adapters Views can be found under **Components** tab > **Middleware** > **TIBCO Adapters**. The displays within the Views will be populated with data once the Solution Package for TIBCO Adapters is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central.

This section contains the following:

- **"All Adapters View"**: The displays in this View allow you to view the current and historical metrics for all adapters in a heatmap or tabular format.
- **"Single Adapter View"**: The displays in this View allow you to view the current and historical metrics for a single adapter in a tabular format.

All Adapters View


These displays provide detailed data for all adapters. Displays in this View are:

- **"All Adapters Heatmap"**: A heatmap view of all adapters in a heatmap format and their associated metrics.
- **"All Adapters Table"**: A tabular view of your adapters and their associated metrics.

All Adapters Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your adapters for each available metric. You can view the adapters in the heatmap based on the following metrics: the current alert severity, the current alert count, the delta messages received, the messages received rate, the messages sent rate, the delta messages sent, and the increase in errors from the previous polling period. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box ☒ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an adapter. Clicking one of the rectangles in the heatmap opens the **"Adapter Summary"** display, which allows you to see additional details for the selected adapter



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

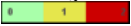
Fields and Data:




- Count** The number of adapters included in the display. This number can change if you toggle the **Expired** check box on and off.
- Show: Expired** Select this check box to display those adapters whose data has not been updated recently (expired).
- Show: Names** Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
Note: Some metrics auto-scale automatically, even when **Auto** is not selected.

Metric


Choose a metric to view in the display.

Alert Severity

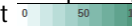
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

Alert Count

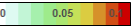
The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

Delta Msgs Rcvd

The increase in the number of messages received (per second) from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of messages received. The middle value in the gradient bar indicates the middle value of the range.


The **Auto** check box does not impact this metric.

Rate Msgs Rcvd

The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TadAdapterMsgsRcvdRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

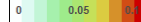
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Delta Msgs Sent

The increase in the number of messages sent (per second) from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of messages sent. The middle value in the gradient bar indicates the middle value of the range.

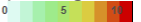
The **Auto** check box does not impact this metric.

Rate Msgs Sent

The number of message sent per second. The color gradient bar  shows the range of the value/color mapping. The current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TadAdapterMsgsSentRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Delta Errors

The increase in the number of errors from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TadAdapterDeltaErrorsHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

All Adapters Table

The table in this display provides a view of all of your adapters and their associated metric data including agent, application instance, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the ["Adapter Summary"](#) display

All Adapters - Table23-Aug-2017 10:56Data OK+Count: 4

Show: ☒ Expired

Agent	Application Instance	Alert Level	Alert Count	Messages Received	Delta Msgs Received	Rate
SLHOST15(sl_qa)	Transfer_Records_Via_BW.FAWriter			0	0	
SLHOST15(sl_qa)	Container_Read_Write.ContainerReader			0	0	
SLHOST15(sl_qa)	Transfer_Records_Via_BW.FAReader			0	0	
SLHOST15(sl_qa)	Container_Read_Write.ContainerWriter			0	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.

6,047

The number of items currently in the display.

Data OK

Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected adapter. Refer to TIBCO Adapter documentation for more information regarding these fields.

Fields and Data:

- Show: Expired

Select this check box to display adapters that have expired data in the table.
- Count




The total number of adapters listed in the All Adapters Table.

All Adapters Table:

- Agent

The name of the agent.
- Application Instance

The name of the application instance.

Alert Level	<p>The current alert severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the host.
Messages Received	The number messages received.*
Delta Messages Received	The increase in the number of messages received (from the previous polling period to the current polling period).*
Rate Messages Received	The number of messages received per second.*
Messages Sent	The total number of messages sent.*
Delta Msgs Sent	The increase in the number of messages sent (from the previous polling period to the current polling period).*
Rate Msgs Sent	The number of messages sent per second.*
New Errors	The number of new errors received since the last polling update.*
Total Errors	The total number of errors.*
Delta Total Errors	The increase in the number total errors (from the previous polling period to the current polling period).*
Rate Total Errors	The number of errors per second.*
Adapter Name	The name of the adapter.*
Last Restart	The date and time the adapter was last restarted.*
Process ID	The process ID of TIBCO Adapter you are running.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Adapters > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

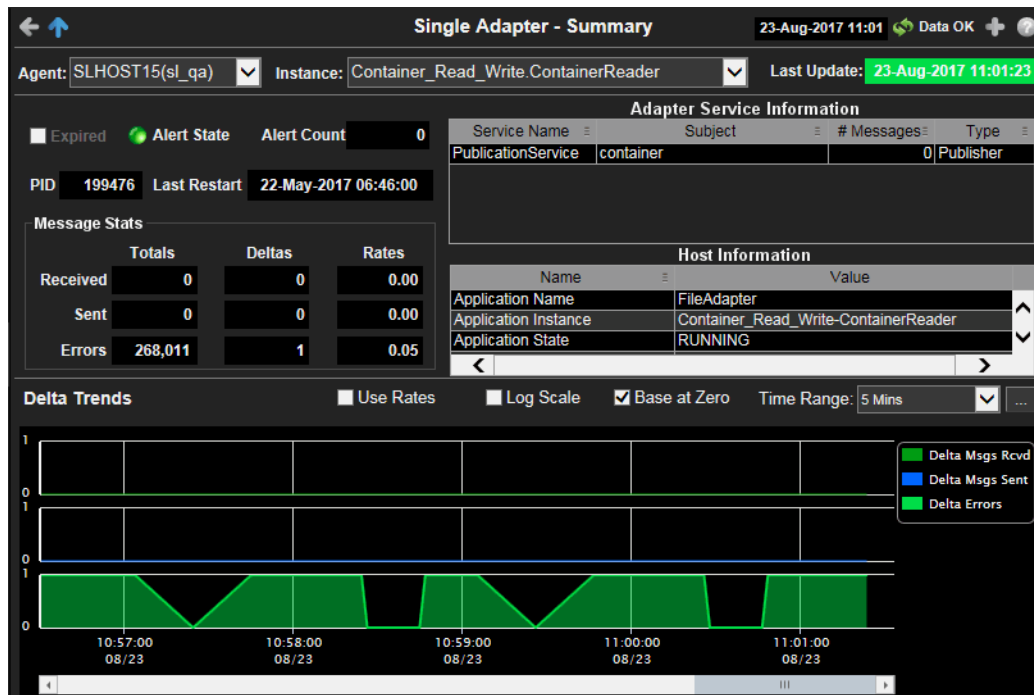
Single Adapter View

This display allows you to view the current and historical metrics for a single adapter. The available display in this View is:

- **"Adapter Summary"**: This display allows you to view current and trending data for a single adapter for a particular agent.

Adapter Summary

This display provides a view of the current and historical metrics for a single adapter. You can view message statistics, adapter service information, and host information for a specific instance. The trend graph in the bottom half of the display traces the current and historical delta messages received, delta messages sent, and delta errors.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by the selected adapter. Refer to TIBCO Adapter documentation for more information regarding these fields.

Filter By:

The display might include these filtering options:

Agent Select the agent for which you want to show data in the display.

Instance Select the instance for which you want to show data in the display.


Fields and Data:


Last Update The date and time in which the data in the display was last updated.

Expired When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **TIBCO Adapters** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

Alert State The current alert severity.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

Alert Count The total number of current alerts.

PID The process ID of the Instance.*

Last Restart The data and time the instance was last restarted.*

Message Stats

Received

Totals -- The total number of messages received.*

Deltas -- The increase in the number of messages received since the last polling update.*

Rates -- The number of messages received per second.*

Sent

Totals -- The total number of messages sent.*

Deltas -- The increase in the number of messages sent since the last polling update.*

Rates -- The number of messages sent per second.*

Errors

Totals -- The total number of errors that have occurred.*

Deltas -- The increase in the number of errors since the last polling update.*

Rates -- The number of errors occurring per second.*

Adapter Service Information

Service Name The name of the service.*

Subject The name of the subject.*

Messages The current number of messages.*

Type The type of adapter service.*

Host Information

Name The name of the host.*

Value The host's value.*

Trends Graph Traces the following:

Delta Msgs Rcvd -- traces the increase in the number of messages received since the last polling update, or the rate of messages received with **Use Rates** selected.

Delta Msgs Sent -- traces the increase in the number of messages sent since the last polling update, or the rate of messages sent with **Use Rates** selected.

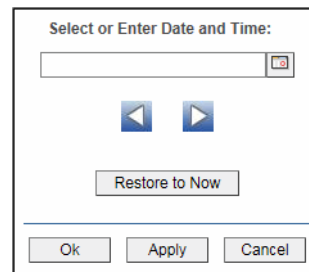
Delta Errors -- traces the increase in the number of errors since the last polling update, or the rate of errors with **Use Rates** selected.

Use Rates Select this toggle to trace the rates (**Msgs Rcvd/sec**, **Msgs Sent/sec**, and **Errors/sec**) instead of the delta numbers (**Delta Msgs Rcvd**, **Delta Msgs Sent**, and **Delta Errors**).


Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.



Base at Zero Select to use zero (**0**) as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field at the top with a calendar icon on its right. Below the input field are two blue navigation arrows, one pointing left and one pointing right. Underneath these arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

TIBCO BusinessEvents

The following TIBCO BusinessEvents Views can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO BusinessEvents is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central.

This section includes:

- ["Clusters / Nodes View"](#)
- ["Events / Concepts View"](#)


Clusters / Nodes View

These displays present performance data for your BusinessEvents system. Displays in this View are:

- ["Clusters"](#)
- ["Cluster Summary"](#)
- ["Cluster Nodes Table"](#)
- ["Cluster Nodes Heatmap"](#)
- ["Inference Node Summary"](#)
- ["Storage Node Summary"](#)

Clusters






Use this display to check event, concept, and backing store metrics for all of your clusters. Consider keeping this display open to monitor your TIBCO BusinessEvents clusters in general. Each row in the table is a different cluster. Click on a cluster row to view additional cluster details (current and historical) in the ["Cluster Summary"](#) display. The summary display includes trend charts so that you can view key metrics over time.

Sort  the table columns when all the rows cannot fit on the screen. For example, sort the **Alert Status** column so that all nodes with red alerts (●) are listed at the top, or sort the **Expired** column so that all expired nodes are listed at the top.



Cluster Name	Alert Severity	Alert Count	Member Count	Num Events Received	Num Events Sent	Events Received Per Sec	Num Asserted From Channel Per Sec	Num Retracted From Channel Per
ckfdcache		0	2	182,406	0	4.52	4.51	
fdcache		0	2	0	0	0.00	0.00	

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 connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.






Open the Alert Views - RTView Alerts Table display.

Clusters Table

Each row in the table is a different cluster, and data in the row columns describe the cluster.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

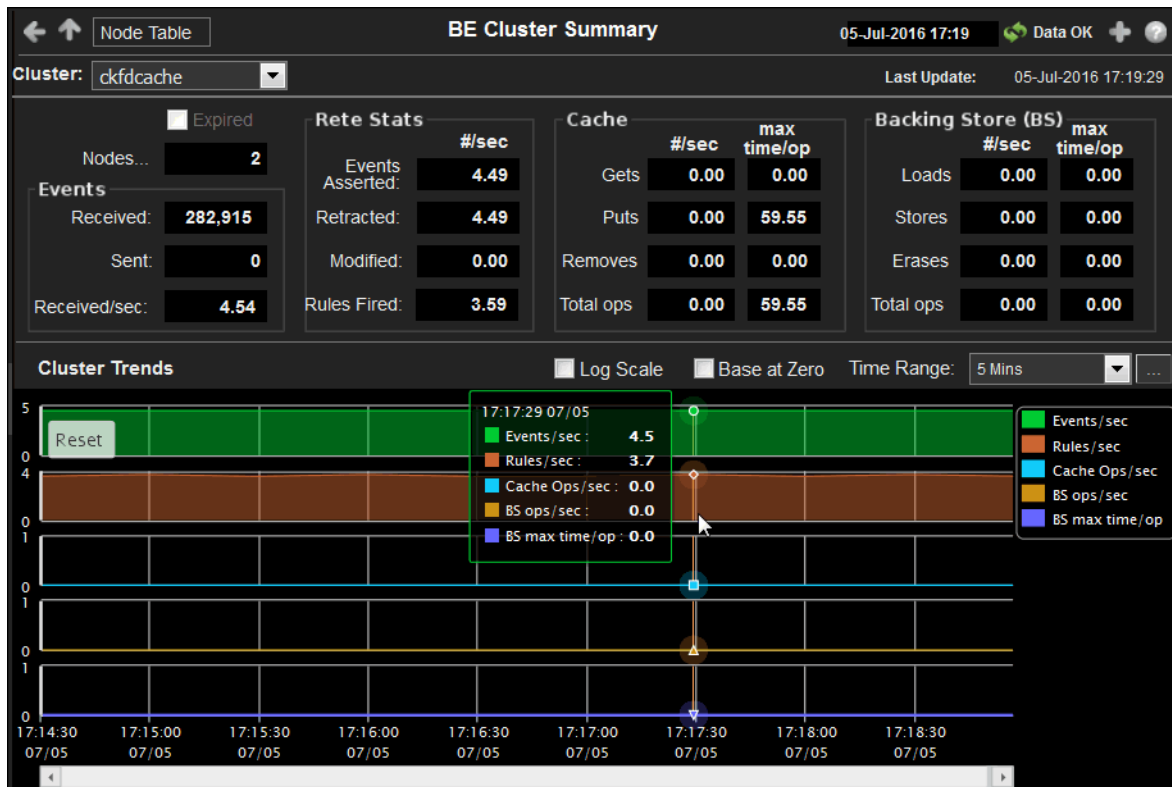
- Count:** The total number of clusters in the table.
- Cluster Name** The name of the TIBCO BusinessEvents cluster.
- Alert Severity** The severity level of open alerts. Values range from **0** to **2**, where **2** is the greatest Severity:
 -  One or more alerts exceeded their ALARM LEVEL threshold.
 -  One or more alerts exceeded their WARNING LEVEL threshold.
 -  No alert thresholds have been exceeded.

Alert Count	The total number of critical and warning alerts.
Member Count	<p>The count of the number of nodes (both cache and inference) that have been collected. For example, for a cluster that has 3 inference nodes and two cache nodes, the Member Count for all 5 rows in the Cluster Table should be 5. If one of the rows shows a member count of one and the others show four, that is a clear indication that a node failed to join the cluster, and the corresponding node should be restarted.</p> <p>Note: The actual number of nodes in the cluster will not match the count in this column if one or more of the nodes do not have connection properties configured in the property file that is read by the data server at startup.</p>
Num Events Received	The total number of events received.*
Num Events Sent	The total number of events sent.*
Events Received Per Sec	The rate of events received in the cluster.
Num Asserted From Channel Per Sec	The rate of events asserted into the Rete network via the channel.
Num Retracted From Channel Per Sec	The rate of events retracted/deleted from the Rete network via the channel.
Num Modified From Channel Per Sec	The rate of events modified in the Rete network via the channel.
Num Rules Fired Rate	The rate of rules fired in the cluster.
Concept Max Get Time	The longest time taken for a "get" operation for any node in the cluster since the cluster was started.*
Concept Max Put Time	The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*
Concept Max Remove Time	The longest time taken for a "remove" operation for any node in the cluster since the cluster was started.*
Concept Max Operation Time	The longest time taken for a concept operation (get/put/remove) for any node in the cluster since the cluster was started.*
Concept Gets/sec	The rate of "get" operations in the cluster.
Concept Puts/sec	The rate of "put" operations in the cluster.
Concept Removes/sec	The rate of "remove" operations in the cluster.
Concept Operations/sec	The rate of operations (gets/puts/removes) in the cluster.
Backing Store Max Erase Time	The longest time taken for an "erase" operation in the Backing Store for any node in the cluster.*
Backing Store Max Load Time	The longest time taken for a "load" operation in the Backing Store for any node in the cluster.*
Backing Store Max Store Time	The longest time taken for a "store" operation in the Backing Store for any node in the cluster.*

Backing Store Max Operation Time	The longest time taken to perform an operation (erase/load/store) in the Backing Store for any node in the cluster.*
Backing Store Erases/sec	The rate of "erases" in the Backing Store.
Backing Store Loads/sec	The rate of "loads" into the Backing Store.
Backing Store Stores/sec	The rate of "stores" into the Backing Store.
Backing Store Operations/sec	The rate of operations (erases/loads/stores) in the Backing Store.
Source	The name of the data server from which the data was collected.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the engine.

Cluster Summary

Use this display to view configuration and utilization data for a single cluster. Select a cluster to view Rete statistics, cache metrics, Backing Store data, and trend data for the cluster.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Cluster** Choose a cluster for which you want to see metrics.
- Last Update** The date and time the data was last updated in the display.

Fields and Data

This display includes:

Note: Fields with an asterisk (*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.


Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.	
Nodes	Lists the number of nodes on the cluster.	
Events	Received	The number of events received since the last data update.*
	Sent	The number of events sent since the last data update.*
	Received/sec	The rate of events received in the cluster.
Rete Stats	Events Asserted (#/sec)	The rate of events asserted into the Rete network.
	Retracted (#/sec)	The rate of events retracted/deleted from the Rete network.
	Modified (#/sec)	The rate of events modified in the Rete network.
	Rules Fired (#/sec)	The rate of rules fired in the Rete network.
Cache	Gets (#/sec)	The rate of "get" operations in the L1 cache.
	Gets (max time/op)	The longest time taken for a "get" operation for any node in the cluster since the cluster was started.*
	Puts (#/sec)	The rate of "put" operations in the cache.
	Puts (max time/op)	The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*
	Removes (#/sec)	The rate of "removes" in the cache.
	Removes (max time/op)	The longest time taken for a "remove" operation for any node in the cluster since the cluster was started.*
	Total ops (#/sec)	The rate of operations (gets/puts/removes) in the cache.
	Total ops (max time/op)	The longest time taken for an operation (get/put/remove) for any node in the cluster since the cluster was started.*
Backing Store (BS)	Loads (#/sec)	The rate of "load" operations into the backing store.
	Loads (max time/op)	The longest time taken for a "load" operation in the backing store for any node in the cluster.*
	Stores (#/sec)	The rate of "store" operations in the backing store.


Stores (max time/op)	The longest time taken for a “store” operation in the backing store for any node in the cluster.*
Erases (#/sec)	The rate of “erase” operations in the backing store.
Erases (max time/op)	The longest time taken for an “erase” operation in the backing store for any node in the cluster.*
Total ops (#/sec)	The rate of operations (loads/stores/erases) in the backing store.
Total ops (max time/op)	The longest time taken to perform an operation (erase/load/store) in the backing store for any node in the cluster.*
Cluster Trends	Shows the following metrics for the selected cluster. <ul style="list-style-type: none"> Events/sec -- Traces the rate of events received in the cluster. Rules/ sec -- Traces the rate of rules in the cluster. Cache Ops/ sec -- Traces the rate of cache operations in the cluster. BS ops/sec-- Traces the rate of backstore operations in the cluster. BS max time/op-- Traces the average maximum time per backstore operation.
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero

When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.





By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.


Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Cluster Nodes Table

Use this display to view configuration and utilization data for nodes in a cluster.

BE Cluster Nodes - Table										
Cluster: ckfdcache										
Node Count: 4										
Cluster Nodes										
Cluster Name	Node	Alert Severity	Alert Count	Member Count	Auto Startup	Backing Store Enabled	Cache Aside	Serialization Optimized	Storage Enabled	Cache Ty
ckfdcache	new51Cache		0	2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	DISTRIBUTE
ckfdcache	new51Inf		0	2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	DISTRIBUTE

**Filter By:**

The display might include these filtering options:

Cluster Choose a cluster for which you want to see metrics.

Cluster Nodes Table

Each row in the table is a different node. Data in the row columns describe the node.

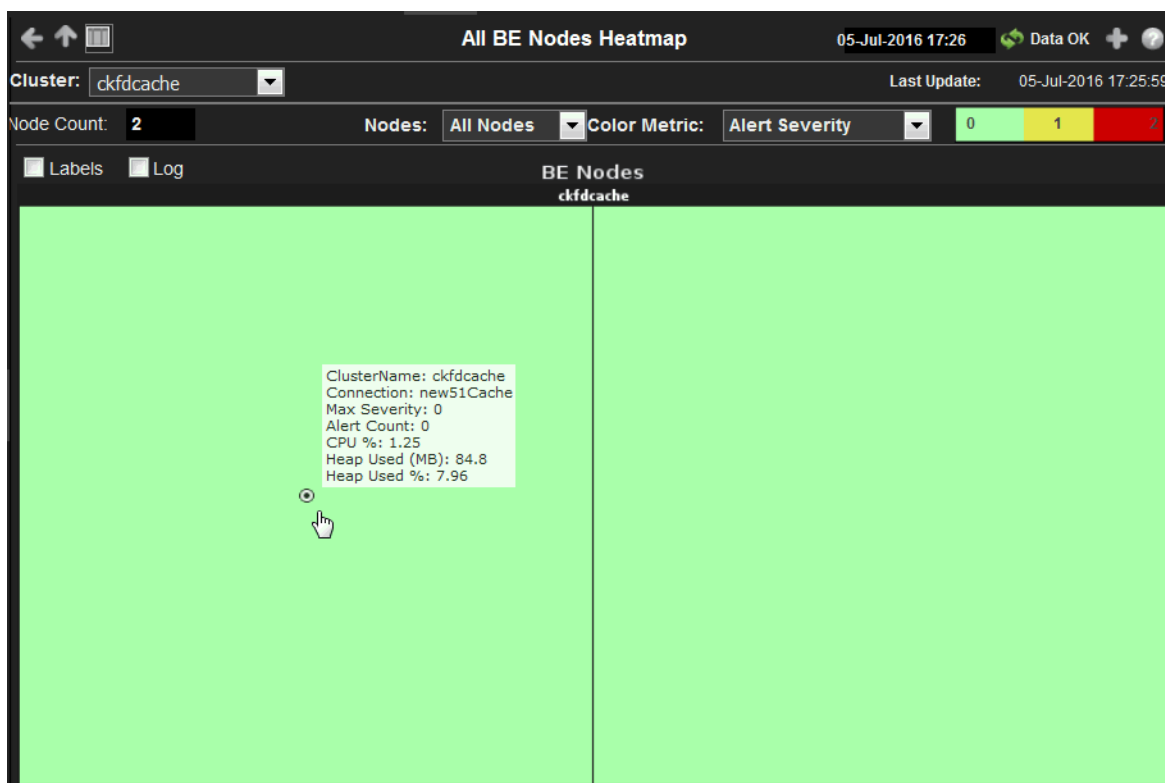
Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Node Count:	The total number of clusters in the table.
Cluster Name	The name of the TIBCO BusinessEvents cluster.
Node	The name of the node.
Alert Severity	<p>The severity level of open alerts. Values range from 0 to 2, where 2 is the greatest Severity:</p> <ul style="list-style-type: none"> ● One or more alerts exceeded their ALARM LEVEL threshold. ● One or more alerts exceeded their WARNING LEVEL threshold. ● No alert thresholds have been exceeded.
Alert Count	The total number of critical and warning alerts.
Member Count	The number of neighbors seen by a given node. This value is obtained directly from each node in the cluster. This value should always match the total "Member Count" in the corresponding row of the Clusters table. If they do not match, the node did not join the cluster properly and, hence, the cluster should be restarted.
Auto Startup	When checked (true), this feature is enabled.
Backing Store Enabled	When checked (true), this feature is enabled.*
Cache Aside	When checked (true), this feature is enabled.*
Serialization Optimized	When checked (true), this feature is enabled.*
Storage Enabled	When checked (true), this feature is enabled.*
Cache Type	The type of TIBCO BusinessEvents cache.*
BE Version	The approximate TIBCO BusinessEvents version, as configured by the connection property. The exact version information is not available via JMX.

Cache Node?	When checked (true), the node is a storage node. Otherwise, it is an inference node. This column is added by the Monitor rather than read from the JMX interface.
Node ID	A unique string that identifies the node.
Host	The IP address of the host to which the node is connected.
Port	The port number of the host to which the node is connected.
URL	Uniform Resource Locator, used as an alternative way to specify a JMX connection. When set, the Host and Port columns are blank (and vice versa).
% CPU Used	The amount of CPU, in percent, used by the node. This value is derived from the java.lang.OperatingSystem MBean.
Heap-Max	The maximum amount of memory, in megabytes, that can be used by the JVM for heap space. This value is provided by standard Java MBeans.
Heap-Used	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
NonHeap Max	The maximum amount of memory, in megabytes, that can be used by the process (not counting heap usage). This value is provided by standard Java MBeans.
NonHeap Used	The current amount of memory, in megabytes, in use by the process (not counting heap usage). This value is provided by standard Java MBeans.
Host OS	The operating system on the host where the node is running.
Connection String	The connection string for the node, which can be the IP address and port of the host that the node is connected to, or the Uniform Resource Locator (which is used as an alternative way to specify a JMX connection).
Connected	When checked (true), the node is currently connected to the Data Server via JMX.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the node.

Cluster Nodes Heatmap

This display allows you to view utilization data for all nodes in a cluster in a heatmap format. You can view heatmap data for **All Nodes**, **Inference** nodes, or **Cache** nodes by selecting the desired option from the **Nodes** drop down list. When you click on the heatmap for one of the nodes, the detailed data for that particular node displays in the “[Inference Node Summary](#)” display if you selected an inference node, or in the “[Storage Node Summary](#)” display if you selected a cache node.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.





23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.



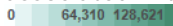
Filter By:

The display might include these filtering options:




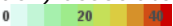
- Cluster** Choose a cluster for which you want to see metrics.
- Last Update** The date and time that the display was last updated.
- Node Count** The total number of nodes in the display.




Nodes	Select the type of nodes for which you want to view metrics. You can select from All Nodes , Inference , and Cache . Your selection in this drop down determines the available options in the Color Metric drop down.
Labels	Select this option to display labels in the heatmap for each of the nodes.
Log	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Color Metric	<p>Select the metric driving the heatmap display. The default is Alert Severity. Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the nodes by cluster, where each rectangle represents a node. Mouse-over any rectangle to display the current values of the metrics for the cluster. Click on a rectangle to drill-down to the associated "Storage Node Summary" display for a detailed view of metrics for that particular server. The available options in this drop down change depending on your selection in the Nodes drop down.</p> <p>Nodes: The following options are available when All Nodes is selected from the All Nodes drop down.</p> <p>Alert Severity The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar , where 2 is the greatest Alert Severity.</p> <p>2 -- Metrics that have exceeded their specified ALARMLEVEL threshold and have an Alert Severity value of 2 are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p>1 -- Metrics that have exceeded their specified WARNINGLEVEL threshold and have an Alert Severity value of 1 are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p>0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of 0 and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p> <p>Alert Count The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>JVM % CPU Used The total percentage of JVM CPU used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of JvmCpuPercentHigh, which is 75. The middle value in the gradient bar indicates the middle value of the range (the default is 38).</p> <p>JVM % Memory Used The total percentage of JVM Memory Used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of JvmMemoryUsedHigh, which is 75. The middle value in the gradient bar indicates the middle value of the range (the default is 38).</p>

Nodes: In addition to **Alert Severity**, **Alert Count**, **JVM % CPU Used**, and **JVM % Memory Used**, the following options are also available when **Inference** is selected from the **Nodes** drop down.

Received Events Rate	The rate of events received in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeChanRecvdRateHigh , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).
Rules Fired Rate	The rate of rules fired in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeRuleFiringRateHigh , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).
Total Rules Fired	The total number of rules fired in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the maximum count of rules fired in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Nodes: In addition to **Alert Severity**, **Alert Count**, **JVM % CPU Used**, and **JVM % Memory Used**, the following options are also available when **Cache** is selected from the **Nodes** drop down.

Backing Store Reads/sec	The rate of reads from the backing store in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeBackingStoreLoadRateHi , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).
Backing Store Writes/sec	The rate of writes to the backing store in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeBackingStoreStoreRateHi , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).
Backing Store Deletes/sec	The rate of deletes from the backing store in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeBackingStoreEraseRateHigh , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).
Concept Gets/sec	The rate of "gets" in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeConceptsGetRateHigh , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).


Concept Puts/sec	The rate of “puts” in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeConceptsPutRateHigh , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).
Concept Removes/sec	The rater of “removes” in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeConceptsRemoveRateHigh , which is 95 . The middle value in the gradient bar indicates the middle value of the range (the default is 48).
Object Table Size	The number of objects maintained in the cache in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of TbeNodeObjectTableSize , which is 10,000 . The middle value in the gradient bar indicates the middle value of the range (the default is 5,000).

Inference Node Summary

Use this display to view configuration and utilization data for a single inference node. View a list of all agents on the node, a Run-To-Completion Transaction Manager Report, and trend graphs tracing the rule execution rate for agents on the node. The rule execution rate is relative to the overall CPU and heap utilization for the engine's JVM.

NOTE: An inference node (also known as an engine or processing unit) is the container where one or more inference agents run. Generally, the agents in a given node implement different rule sets, and distributing nodes on different hosts provides fault tolerance and load balancing for the cluster. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Change the trend graph **Time Range** to “zoom in” on the graph and see more detail or “zoom out” from the graph to see larger trends over time. To change the time range, click Open Time Range , choose the date and time, and then click **OK**.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- 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.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

- Cluster** Choose a cluster for which you want to view metrics.
- Node** Choose a node for which you want to view metrics.

Fields and Data:

- Last Update** The date and time the data in the display was last updated.
- Cluster Name:** The name of the TIBCO BusinessEvents cluster with which the node is a member.
- BE Version:** The approximate TIBCO BusinessEvents version, as configured by the **connection** property. The exact version information is not available via JMX.
- Node ID:** A unique string that identifies the node.

Connection:	The JMX connection method specified in the connection property for a given engine. It is displayed as either a combination of the host and port fields (<host>:<port>), or the URL. This convention saves space on the display by avoiding empty fields. This information is provided as a convenience for those rare occasions where a user might wish to view the data directly in jconsole.
% CPU:	The percent of CPU used by the engine process. This value is provided by standard Java MBeans.
Heap used:	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
Heap max	The maximum amount of memory, in megabytes, that can be used by the JVM for heap space. This value is provided by standard Java MBeans.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Auto Startup	When checked (true), this feature is enabled.
Cache Aside	When checked (true), this feature is enabled.
Backing Store Enabled	When checked (true), this feature is enabled.
Storage Enabled	When checked (true), this feature is enabled.
Serialization Optimized	When checked (true), this feature is enabled.

RTC TXN Manager Report

Note: Fields in this display with an asterisk (*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Avg Action Txn Millisec	The average amount of time taken for an action transaction, in milliseconds.*
Avg Cache Queue Wait Time Millisec	The average cache queue wait time, in milliseconds.*
Avg Cache Txn Millisec	The average amount of time taken for a cache transaction, in milliseconds.*
Avg DB Ops Batch Size	The average database operation batch size.*
Avg DB Queue Wait Time Millisec	The average database queue wait time, in milliseconds.*
Avg DB Txn Millisec	The average amount of time taken for a database transaction, in milliseconds.*

Avg Successful Txn Time Millisec	The average amount of time taken for a successful transaction, in milliseconds.*
Last DB Batch Size	The size of the last database batch.*
Pending Actions	The total number of pending actions.*
Pending Cache Writes	The total number of pending cache writes.*
Pending DB Writes	The total number of pending database writes.*
Pending Events to Ack	The total number of pending events that need to be acknowledged.*
Pending Locks to Release	The total number of pending locks that need to be released.*
Total DB Txns Completed	The total number of database transactions that have been completed.*
Total Errors	The total number of errors.*
Total Successful Txns	The total number of successful transactions.*

Agents for this Node Table

Each row in the table is an agent associated with the node, with data in the row columns describing the agent.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Agent Count:	The number of agents currently in the table.
Agent ID	The agent's ID.
Agent Class	The agent's class. See TIBCO documentation for more information.
Type	The type of agent (Inference, Cache, Query, or Dashboard).*
Current State	The current state of the agent.*
Started	When checked, denotes that the agent is started.*
Suspended	When checked, denotes that the agent is suspended.*
Concurrent	When checked, denotes that it is a concurrent agent.*
Queue Capacity	The queue capacity for the agent.*
Queue Size	The queue size for the agent.*
Thread Count	The total number of threads for the agent.*

Total # Rules Fired	The total number of rules fired for the agent.*
Rules Fired	The number of rules fired.*
Rules/sec	The rate of rules fired for the agent.
Avg Receive Time	See TIBCO documentation for more information.*
Avg Txn Commit Time	The average amount of time taken to commit a transaction.*
Cache Queue Remaining	The total amount of remaining space on the cache queue.*
DB Ops Queue Remaining	The total amount of remaining space on the DB Operations queue.*
Hit Ratio	See TIBCO documentation for more information.*
Job Rate	See TIBCO documentation for more information.*
L1 Cache Max Size	The maximum size of the L1 cache.*
L1 Cache Size	The current size of the L1 cache.*
Max Active	See TIBCO documentation for more information.*
# Event Threads	The total number of currently active event threads.*
# Jobs	The total number of currently active jobs.*
Priority	See TIBCO documentation for more information.*
Read Only	See TIBCO documentation for more information.*
Txn Commit Count	The number of transactions committed by the agent.*
Txn Receive Count	The number of transactions received by the agent.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the agent.

Trend Graph

Shows metrics for the selected node.

% CPU -- Traces the amount of CPU used, in percent, by the node.

Rules/sec -- Traces the number of rules processed, per second, by the agent.

Heap-max -- Traces the maximum amount of heap space, in bytes, used by the node since the agent was started.

Heap-used -- Traces the current amount of heap space, in bytes, used by the agent.

Rules/sec for Agent Choose an agent from the drop-down menu.

Log Scale


This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero

When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.


Click **Restore to Now** to reset the time range end point to the current time.

Storage Node Summary

Use this display to view configuration details for a single cache node, the database connection pool status, as well as a list of all caches that are backed by the backing store (database). Also view trend graphs that trace utilization metrics such as CPU and heap memory usage.








NOTE: A storage node (also known as a cache node) provides fast access to events and concepts required during each RTC by the inference engines. Storage nodes also serve as buffers for reads and writes between the cluster and the backing store. For details, refer to TIBCO documentation.


Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Change the trend graph **Time Range** to “zoom in” on the graph and see more detail or “zoom out” from the graph to see larger trends over time. To change the time range click Open Time Range , choose the date and time, then click **OK**.



Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.
-  6,047 The number of items currently in the display.



Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.



Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

Cluster: Choose a cluster to see metrics for.

Node: Choose a node to see metrics for.

Fields and Data

Last Update The date and time the data was last updated in the display.

Cluster Name:	The name of the TIBCO BusinessEvents cluster with which the node is a member.
BE Version:	The approximate TIBCO BusinessEvents version, as configured by the connection property. The exact version information is not available via JMX.
Node ID:	A unique string that identifies the node.
Connection:	The JMX connection method specified in the connection property for a given engine. It is displayed as either a combination of the host and port fields (<host>:<port>), or the URL. This convention saves space on the display by avoiding empty fields. This information is provided as a convenience for those rare occasions where a user might wish to view the data directly in jconsole.
% CPU:	The amount of CPU, in percent, used by the node. This value is provided by standard Java MBeans.
Heap used:	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
Heap Max:	The maximum amount of memory, in megabytes, that can be used by the JVM for heap space. This value is provided by standard Java MBeans.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Auto Startup	When checked (true), this feature is enabled.
Cache Aside	When checked (true), this feature is enabled.
Backing Store Enabled	When checked (true), this feature is enabled.
Storage Enabled	When checked (true), this feature is enabled.
Serialization Optimized	When checked (true), this feature is enabled.

DB Connection Pool

Values describe status of the pool of database connections used by the cache agent to move data between the local caches and the database.

Note: Fields in this region with an asterisk (*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

Pool State	The state of the database connection pool.*
Auto Failover	The number of times auto failover has occurred.*
Failover Interval	The number of seconds taken for failover to take place.*
Cache Size	The cache size.*
# Connections Available	The total number of connections available.*
# Connections in Use	The total number of connections currently in use.*

Backing StoreTable

A cache node manages access to current events and concepts, buffering as necessary between local memory and a database. The Backing Store table provides a list of caches and the database select/insert/delete statistics for each cache.


Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.


Cache Name	The name of the cache.*
Active	When checked, denotes that the cache is active.*
Delete Avg Time	The average amount of time taken for a "delete" ("erase") operation in the Backing Store for the cache.*
Load Avg Time	The average amount of time taken for a "load" operation in the Backing Store for the cache.*
Store Avg Time	The average amount of time taken for a "store" operation in the Backing Store for the cache.*
Delete Total	The total number of "delete" operations in the Backing Store for the cache.*
Load Total	The total number of "load" operations in the Backing Store for the cache.*
Store Total	The total number of "store" operations in the Backing Store for the cache.*
Deletes	The number of "delete" operations during the last polling interval.*
Loads	The number of "load" operations during the last polling interval.*
Stores	The number of "store" operations during the last polling interval.*
Deletes/sec	The rate of "delete" operations in the node.
Loads/sec	The rate of "load" operations in the node.
Stores/sec	The rate of "store" operations in the node.
Object Table Trends	Shows metrics for the selected cluster/node combination: <ul style="list-style-type: none"> % CPU -- Traces the amount of CPU used, in percent, by the engine. Table Size -- Traces the number of unique objects cached in the local index table. Ext ID Tbl Size -- Traces the number of entries in the table of external IDs used as indexes by the backing store. Max Heap (MB) -- Traces the maximum amount of memory, in megabytes, that can be used by the JVM for heap space. Heap (MB) -- Traces the current heap space, in megabytes, in use by the JVM.
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero


When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Events / Concepts View



These displays present performance data for your BusinessEvents system. Displays in this View are:

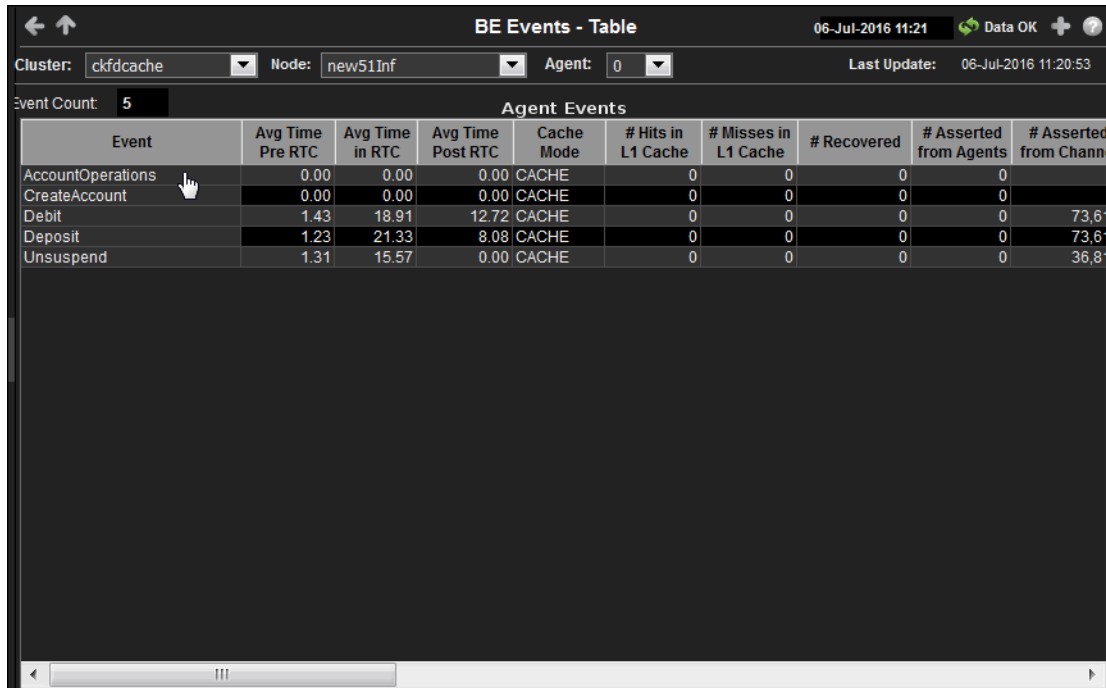
- ["Agent Events"](#)
- ["Agent Event Summary"](#)
- ["Event Cache Hits"](#)
- ["Event Hit Summary"](#)
- ["Concept Cache Hits"](#)
- ["Concept Hit Summary"](#)
- ["Channels"](#)
- ["All Inference Agents"](#)
- ["All RTC Reports"](#)

Agent Events

View run-time statistics for a selected group of agents. With TIBCO BusinessEvents, events are cached when they are out of scope, and deleted or persisted to the backing store when they are no longer useful. Clicking on a row in the table displays access patterns over time for the event in the ["Agent Event Summary"](#) display.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Sort  the table columns when all the rows cannot fit on the screen. For example, sort  the **Expired** column so that all expired nodes are listed at the top.



BE Events - Table 06-Jul-2016 11:21 Data OK + ?




Cluster: Node: Agent: Last Update: 06-Jul-2016 11:20:53



Event Count: 5

Agent Events

Event	Avg Time Pre RTC	Avg Time in RTC	Avg Time Post RTC	Cache Mode	# Hits in L1 Cache	# Misses in L1 Cache	# Recovered	# Asserted from Agents	# Asserted from Chann
AccountOperations	0.00	0.00	0.00	CACHE	0	0	0	0	
CreateAccount	0.00	0.00	0.00	CACHE	0	0	0	0	
Debit	1.43	18.91	12.72	CACHE	0	0	0	0	73.6
Deposit	1.23	21.33	8.08	CACHE	0	0	0	0	73.6
Unsuspend	1.31	15.57	0.00	CACHE	0	0	0	0	36.8

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.
- open commonly accessed displays.
- 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.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
-  Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Cluster:** Select the cluster containing the node and agent for which you want to view metrics.
- Node:** Select a node containing the agent for which you want to view metrics.
- Agent** Select the agent for which you want to view metrics.

Fields and Data:

Last Update: The date and time the data on the display was last updated.

Agent Events Table:

Each row in the table is a different event. Data in the row columns describe the event. The following fields are added by Monitor collection. The assertions/sec, modified/sec, and retracted/sec metrics are calculated from the corresponding counters as the delta between two successive samples divided by the polling interval.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.


Event Count:	The total number of events in the table.
Event	The name of the event.
Avg Time Pre RTC	The average amount of time taken for the event to begin its run to completion cycle.*
Avg Time in RTC	The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*
Avg Time Post RTC	The average amount of time taken by the event after its run to completion cycle has ended.*
Cache Mode	Lists the mode used by the event, which can be either CACHE (only) or MEMORY (only).*
# Hits in L1 Cache	The number of times data has been searched for in the L1 cache since the last data update.*
# Misses n L1 Cache	The number of times data has been searched for in the L1 cache, but was not found, since the last data update.*
# Recovered	The number of times data is not found in the L1 cache, but is found in a different cache, since the last data update.*
# Asserted from Agents	The number of times the event was asserted by an agent into the Rete network.*
# Asserted from Channel	The number of times the event was asserted into the Rete network via the channel.*
# Modified from Agents	The number of times the event was modified by an agent in the Rete network.*
# Modified from Channel	The number of times the event was modified in the Rete network via the channel.*
# Retracted from Agents	The number of times the event was retracted/deleted by an agent from the Rete network.*
# Retracted from Channel	The number of times the event was retracted/deleted from the Rete network via the channel.*
L1 Cache Hits/sec	The rate of L1 cache hits.
L1 Cache Misses/sec	The rate of L1 cache misses.
# Recovered / sec	The rate of recovered data.

Assertions/ sec (Agent)	The rate of event assertions into the Rete network by the agent.
Assertions/ sec (Channel)	The rate of event assertions into the Rete network via the channel.
Modifies/ sec (Agent)	The rate of event modifications in the Rete network by the agent.
Modifies/ sec (Channel)	The rate of event modifications in the Rete network via the channel.
Retractions /sec (Agent)	The rate of event retractions/deletions from the Rete network by the agent.
Retractions /sec (Channel)	The rate of event retractions/deletions from the Rete network via the channel.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the engine.

Agent Event Summary

View detailed performance metrics for an agent's event. You can view cache, RTC, event statistics by channel, and event trend data over a specified period of time.



**Filter By:**

The display might include these filtering options:

Note: Fields in this display with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.


Cluster:	Select the cluster for which you want to see metrics.
Node:	Select the node for which you want to see metrics.
Agent	Select the agent for which you want to see metrics.
Last Update	The date and time in which the data was last updated.
Expired	When checked (true), the Monitor has not received a response from the event for the amount of time specified by the \$tbeRowExpirationTime property (the default is 120 seconds). When the amount of time specified by the \$tbeRowExpirationTimeForDelete property elapses (the default is one day), the event data is deleted from the cache and display.
Event	The name of the event.
Cache Mode	Lists the mode used by the event, which can be either CACHE (only) or MEMORY (only).*
Cache Hits	The number of times data has been searched for in the L1 cache since the last data update.*
Cache Misses	The number of times data has been searched for in the L1 cache, but was not found, since the last data update.*
Avg Time	in RTC The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*
	Pre RTC The average amount of time taken for the event to begin its run to completion cycle.*
	Post RTC The average amount of time taken by the event after its run to completion cycle has ended.*
Event Stats by Channel	Asserted Total The total number of times the event was asserted into the Rete network via the channel.*
	Asserted Per Interval The number of times the event was asserted into the Rete network via the channel since the last data update.*
	Asserted Per Second The rate of event assertions into the Rete network via the channel.


Retracted Total	The total number of times the event was retracted/deleted from the Rete network via the channel.*
Retracted Per Interval	The number of event retractions/deletions from the Rete network.
Retracted Per Second	The rate of event retractions/deletions from the Rete network via the channel.
Modified Total	The total number of times the event was modified in the Rete network via the channel.*
Modified Per Interval	The number of event modifications in the Rete network via the channel.
Modified Per Second	The rate of event modifications in the Rete network via the channel.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Event Trends	Shows metrics for the selected event: <ul style="list-style-type: none"> Time in RTC-- Traces the event spends in the run to completion cycle. Asserted(/sec)-- Traces the number of events asserted into the Rete network (or the rate of event assertions per second depending on Use Rates setting). Retracted(/sec)-- Traces the number events retracted from the Rete network (or rate of event retractions per second depending on Use Rates setting). Modified(/sec)-- Traces the number of events modified in the Rete network (or rate of events modified per second depending on Use Rates setting).
Use Rates	When selected, this toggle allows you to view data in the trend graph in counts per second (asserted count per second, retracted count per second, and modified count per second) instead of the default counts per selected interval (asserted count, retracted count, modified count).
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero

When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.



Click **Restore to Now** to reset the time range end point to the current time.

Event Cache Hits

View cache performance metrics per event for a single cluster or **All Clusters**.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Sort  the table columns when all the rows cannot fit on the screen. For example, sort  the **Expired** column so that all expired nodes are listed at the top.

BE Event Cache Hits - Table 06-Jul-2016 10:39 Data OK +






Cluster: All Clusters Node: new51Inf Last Update: 06-Jul-2016 10:39:19


Event Count: 5

Node-Level Event Cache Hits

Event	Node	Cache Size	Get Avg Time	Put Avg Time	Remove Avg Time	Gets / sec	Puts / sec	Removes / sec	Get Count
AccountOperations	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
CreateAccount	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
Debit	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
Deposit	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
Unsuspend	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	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.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
-  Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Cluster:** Select a cluster for which you want to see metrics.
- Node:** Select a node for which you want to see metrics.
- Last Update** The date and time the data was last updated.

Node-Level Event Statistics Table:

Each row in the table is a different event, with data in the row columns describing the event.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Event Count:	The total number of events in the table.
Event	The name of the event.
Node	The name of the node.
Cache Size	The size of the event's cache.*
Get Avg Time	The average time taken for a "get" event for the node.*
Put Avg Time	The average time taken for a "put" event for the node.*
Remove Avg Time	The average time taken for a "remove" event for the node.*
Gets/sec	The rate of "get" operations for the event.
Puts/sec	The rate of "put" operations for the event.
Removes/sec	The rate of "remove" operations for the event.
Get Count	The total number of "get" operations for the event.*
Put Count	The total number of "put" operations for the event.*
Remove Count	The total number of "remove" operations for the event.*
Num Handles In Store	The number of handles in the Backing Store for the event.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the engine.

Event Hit Summary

View detailed event performance metrics for a single cluster or **All Clusters**, a node, and an event.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters**, a node and an event from the drop-down menus.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

Note: Fields in this display with an asterisk (*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

- Cluster:** Select a cluster containing the node and event for which you want to see metrics.
- Node:** Select a node containing the event for which you want to see metrics.
- Event** Select the event for which you want to see metrics.
- Last Update** The date and time in which the data was last updated.

Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.	
Cache Size	The size of the cache.*	
Handles in Store	The number of handles in the Backing Store for the event.*	
Avg Cache Access Time	Get	The average time taken for a "get" operation.*
	Put	The average time taken for a "put" operation.*
	Remove	The average time taken for a "remove" operation.*
Cache Access Stats	Get -- Total Hits	The total number of "get" operations for the event.*
	Get-- Hits	The number of "get" operations for the event since the last data update.*
	Get-- Hits/sec	The rate of "get" operations for the event.
	Put-- Total Hits	The total number of "put" operations for the event.*
	Put-- Hits	The number of "put" operations for the event since the last data update.*
	Put-- Hits/sec	The rate of "put" operations for the event.
	Remove --Total Hits	The total number of "remove" operations for the event.*
	Remove-- Hits	The number of "remove" operations for the event since the last data update.*
	Remove --Hits/sec	The rate of "remove" operations for the event.
Cache Access Trends	Shows metrics for the selected cluster/node/event combination:	
	Gets(/sec) -- Traces the number of "gets" (or rate of "gets" per second depending on Use Rates setting) for the event.	
	Puts(/sec) -- Traces the number of "puts" (or rate of "puts" per second depending on Use Rates setting) for the event.	
	Removes(/sec) -- Traces the number of "removes" (or rate of "removes" per second depending on Use Rates setting) for the event.	
	Use Rates	When selected, this toggle allows you to view data in the trend graph in counts per second ("get" operations count per second, "put" operations count per second, and "remove" operations count per second) instead of the default counts per selected interval ("get" operations count, "put" operations count, "remove" operations count).

Log Scale


This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero

When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.





By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

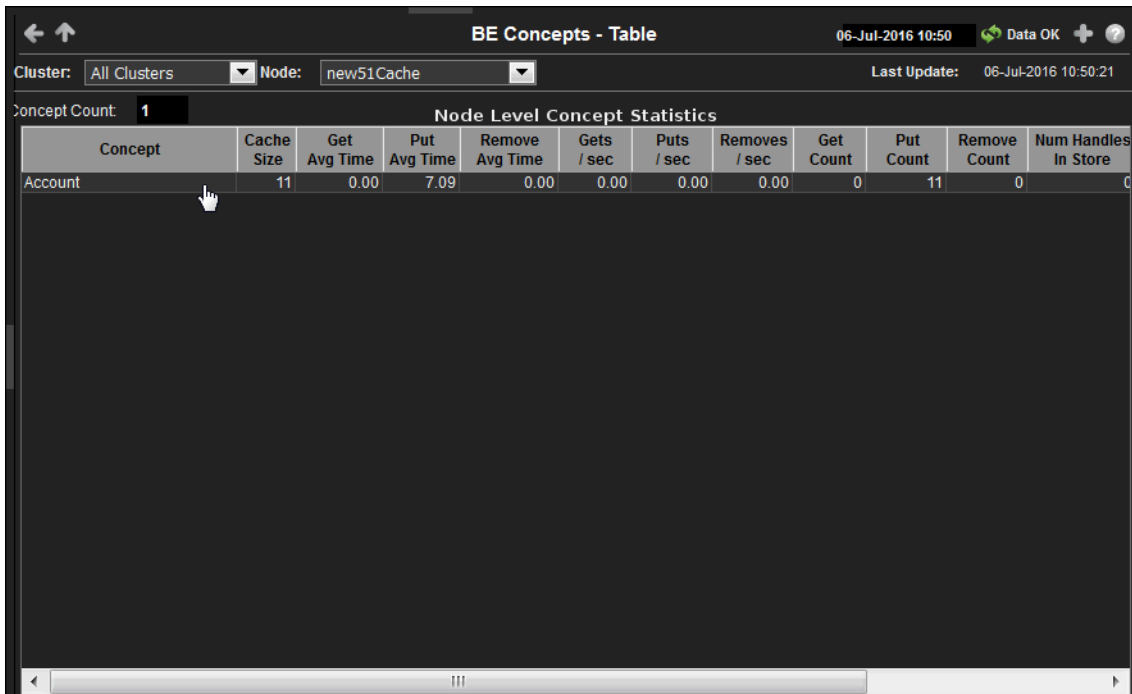
Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Concept Cache Hits








View a list of concepts and their run-time statistics. Choose a single cluster or **All Clusters** and a node from the drop-down menus.




Sort  the table columns when all the rows cannot fit on the screen. For example, sort  the **Expired** column so that all expired nodes are listed at the top.



Concept	Cache Size	Get Avg Time	Put Avg Time	Remove Avg Time	Gets / sec	Puts / sec	Removes / sec	Get Count	Put Count	Remove Count	Num Handles In Store
Account	11	0.00	7.09	0.00	0.00	0.00	0.00	0	11	0	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.
-   open commonly accessed displays.
-  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.
-  Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
-  Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Cluster:** Choose a cluster to see metrics for.
- Node:** Choose a node to see metrics for.
- Last Update** The date and time the data was last updated.

Node-Level Concept Statistics Table:

Each row in the table provides statistics regarding data access for a given BusinessEvents concept.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.


Concept Count:	The total number of concepts in the table.
Concept	The name of the concept.
Cache Size	The size of the concept's cache.*
Get Avg Time	The average time taken for a "get" operation.*
Put Avg Time	The average time taken for a "put" operation.*
Remove Avg Time	The average time taken for a "remove" operation.*
Gets/sec	The rate of "gets" for the concept.
Puts/sec	The rates of "puts" for the concept.
Removes/sec	The rate of "removes" for the concept.
Get Count	The total number of "gets" for the concept.*
Put Count	The total number of "puts" for the concept.*
Remove Count	The total number of "removes" for the concept.*
Num Handles In Store	The number of handles in the Backing Store for the concept.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the concept.

Concept Hit Summary

Use this display to view current and historic data for a single concept. Data in this display can be useful if your BusinessEvents system uses Cache object management. When Cache object management is used, concepts with a sufficiently long time to live (TTL) setting are cached.

Cache reference patterns for certain concepts may be related to incoming events (for example, customer purchase orders with associated inventory queries). The trend charts show the cache activity of such concepts, and might be useful in diagnosing the behavior of your application over time.

Choose a single cluster or **All Clusters**, a node and a concept from the drop-down menus. Change the trend graph **Time Range** to “zoom in” on the graph and see more detail or “zoom out” from the graph to see larger trends over time.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:


Fields in this table with an asterisk (*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields. The display might include these filtering options:


- Cluster:** Select a cluster containing the node and concept for which you want to see metrics.
- Node:** Select a node containing the concept for which you want to see metrics.
- Concept:** Select the concept for which you want to see metrics.


Fields and Data:

Note: Fields in this table with an asterisk (*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

Last Update	The date and time in which the data was last updated in the display.	
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.	
Cache Size	The size of the cache.*	
Handles in Store	The number of handles in the Backing Store.*	
Avg Cache Access Time	Get	The average time taken for a "get" operation.*
	Put	The average time taken for a "put" operation.*
	Remove	The average time taken for a "remove" operation.*
Cache Access Stats	Get -- Total Hits	The total number of "get" operations for the concept.*
	Get-- Hits	The number of "get" operations for the concept since the last data update.*
	Get-- Hits/sec	The rate of "get" operations for the concept.
	Put-- Total Hits	The total number of "put" operations for the concept.*
	Put--Hits	The number of "put" operations for the concept since the last data update.*
	Put-- Hits/sec	The rate of "put" operations for the concept.
	Remove --Total Hits	The total number of "remove" operations for the concept.*
	Remove--Hits	The number of "remove" operations for the concept since the last data update.*
	Remove --Hits/sec	The rate of "remove" operations for the concept.
Cache Access Trends	Shows metrics for the selected cluster/node/concept combination: Gets(/sec) -- Traces the number of "get" operations (or rate of "get" operations depending on Use Rates setting) for the concept. Puts(/sec) -- Traces the number of "put" operations (or rate of "put" operations depending on Use Rates setting) for the concept. Removes(/sec) -- Traces the number of "remove" operations (or rate of "remove" operations depending on Use Rates setting) for the concept.	

Use Rates	When selected, this toggle allows you to view data in the trend graph in counts per second ("get" operations count per second, "put" operations count per second, and "remove" operations count per second) instead of the default counts per selected interval ("get" operations count, "put" operations count, "remove" operations count).
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
Base at Zero	When this option is checked, zero is set as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

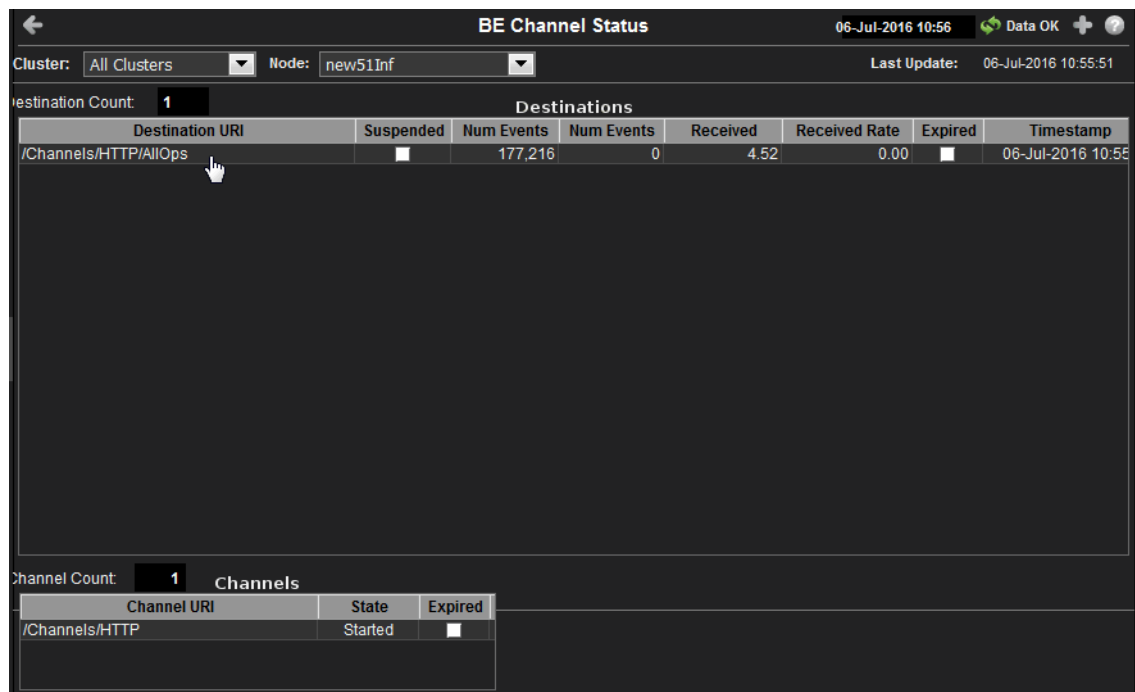
Click **Restore to Now** to reset the time range end point to the current time.

Channels

Use this display to view a list of destinations, which are sources and sinks of events. Destinations are potentially bi-directional, and the table indicates whether events are sent or received.

NOTE: Channels provide a class wrapper for destinations, and make it possible to enable or disable a group of destinations with one operation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus. Each row in the table is a different destination URI. Click a row to view channel details in the **Channels** table.



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.

Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Cluster:** Choose a cluster to see metrics for.
- Node:** Choose a node to see metrics for.

Destinations Table

Each row in the table provides data for a particular destination.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

- Destination Count:** The total number of destinations in the table.
- Destination URI** The Uniform Resource Identifier (URI) for the destination.*

Suspended	Denotes whether the destination is suspended.*
Num Events Received	The number of events received by the destination.*
Number of Events Sent	The number of events sent by the destination.*
Received Events Rate	The rate of events received by the destination.
Received Rate Last Interval	The rate of events received.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the destination.

Channels Table

Each row in the table provides data for a particular channel.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Channel Count:	The total number of channels in the table.
Channel URI	The Uniform Resource Identifier (URI) for the channel.*
State	The current state of the channel.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

All Inference Agents

Use this display to compare agent metrics across deployed engines and verify that the cluster is properly load-balanced. View a list of all the inference agents deployed in each cluster. You can view agent data for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the ["Cluster Summary"](#) display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All Clusters** from the drop-down menus. Each row in the table is a different agent.

Cluster	Node	Agent	Agent Name	Type	Current	Started	Suspended	Concurrent	Queue
ckddcache	new51inf	0	inference-class	INFERENCE	ACTIVATED	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,0
fdcache	newbe4inference	1	inference-class	Inference	Activated	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,0

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.
- 6,047** The number of items currently in the display.



Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.



Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

Cluster: Select the cluster for which you want to see metrics, or select **All Clusters** to see metrics for all clusters.

Table

Each row in the table provides details for an agent.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Agent Count: The number of agents currently in the table.

Cluster The name of the TIBCO BusinessEvents cluster.

Node The name of the node.

Agent ID	A unique string that identifies the agent.
Agent Name	The name of the agent.
Type	The type of agent (Inference, Cache, Query, or Dashboard).*
Current State	The current state of the agent.*
Started	When checked, denotes that the agent is started.*
Suspended	When checked, denotes that the agent is suspended.*
Concurrent	When checked, denotes that it is a concurrent agent.*
Queue Capacity	The queue capacity for the agent.*
Queue Size	The queue size for the agent.*
Thread Count	The total number of threads for the agent.*
Total # Rules Fired	The total number of rules fired for the agent.*
Rules/sec	The rate of rules fired for the agent.
Avg Receive Time	See TIBCO documentation for more information.*
Avg Txn Commit Time	The average amount of time taken to commit a transaction.*
Cache Queue Remaining	The total amount of remaining space on the cache queue.*
DB Ops Queue Remaining	The total amount of remaining space on the DB Operations queue.*
Hit Ratio	See TIBCO documentation for more information.*
Job Rate	See TIBCO documentation for more information.*
L1 Cache Max Size	The maximum size of the L1 cache.*
L1 Cache Size	The current size of the L1 cache.*
Max Active	See TIBCO documentation for more information.*
# Event Threads	The total number of currently active event threads.*
# Jobs	The total number of currently active jobs.*
Priority	See TIBCO documentation for more information.*
Read Only	See TIBCO documentation for more information.*

Txn Commit Count	The number of transactions committed by the agent.*
Txn Receive Count	The number of transactions received by the agent.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the destination.


All RTC Reports

Use this display to compare RTC metrics across deployed engines. View a list of all the inference engine RTC reports. You can view reports for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the [“Cluster Summary”](#) display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All Clusters** from the drop-down menus. Each row in the table is a different node.

Cluster	Node	Avg Action	Avg Cache Queue	Avg Cache	Avg DB Ops	Avg DB Qi	Avg DB	Avg Successful
ckfdcache	new51Inf	0.00	0.00	0.00	1.00	0.00	8.45	8.33
fdcache	newbe4infern	0.00	0.00	0.00	0.00	0.00	0.00	0.00

**Filter By:**

The display might include these filtering options:

Cluster: Select the cluster for which you want to see metrics, or select **All Clusters** to see metrics for all clusters.

RTC Txn Manager Reports Table

Each row in the table is a different report. Data in the row columns describe the report.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Report Count:	The number of reports currently in the table.
Cluster	The name of the TIBCO BusinessEvents cluster.
Node	The name of the node.
Avg Action Txn Millisec	The average amount of time taken for an action transaction, in milliseconds.*
Avg Cache Queue Wait Time Millisec	The average cache queue wait time, in milliseconds.*
Avg Cache Txn Millisec	The average amount of time taken for a cache transaction, in milliseconds.*
Avg DB Ops Batch Size	The average database operation batch size.*
Avg DB Queue Wait Time Millisec	The average database queue wait time, in milliseconds.*
Avg DB Txn Millisec	The average amount of time taken for a database transaction, in milliseconds.*
Avg Successful Txn Time Millisec	The average amount of time taken for a successful transaction, in milliseconds.*
Last DB Batch Size	The size of the last database batch.*
Pending Actions	The total number of pending actions.*

Pending Cache Writes	The total number of pending cache writes.*
Pending DB Writes	The total number of pending database writes.*
Pending Events to Ack	The total number of pending events that need to be acknowledged.*
Pending Locks to Release	The total number of pending locks that need to be released.*
Total DB Txns Completed	The total number of database transactions that have been completed.*
Total Successful Txns	The total number of successful transactions.*
Total Errors	The total number of errors.*
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO BusinessEvents > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the destination.

TIBCO BusinessWorks

The following TIBCO BusinessWorks Views can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO BusinessWorks is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central.

- [“BW Applications”](#)
- [“BW Containers”](#)
- [“BW AppNodes”](#)
- [“BW AppSlices”](#)
- [“BW Processes”](#)
- [“BW5 Engines”](#)
- [“BW5 Processes”](#)
- [“BW5 Activities”](#)
- [“BW5 Servers”](#)

BW Applications

These displays present process performance data for your BusinessWorks applications and AppSpaces across BusinessWorks Domains. Process metrics are totaled by application. Use these displays to monitor critical alerts for all your BusinessWorks applications, and investigate those alerts in lower-level displays. Displays in this View are:

- [“BW All Applications Heatmap” on page 941](#): A color-coded heatmap view of selected application performance metrics.
- [“BW All Applications Table” on page 944](#): A tabular view of all available application performance data in this BusinessWorks View.
- [“BW Single Application Summary” on page 947](#): Current and historical metrics for a single application.

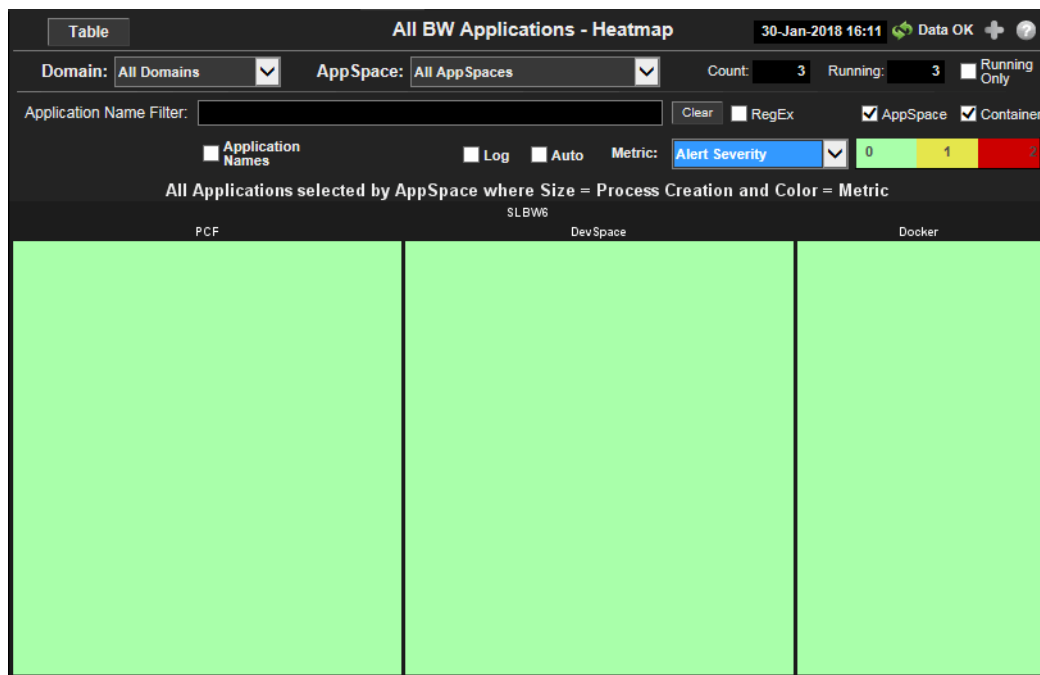
BW All Applications Heatmap

View the most critical BusinessWorks application alert states pertaining to process creation and execution for all nodes on which the applications are deployed. Use this display to quickly identify applications with critical alerts.

Each rectangle in the heatmap represents an application. The rectangle color indicates the most critical alert state associated with the application. The rectangle size represents process creation across applications; a larger size is a larger value.

Choose a domain and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **Application Name Filter** field to limit data shown in the display. Use the **Application Names** check-box ☒ to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate an application by clicking a rectangle in the heatmap to view details in the [“BW Single Application Summary”](#) display.












Title Bar (possible features are):








- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Domain:** Select the domain for which you want to view data in the display.
- AppSpace** Select the AppSpace for which you want to view data in the display.
- Application Name Filter** Enter a string (all or part of an application name) to filter the data shown in the display. If you enter part of an application name, you must enter "*" before and/or after the string. For example, if you have an application named AppNameOne, you could filter using *Name*, *NameOne, or AppName*.
- Clear** Clears the **Application Name Filter** entries from the display.
- RegEx** Toggles the **Application Name Filter** to accept Regular Expressions for filtering. For example, if your application name is AppNameOne and this option was toggled on, you could enter "Name" (without using "*" to display the application in the heatmap).
- AppSpace** When selected, those AppNodes deployed in an AppSpace display in the heatmap.
- Container** When selected, those AppNodes deployed in a container display in the heatmap.

Application Names	Check to include labels in the heatmap.
Fields and Data:	
Count:	The total number of AppSpaces currently shown in the display.
Running	The total number of AppSpaces currently running in the display.
Running Only	Select to show only running applications in the display.
Log	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Active Count	<p>The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Completed Count	<p>The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Suspended Count	<p>The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Failed Count	<p>The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>

Created / sec	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Suspended / sec	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Failed / sec	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Exec Time / sec	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Most Recent Exec Time	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Exec Time	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Elapsed Time	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

BW All Applications Table

View BusinessWorks data shown in the [“BW All Applications Heatmap”](#), and additional details, in a tabular format.

Each row in the table is an application. Choose a domain and AppSpace from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Application Name Filter** field to limit data shown in the display.

Drill-down and investigate by clicking a row to view details for the selected application in the [“BW Single Application Summary”](#) display.

Heatmap All BW Applications - Table 16-Nov-2017 13:41 Data OK +

Domain: All Domains AppSpace: All AppSpaces Count: 6 Running: 6 ☐ Running Only

Application Name Filter: Clear ☐ RegEx ☒ Appspace ☒ Container

Domain	AppSpace	Name	Alert Level	Alert Count	State	Deployment	AppSource
SLBW6	DevSpace	tibco.bw.sample.binding.rest.BookStore.applicati			Running	Appspace	
SLBW6	Docker	tibco.bwce.sample.binding.rest.BookStore.applic:			Running	Container	
SLBW6	PCF	tibco.bwce.sample.binding.rest.BookStore.applic:			Running	Container	
standalone	standalone-3	tibco.bwce.sample.binding.rest.BookStore.applic:			Running	Container	
standalone	standalone-9	tibco.bwce.sample.binding.rest.BookStore.applic:			Running	Container	
standalone	standalone-c	tibco.bwce.sample.binding.rest.BookStore.applic:			Running	Container	

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.
- open commonly accessed displays.
- 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.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:




- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- Application Name Filter** Enter a string (all or part of a application name) to filter the data shown in the display. If you enter part of an application name, you must enter "*" before and/or after the string. For example, if you have an application named AppNameOne, you could filter using *Name*, *NameOne, or AppName*.
- Clear** Clears the **Application Name Filter** entries from the display.
- RegEx** Toggles the **Application Name Filter** to accept Regular Expressions for filtering. For example, if your application name is AppNameOne and this option was toggled on, you could enter "Name" (without using "*" to display the application in the table).
- AppSpace** When selected, those AppNodes deployed in an AppSpace display in the table.
- Container** When selected, those AppNodes deployed in a container display in the table.

Fields and Data:

Count:	The total number of applications in the AppSpace.
Running	The total number of applications currently running in the AppSpace.
Running Only	Select to show only running applications in the display.

Table:

Each row in the table is a different application.

Domain	The domain in which the application resides.
AppSpace	The AppSpace in which the application resides.
Name	The name of the application.
Alert Level	<p>The most critical alert state for alerts in the row:</p> <p> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</p> <p> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</p> <p> Green indicates that no metrics have exceeded their alert thresholds.</p>
Alert Count	The total number of active alerts for the application.
State	The current status of the application. Valid values are Running and Stopped .
AppNodes	The total number of AppNodes associated with the application.
Active Processes	The number of currently active application processes.
Suspended Processes	The number of suspended application processes.
Failed Processes	The number of failed application processes.
Completed Processes	The number of completed application processes.
Created /sec	The number of application processes created per second.
Suspended / sec	The number of application process suspensions per second.
Failed /sec	The number of application process failures per second.
Exec Time/ sec	The number of processes executed per second.
Recent Exec Time	The number of seconds for the most recently executed process.
Average Exec Time	The average number of seconds for all processes to execute.
Version	The application version.
Module	The application module.
Shared Module	The shared module, if any.
Time Stamp	The date and time the row data was last updated.

Source	Name of RTView Data Server sending this data (or localhost).
Expired	When checked, data has not been received from this host in the specified amount of time.

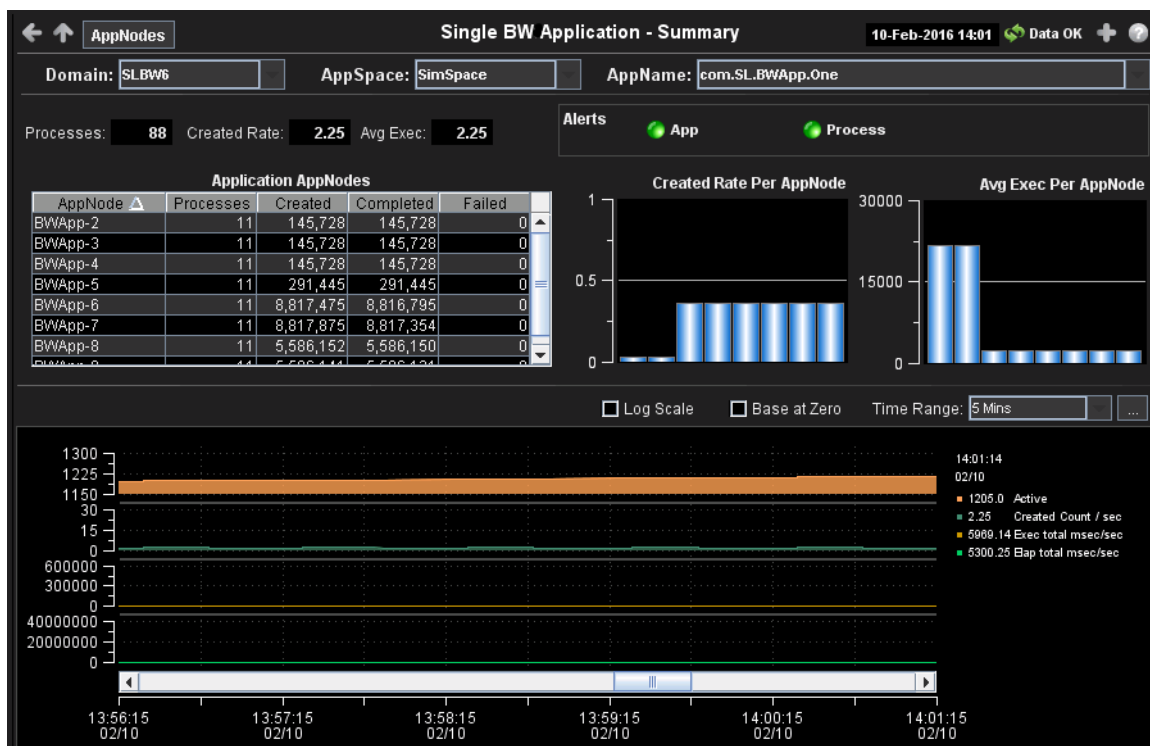
BW Single Application Summary

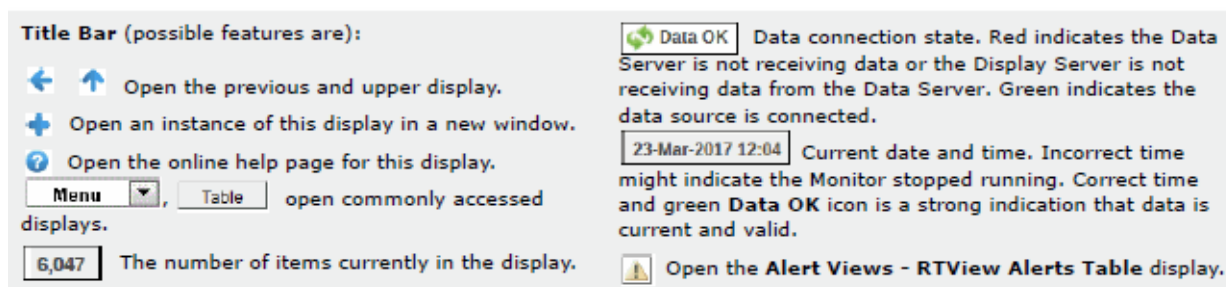
View current and historical metrics for a single BusinessWorks application across multiple nodes. Use this display to investigate performance issues of application AppNodes within an AppSpace. Use this display to view all available data for each AppNode by Domain and AppSpace.

This display includes a list of AppNodes with their host names and memory metrics, bar graphs per AppNode for process creation and execution, and trend graphs of process creation and execution metrics.

The summary display also shows the AppNodes of the deployment and process metrics totaled by AppNode. This is useful to see the deployment and load balancing of the Application in current and historical time.

Choose a domain, AppSpace and Application from the drop-down menus. Drill-down and investigate by clicking an AppNode in the table to view details in the “[BW Single AppNode Summary](#)” display.



**Filter By:**

The display might include these filtering options:

- Domain:** Select the domain for which you want to view data in the display.
- AppSpace:** Choose the AppSpace for which you want to view data in the display.
- AppName:** Choose the AppName for which you want to view data in the display.

Fields and Data:

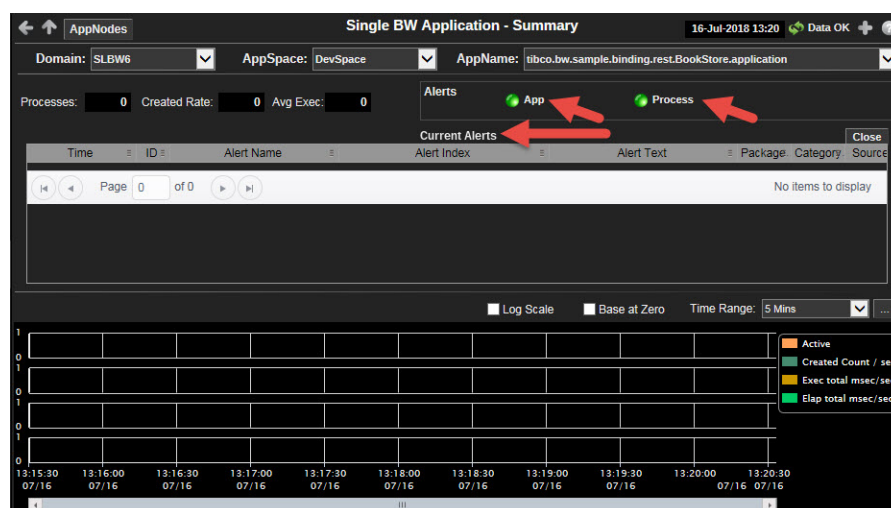
- Processes:** The number of processes currently running for the selected application.
- Created Rate:** The number of processes created per second for the selected application.
- Avg Exec:** The average number of seconds for processes to execute for the selected application.

Alerts

Indicates the greatest severity level and the number of open **App** and **Process** alerts for the selected application. Values range from **0** to **2**, where **2** is the greatest Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

Click on the alert indicator to display a table listing the current alerts for the selected application. Click the **Close** button (for the current alerts table) to close the table.



Application Appnodes Table

Each row in the table is a different AppNode. Column values describe processes for the selected application on that AppNode. Click a row to view AppNode details in the ["BW Single AppNode Summary"](#) display.

AppNode	The name of the AppNode.
Processes	The number of processes currently running on the AppNode.
Created	The total number of processes created on the AppNode.
Completed	The total number of completed processes on the AppNode.
Failed	The total number of failed processes on the AppNode.

Created Rate Per AppNode Bar Graph

The bar graph shows the current process creation rate per AppNode. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

AvgExec Per AppNode Bar Graph

The bar graph shows the current average process execution rate per AppNode for the selected application. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

Trend Graphs


Traces the sum of process metrics across all processes in all slices of the selected application.


Active	Traces the number of currently active application processes.
Created Count / sec	Traces the number of created application processes.
Exec total msec/sec	Traces the rate at which the application is accumulating process execution time, in milliseconds per second.
Elap total msec/sec	Traces the rate at which the application accumulates process elapsed time, in milliseconds per second.
Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Base at Zero



Select to use zero (0) as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

BW Containers

These displays present process performance data for your BusinessWorks containers across BusinessWorks Domains. Process metrics are totaled by container. Use these displays to monitor critical alerts for all your BusinessWorks containers, and investigate those alerts in lower-level displays. Displays in this View are:

- **"All Containers Heatmap" on page 950:** A color-coded heatmap view of selected container performance metrics.
- **"All Containers Table" on page 953:** A tabular view of all available container performance data in this BusinessWorks View.
- **"Single Container Summary" on page 956:** Current and historical metrics for a single container.

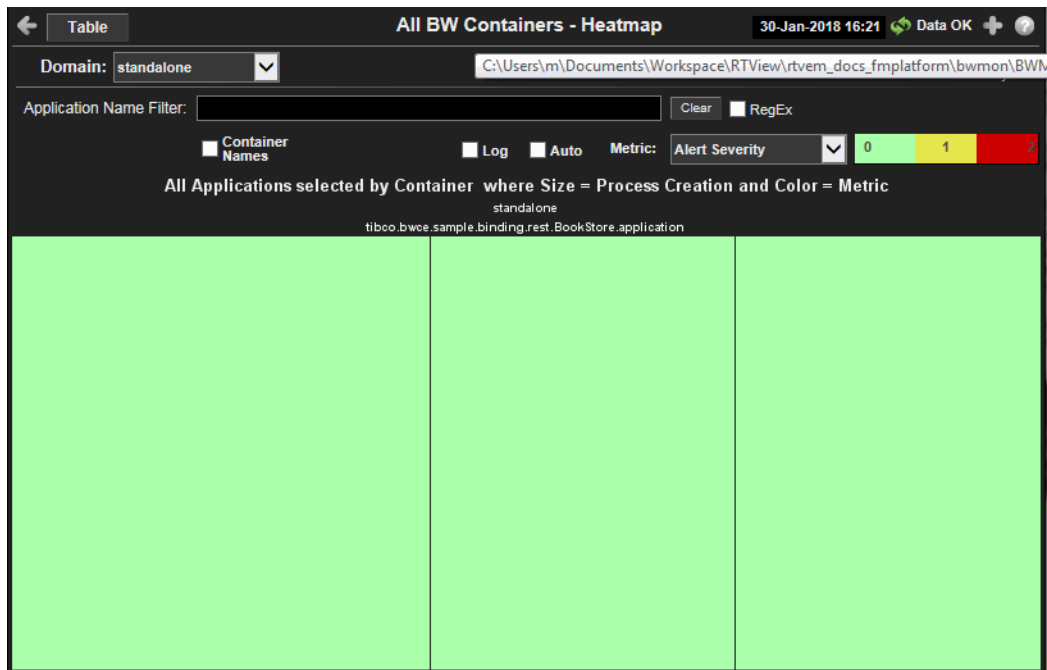
All Containers Heatmap

View the most critical BusinessWorks container alert states pertaining to process creation and execution for all nodes on which the containers are deployed. Use this display to quickly identify containers with critical alerts.

Each rectangle in the heatmap represents a container. The rectangle color indicates the most critical alert state associated with the container. The rectangle size represents process creation across containers; a larger size is a larger value.

Choose a domain and AppSpace from the drop-down menus containing the containers for which you want to view metrics. By default, this display shows the heatmap based on the **Alert Severity** metric, but you can select a different metric from the **Metric** drop-down menu to view the heatmap based on the selected metric. To view data shown for a specific container(s) in the display, enter a string in the **Container Name Filter** field. Use the **Container Names** check-box ☒ to include or exclude labels in the heatmap. You can mouse over a rectangle to see additional metrics.

Drill-down and investigate a container by clicking a rectangle in the heatmap, which opens the details for the selected container in the [“Single Container Summary”](#) display.











Title Bar (possible features are):








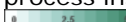
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Domain:** Select the domain for which you want to view data in the display.
- Application Name Filter** Enter a string (all or part of an application name) to filter the data shown in the display. If you enter part of a application name, you must enter "*" before and/or after the string. For example, if you have an application named AppNameOne, you could filter using *Name*, *NameOne, or AppName*. You can also enable the **Regex** toggle to just enter a portion of the application name.

	Clear	Clears the Application Name Filter entries from the display.
RegEx		Toggles the Application Name Filter to accept Regular Expressions for filtering. For example, if your application name is AppNameOne and this option was toggled on, you could enter "Name" (without using "*" to display the application in the heatmap).
Container Names		Check to include container name labels in the heatmap.
Fields and Data:		
Count:		The total number of containers currently shown in the display.
Running		The total number of containers currently running in the display.
Running Only		Select to show only running containers in the display.
Log		Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto		Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric		Choose a metric to view in the display.
	Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
	Alert Count	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
	Active Count	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	Completed Count	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	Suspended Count	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Failed Count	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Created / sec	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Suspended / sec	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Failed / sec	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Exec Time / sec	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Most Recent Exec Time	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Exec Time	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Elapsed Time	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

All Containers Table

This display provides a view of the most critical BusinessWorks container alert states pertaining to process creation and execution for all nodes on which the containers are deployed in a tabular format. Use this display to quickly identify containers with critical alerts. Each row in the table is a container in the selected domain.

Select a domain from the drop-down menu to view associated containers and, optionally, enter a string in the **Container Name Filter** field to further limit the list of containers shown in the display. You can click a column header to sort column data in numerical or alphabetical order.

To view additional details for a specific container, drill-down and investigate by clicking the row in the table for the desired container, which opens the “Single Container Summary” display.

Heatmap

All BW Containers - Table

30-Jan-2018 16:36

Data OK

Domain: standalone

Count: 3

Running: 3

Running Only

Container Name Filter: tibco

Clear

RegEx

Name	Alert Level	Alert Count	State	Active Processes	Container	Module
tibco.bwce.sample.binding.rest.BookStore.application			Running	0	standalone-1	tibco.bwce.sa
tibco.bwce.sample.binding.rest.BookStore.application			Running	0	standalone-a	tibco.bwce.sa
tibco.bwce.sample.binding.rest.BookStore.application			Running	0	standalone-e	tibco.bwce.sa

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Domain:

Choose a domain to show data for in the display.
- Container Name Filter

Enter a string (all or part of a container name) to filter the data shown in the display. If you enter part of a container name, you must enter “*” before and/or after the string. For example, if you have a container named ContNameOne, you could filter using *Name*, *NameOne, or ContName*.
- Clear

Clears the **Container Name Filter** entries from the display.
- RegEx

Toggles the **Container Name Filter** to accept Regular Expressions for filtering. For example, if your application name is ContNameOne and this option was toggled on, you could enter “Name” (without using “*” to display the container in the table).

Fields and Data:

- Count:




The total number of containers listed in the table.

Running The total number of containers that are currently running.

Running Only Select to show only running containers in the display.

Table:

Each row in the table is a different application.


Name	The name of the container.
Alert Level	<p>The most critical alert state for alerts in the row:</p> <p> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</p> <p> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</p> <p> Green indicates that no metrics have exceeded their alert thresholds.</p>
Alert Count	The total number of active alerts in the container.
State	The current status of the application. Valid values are Running and Stopped .
Active Processes	The number of currently active processes in the container.
Container	The name of the container.
Module	The name of the container module.
Version	The version of the container.
Average Exec Time	The average number of seconds for all processes to execute.
Completed Processes	The number of completed processes in the container.
Expired	When checked, data has not been received from this host in the specified amount of time.
Failed Processes	The number of failed processes in the container.
Recent Exec Time	The number of seconds for the most recently executed process.
Created /sec	The number of processes created per second in the container.
Failed /sec	The number of process failures per second in the container.
Suspended / sec	The number of process suspensions per second in the container.
Exec Time/ sec	The number of processes executed per second in the container.
Suspended Processes	The number of suspended application processes in the container.
Time Stamp	The date and time the row data was last updated.
Source	Name of RTView Data Server sending this data (or localhost).

Single Container Summary



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.



Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.



Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

Domain: Select the domain for which you want to view data in the display.

Container Choose the container for which you want to view data in the display.

Fields and Data:

Processes: The number of processes currently running on the selected container.

Created Rate: The number of processes created per second on the selected container.

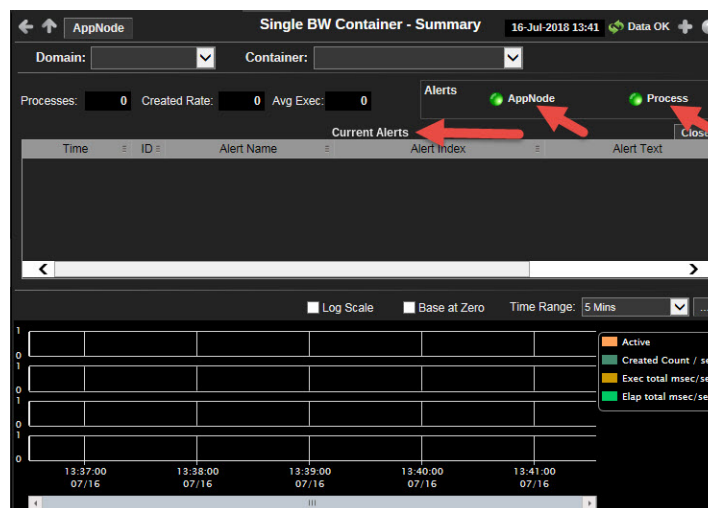
Avg Exec: The average number of seconds for processes to execute on the selected container.

Alerts

Indicates the greatest severity level and the number of open **AppNode** and **Process** alerts for the selected container. Values range from **0** to **2**, where **2** is the greatest Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

Click on the alert indicator to display a table listing the current alerts for the selected container. Click the **Close** button (for the current alerts table) to close the table.



Application Name The name of the application running on the container.

AppNode Information

CPU %	The percentage of CPU used by the AppNode.
Used %	The percentage of memory used by the AppNode.
Free MB	The amount of free memory, in megabytes.
Total MB	The total amount of used and free memory, in megabytes.

Created Rate Bar Graph

The bar graph shows the current process creation rate per AppNode. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.


AvgExec Bar Graph

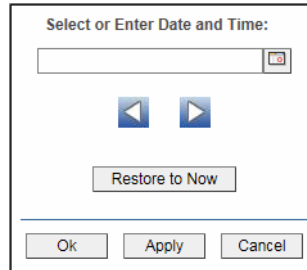
The bar graph shows the current average process execution rate per AppNode for the selected application. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

Trend Graphs


Traces the sum of process metrics across all processes in all slices of the selected container.



Active	Traces the number of currently active application processes on the container.
Created Count / sec	Traces the number of created application processes on the container.
Exec total msec/sec	Traces the rate at which the application is accumulating process execution time, in milliseconds per second, on the container.

- Elap total msec/sec** Traces the rate at which the application accumulates process elapsed time, in milliseconds per second, on the container.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath these arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

BW AppNodes

These displays present internal JVM memory and host CPU utilization for BusinessWorks AppNodes and their resources. This is useful because the AppNode performance is dependent on both internal and external factors and they sometimes interact. Displays in this View are:

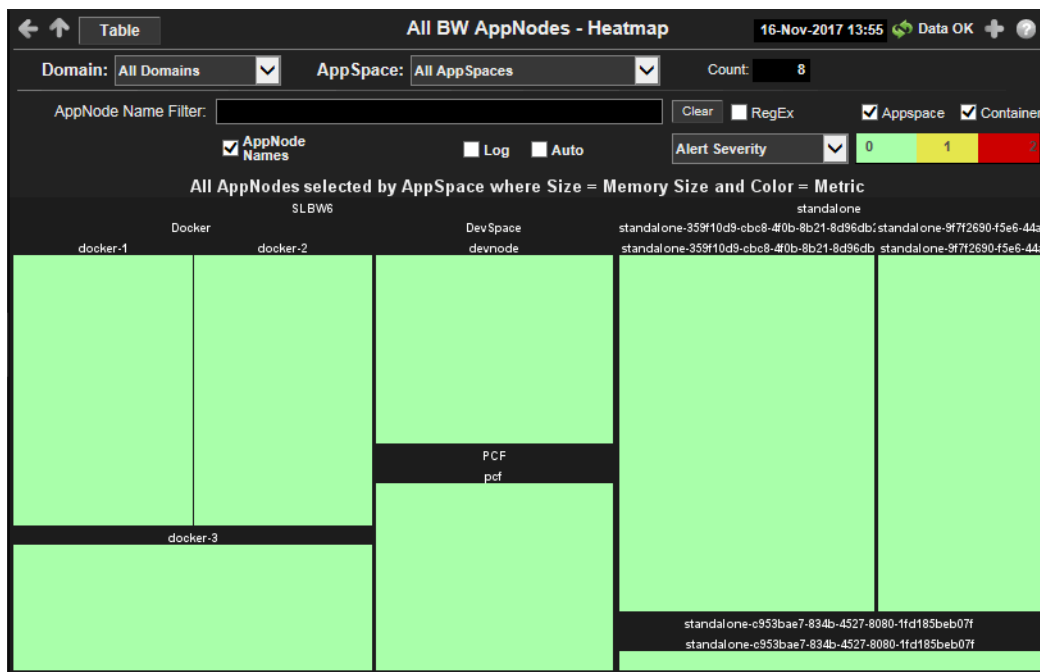
- ["BW All AppNodes Heatmap" on page 958](#): A color-coded heatmap view of utilization metrics.
- ["BW All AppNodes Table" on page 961](#): A tabular view of all available utilization data in this BusinessWorks View.
- ["BW Single AppNode Summary" on page 963](#): Current and historical metrics for a single AppNode.

BW All AppNodes Heatmap

View the most critical JVM memory and host resource utilization for BusinessWorks AppNodes. Use this display to quickly identify AppNodes with critical alerts.

Each rectangle in the heatmap represents an AppNode. The rectangle color indicates the most critical alert state associated with the AppNode. The rectangle size represents the maximum memory used in the rectangle; a larger size is a larger value.

Choose a domain and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **AppNode Name Filter** field to limit data shown in the display. Use the **AppNode Names** check-box ☒ to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the “[BW Single AppNode Summary](#)” display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.









23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.


Filter By:

The display might include these filtering options:


- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- Count:** The total number of AppNodes in the AppSpace.

AppNode Name Filter	Enter a string to limit data shown in the display.
Clear	Clears the Application Name Filter entries from the display.
RegEx	Toggles the Search Text field to accept Regular Expressions for filtering.
AppSpace	When selected, those AppNodes deployed in an AppSpace display in the heatmap.
Container	When selected, those AppNodes deployed in a container display in the heatmap.
AppNode Names	Check to include labels in the heatmap.
Log	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
CPU Used%	The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Memory Used%	The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Active Processes	The number of currently active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Created Processes

The number of processes created in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Created /sec

The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

BW All AppNodes Table

View BusinessWorks data shown in the [“BW All AppNodes Heatmap”](#) display, and additional details, in a tabular format. Use this display to view all available data for each AppNode by Domain and AppSpace.









Each row in the table is an AppNode. Choose a domain and AppSpace from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Application Name Filter** field to limit data shown in the display.


Drill-down and investigate by clicking a row to view details for the selected AppNode in the [“BW Single AppNode Summary”](#) display.

Heatmap **All BW AppNodes - Table** 16-Nov-2017 14:13 Data OK

Domain: All Domains AppSpace: All AppSpaces Count: 8

AppNode Name Filter: Clear ☐ RegEx ☒ Appspace ☒ Container

Domain	AppSpace	AppNode	Alert Level	Alert Count	State	Host	Deployment	Process Count	Active Thread
SLBW6	DevSpace	devnode			ACTIVE	qawin5(SLBW6)	Appspace	2	
SLBW6	Docker	docker-1			ACTIVE	oi7-20-1(SLBW6)	Container	2	
SLBW6	Docker	docker-2			ACTIVE	oi7-20-2(SLBW6)	Container	2	
SLBW6	Docker	docker-3			ACTIVE	oi7-20-3(SLBW6)	Container	1	
SLBW6	PCF	pcf			ACTIVE	sihost44(SLBW6)	Container	2	
standalone	standalone-35	standalone-359f			ACTIVE	oi7-20-4	Container	2	
standalone	standalone-9f	standalone-9f7f2			ACTIVE	oi7-20-5	Container	2	
standalone	standalone-c5	standalone-c553			ACTIVE	oi7-20-6	Container	2	

**Filter By:**

The display might include these filtering options:

Domain:	Choose a domain to show data for in the display.
AppSpace	Choose an AppSpace to show data for in the display.
Count:	The total number of rows in the table.
AppNode Name Filter	Enter a string to limit data shown in the display.
Clear	Clears the Application Name Filter entries from the display.
RegEx	Toggles the Search Text field to accept Regular Expressions for filtering.
AppSpace	When selected, those AppNodes deployed in an AppSpace display in the AppNodes table.
Container	When selected, those AppNodes deployed in a container display in the AppNodes table.

Count: The total number of rows in the table.

Table:

Column values describe the AppNode.

Domain	The domain in which the AppNode resides.
AppSpace	The AppSpace in which the AppNode resides.
AppNode	The name of the AppNode.
Alert Level	The most critical alert state for alerts in the row: <div style="display: flex; align-items: center;"> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. </div> <div style="display: flex; align-items: center;"> ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. </div> <div style="display: flex; align-items: center;"> ● Green indicates that no metrics have exceeded their alert thresholds. </div>
Alert Count	The total number of active alerts for the AppNode.
Host	The host on which the AppNode resides.
Process Count	The number of processes running.
Active Threads	The number of currently active threads.
Total Memory	The total amount of used and free memory, in megabytes.
Used Memory	The amount of used memory, in megabytes.

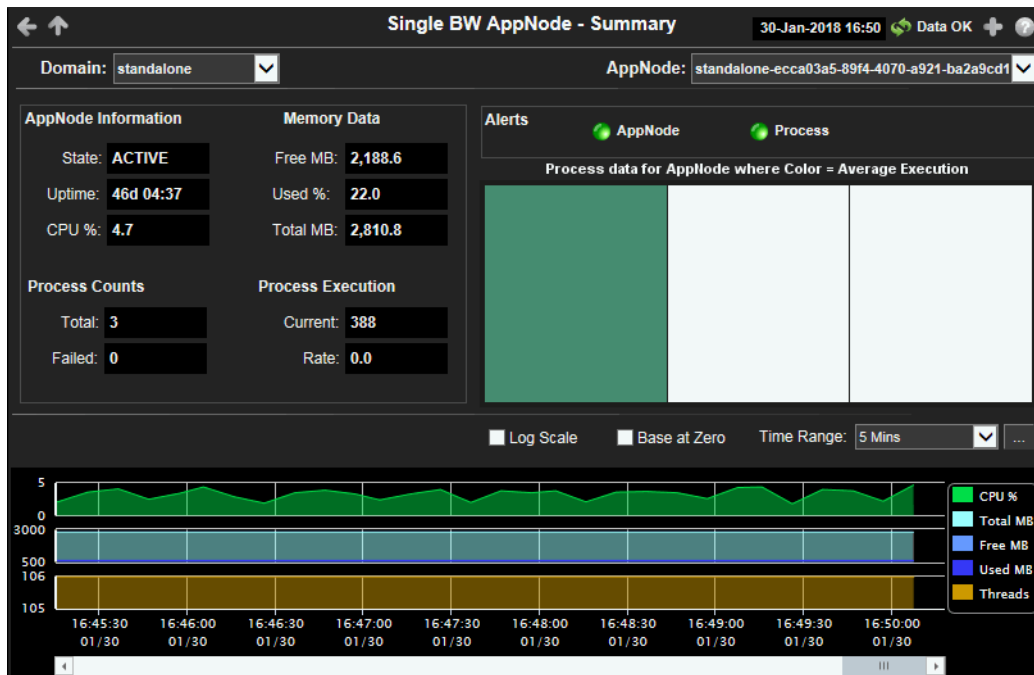
Free Memory	The amount of free memory, in megabytes.
Used Memory%	The percent (%) used memory.
Used CPU%	The percent (%) used CPU.
System Process ID	A unique string identifier for the process.
Up Since	The date and time the AppNode was last started.
Active Processes	The number of currently active processes.
Suspended Processes	The number of suspended application processes.
Failed Processes	The number of failed application processes.
Completed Processes	The number of completed application processes.
Created /sec	The number of application processes created per second.
Suspended / sec	The number of application processes suspended per second.
Failed /sec	The number of failed application processes per second.
Exec Time / sec	The number of application processes executed per second.
Recent Exec Time	The number of seconds for the most recently executed process.
Average Exec Time	The average number of seconds for all processes to execute.
Time Stamp	The date and time the row data was last updated.
Source	Name of RTView Data Server sending this data (or localhost).

BW Single AppNode Summary

View current and historical utilization and performance metrics for a single BusinessWorks AppNode. Use this display to investigate performance issues on an AppNode.

This display includes a heatmap showing most critical alerts pertaining to process execution, and trend graphs tracing CPU utilization and thread count.

Choose a domain, AppSpace and AppNode from the drop-down menus. Use the **Time-Range** to “zoom-in” or “zoom-out” on a specific time frame in the trend graph. Drill-down and investigate by clicking an AppNode in the table to view details in the [“BW Single AppNode Summary”](#) display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.




Filter By:

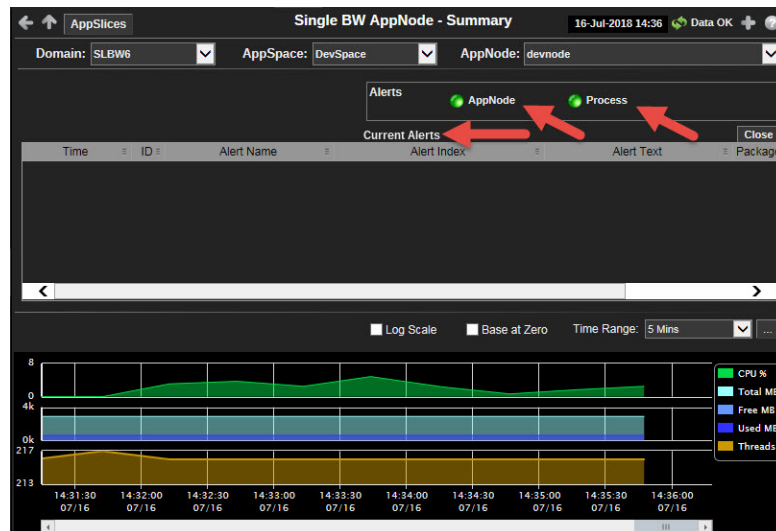
The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.

Fields and Data:

- | | | |
|----------------------------|-----------------|--|
| AppNode Information | Uptime: | The number of days, hours and minutes since the AppNode started. |
| | CPU% | The percent (%) CPU used on the AppNode. |
| | Threads: | The number of currently active threads for the AppNode. |
| Memory Data | Free: | The amount of available memory on the AppNode. |
| | Used% | The percent (%) memory used on the AppNode. |

	Total	The total amount of memory on the AppNode.
Process Counts	Total:	The number of currently active processes for the AppNode.
	Failed:	The number of failed processes for the AppNode.
Process Execution	Current	The number of processes executed by the AppNode.
	Rate:	The number of processes executed per second.
Alerts	Indicates the greatest severity level and the number of open AppNode and Process alerts for the selected AppNode. Values range from 0 to 2 , where 2 is the greatest Severity:	
	<div><div> One or more alerts exceeded their ALARM LEVEL threshold.</div><div> One or more alerts exceeded their WARNING LEVEL threshold.</div><div> No alert thresholds have been exceeded.</div></div>	
	Click on the alert indicator to display a table listing the current alerts for the selected AppNode. Click the Close button (for the current alerts table) to close the table.	




Heatmap

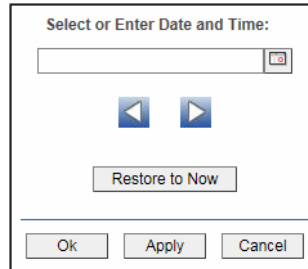
Each rectangle in the heatmap represents an AppSlice. The rectangle color indicates the most critical **Average Exec Time** alert state associated with the AppSlice. The rectangle size represents the maximum number of processes executed in the rectangle; a larger size is a larger value. Click a rectangle to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

Trend Graphs


Traces the sum of process metrics across all processes for all applications on the AppNode.



- **CPU%:** The percent (%) CPU used on the AppNode.
- **Total MB:** The amount of memory used.
- **Free MB:** The amount of available memory.
- **Used MB:** The amount of used memory.
- **Threads:** The number of threads.

- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

BW AppSlices

These displays present process metrics totaled by Application and AppNode for AppSlices. This is useful to see how the application is distributed and how each part of it is performing. The AppSlice is the part of an application running on a specific AppNode when the application is deployed to multiple AppNodes. Displays in this View are:

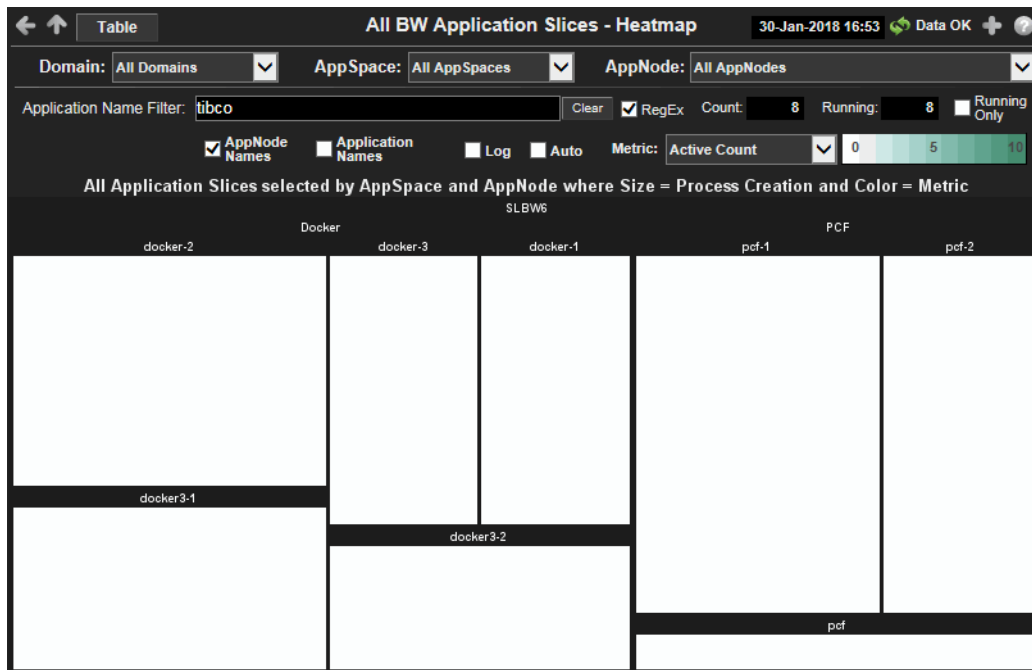
- ["BW All AppSlices Heatmap" on page 966](#): A color-coded heatmap view of process creation and execution metrics.
- ["BW All AppSlices Table" on page 969](#): A tabular view of all available data in this BusinessWorks View.
- ["BW Single AppSlice Summary" on page 971](#): Current and historical metrics for a single AppSlice.

BW All AppSlices Heatmap

View the most critical performance metrics for BusinessWorks AppSlices. Use this display to quickly identify AppSlices with high process execution numbers.

Each rectangle in the heatmap represents an AppSlice. The rectangle color indicates the process execution numbers for the AppSlice. The rectangle size represents the number of processes created in the rectangle; a larger size is a larger value.

Choose a domain, AppSpace and AppNode from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **Application Name Filter** field to limit data shown in the display. Use the **AppNode Names** and **Application Names** check-boxes ☒ to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Active Count**. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the “[BW Single AppSlice Summary](#)” display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.







Filter By:




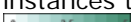
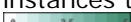
The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.

Fields and Data:

- Application Name Filter** Enter a string to limit data shown in the display.

	Clear	Clears the Application Name Filter entries from the display.
RegEx		Toggles the Search Text field to accept Regular Expressions for filtering.
Count		The number of AppNodes in the display.
Running		The total number of AppSpaces currently running in the display.
Running Only		Select to show only running applications in the display.
AppNode Names		Check to include labels in the heatmap.
Application Names		Check to include labels in the heatmap.
Log		Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto		Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric		Choose a metric to view in the display.
	Active Count	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	Completed Count	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	Suspended Count	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	Failed Count	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	Created / sec	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	Suspended / sec	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Failed / sec	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Exec Time / sec	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Most Recent Exec Time	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Exec Time	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Elapsed Time	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

BW All AppSlices Table

View BusinessWorks data shown in the ["BW All AppSlices Heatmap"](#), and additional details, in a tabular format.

Each row in the table is an AppSlice. Choose a domain (or **All Domains**), an AppSpace (or **All AppSpaces**) and an AppNode (or **All AppNodes**) from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Application Name Filter** field to limit data shown in the display.

Drill-down and investigate by clicking a row to view details in the ["BW Single AppSlice Summary"](#) display.

Heatmap

All BW Application Slices - Table

31-Jan-2018 09:28 Data OK

Domain: All Domains AppSpace: All AppSpaces AppNode: All AppNodes

Application Name Filter: Clear RegEx Count: 8 Running: 8 Running Only

Domain	AppSpace	AppNode	Name	Version	State
SLBW6	Docker	docker-1	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker-2	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker-3	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker3-1	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker3-2	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	PCF	pcf	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	PCF	pcf-1	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	PCF	pcf-2	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

Domain: Choose a domain to show data for in the display.

AppSpace Choose an AppSpace to show data for in the display.

AppNode Choose an AppNode to show data for in the display.

Application Name Filter Enter a string to limit data shown in the display.

Clear Clears the **Application Name Filter** entries from the display.

RegEx Toggles the **Application Name Filter** to accept Regular Expressions for filtering.

Fields and Data:

Count: The total number of rows in the table.

Running The total number of applications currently running in the AppSpace.

Running Only Select to show only running applications in the display.

Table:

Each row in the table is a different AppNode.

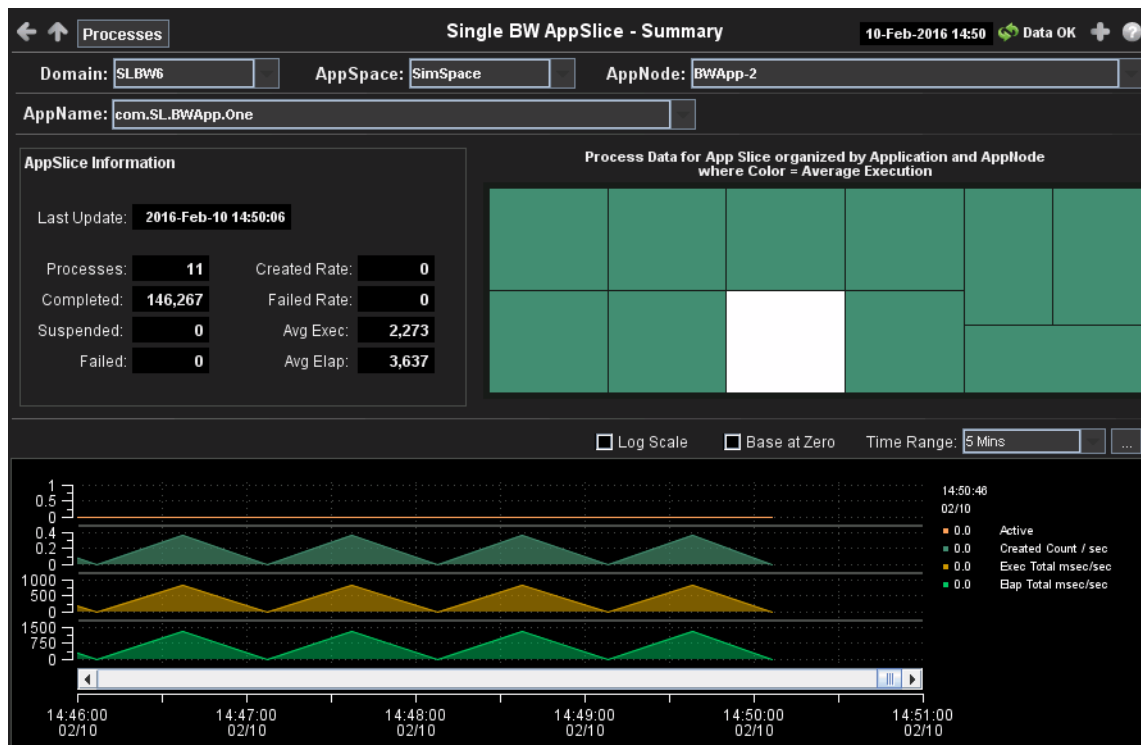
Domain	The domain in which the AppSpace resides.
AppSpace	The AppSpace the AppNode is associated with.
AppNode	The name of the selected AppNode.
Name	The name of the application.
Version	The application version.
State	The current status of the application. Valid values are Running and Stopped .
Module	The application module.
Shared Module	The shared module, if any.
Active Processes	The number of currently active application processes.
Suspended Processes	The number of suspended application processes.
Failed Processes	The number of failed application processes.
Completed Processes	The number of completed application processes.
Created /sec	The number of application processes created per second.
Suspended / sec	The number of application process suspensions per second.
Failed /sec	The number of application process failures per second.
Exec Time / sec	The number of processes executed per second.
Recent Exec Time /sec	The number of seconds for the most recently executed process.
Average Exec Time	The average number of seconds for all processes to execute.
Time Stamp	The date and time the row data was last updated.
Source	Name of RTView Data Server sending this data (or localhost).

BW Single AppSlice Summary

View current and historical utilization and performance metrics for a single BusinessWorks AppSlice. Use this display to investigate performance issues on an AppSlice level.

This display includes a heatmap showing the most critical process execution alerts for AppSlices on the selected application, and trend graphs tracing process execution times.

Choose a domain, AppSpace, AppNode and AppNode from the drop-down menus. Use the **Time-Range** to “zoom-in” or “zoom-out” on a specific time frame in the trend graph. Drill-down and investigate by clicking a process in the heatmap to view details in the “[BW Single Process Summary](#)” display.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.
- AppName:** Choose an AppName to show data for in the display.

Fields and Data:

- AppSlice Information** **Last Update:** The date and time the data was last updated.

Processes	The number of active processes.
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 for processes, in seconds.

Heatmap

Each rectangle in the heatmap represents one process in an AppSlice. The rectangle color indicates the most critical **Average Exec Time** alert state associated with the AppSlice. The rectangle size represents the processes execution time in the rectangle; a larger size is a larger value. Click a rectangle to drill-down and investigate in the ["BW Single Process Summary"](#) display.


Trend Graphs

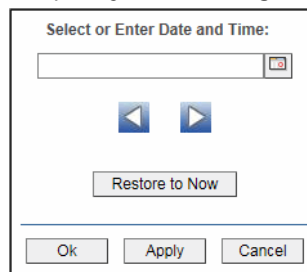
Traces the sum across all processes in one AppSlice of the application.


- **Active:** Traces the number of active processes.
- **Created Count:** Traces the number of processes created.
- **Exec Total msec/sec:** Traces the rate at which the application accumulates process execution time, in milliseconds per second.
- **Elap Total msec/sec:** Traces the rate at which the application is accumulating process elapsed time, in milliseconds per second.



Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

BW Processes

These displays present performance data for BusinessWorks processes. Use these displays to verify that individual BusinessWorks processes are executing and using resources as expected. Displays in this View are:

- [“BW All Processes Heatmap” on page 974](#): A color-coded heatmap view of selected process performance metrics.
- [“BW All Processes Table” on page 977](#): A tabular view of all available process performance data in this BusinessWorks View.
- [“BW Single Process Summary” on page 980](#): Current and historical metrics for a single process.

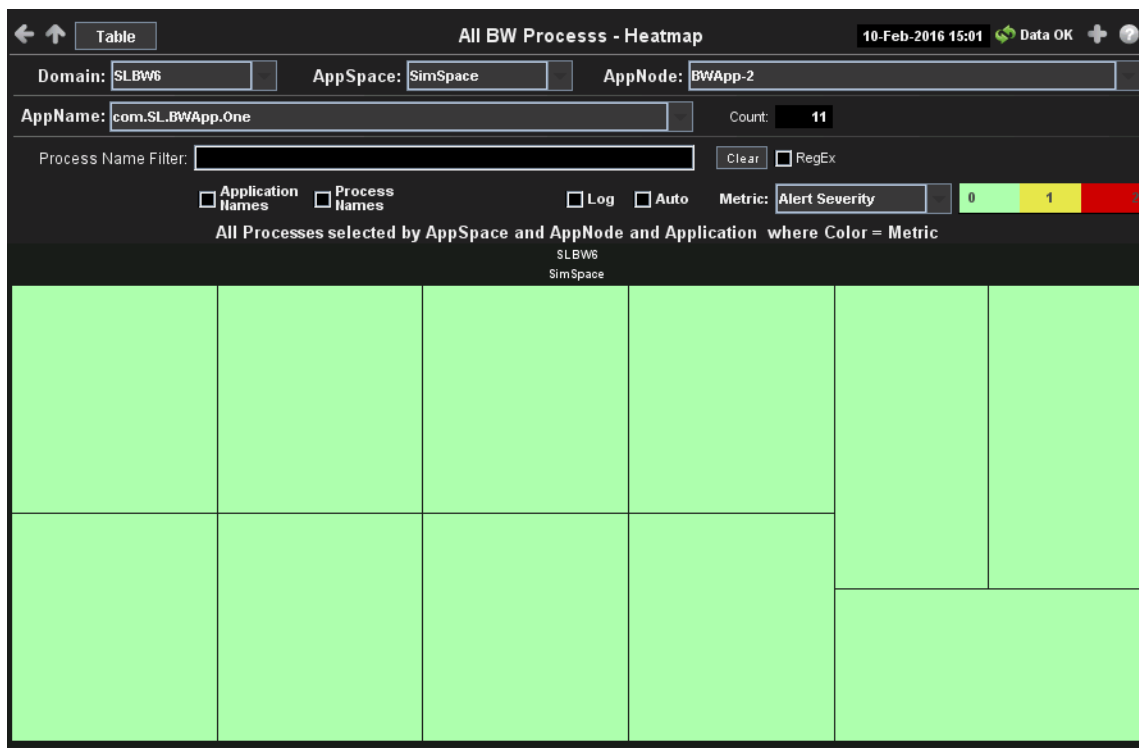
BW All Processes Heatmap


View the most critical BusinessWorks alerts pertaining to process creation and execution. Use this display to quickly identify processes with critical alerts.

Each rectangle in the heatmap represents a process. The rectangle color indicates the most critical alert state associated with the processes (the rectangle size is uniform for all processes.)

Choose a domain, applications, AppNode and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **Process Name Filter** field to limit data shown in the display. Use the **Application Names** and **Process Names** check-boxes ☒ to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate by clicking a rectangle in the heatmap to view details in the [“BW Single Process Summary”](#) display.














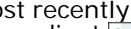
Filter By:

The display might include these filtering options:


- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.
- AppName** Choose an AppName to show data for in the display.
- Count:** The total number of processes currently shown in the display.

Fields and Data:


- Process Name Filter** Enter a string to limit data shown in the display.
 - Clear** Clears the **Processes Name Filter** entries from the display.
- RegEx** Toggles the **Processes Name Filter** to accept Regular Expressions for filtering.
- Application Names** Check to include labels in the heatmap.
- Process Names** Check to include labels in the heatmap.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.
- Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:
 - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.

Alert Count	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
Active Count	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Completed Count	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Suspended Count	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Failed Count	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Created / sec	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Suspended / sec	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Failed / sec	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Exec Time / sec	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Most Recent Exec Time	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.


**Average
Exec Time**

The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Most
Recent
Elapsed
Time**

The elapsed time for the most recent process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Average
Elapsed
Time**

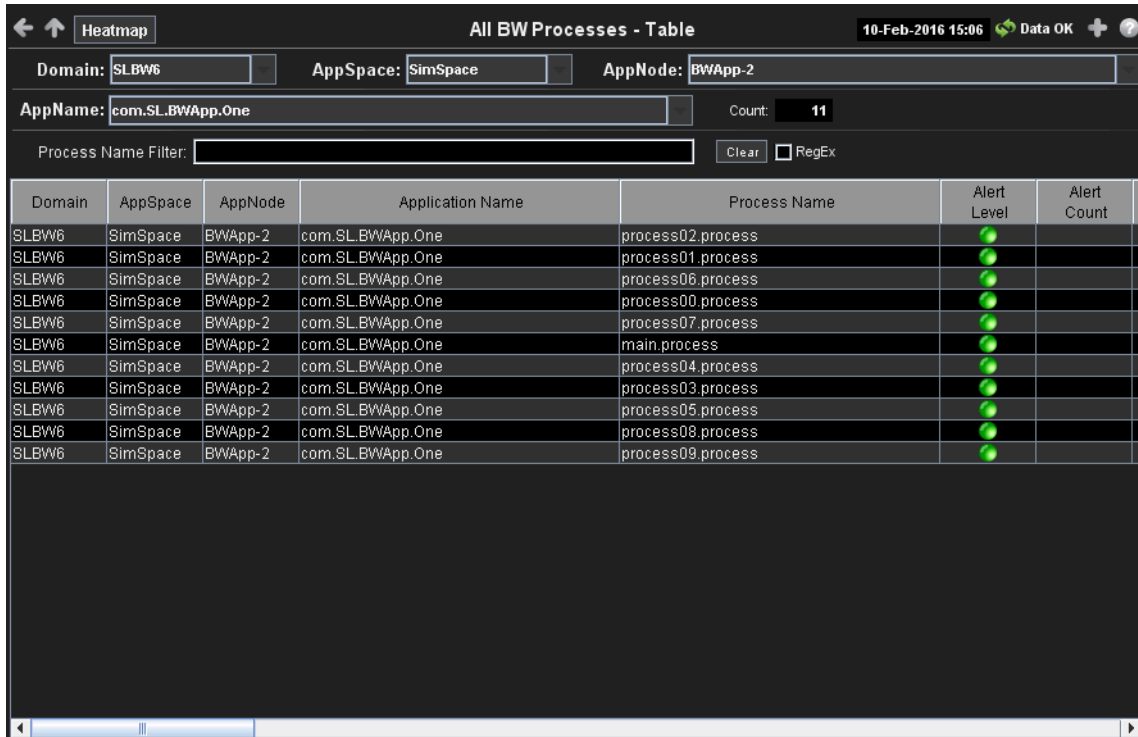
The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

BW All Processes Table


View BusinessWorks data shown in the “[BW All Applications Heatmap](#)”, and additional details, in a tabular format.

Each row in the table is a process. Choose a domain, applications, AppNode and AppSpace from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Process Name Filter** field to limit data shown in the display.

Drill-down and investigate by clicking a row to view details for the selected process in the “[BW Single Process Summary](#)” display



Domain	AppSpace	AppNode	Application Name	Process Name	Alert Level	Alert Count
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process02.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process01.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process06.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process00.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process07.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	main.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process04.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process03.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process05.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process08.process	1	
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process09.process	1	

**Filter By:**

The display might include these filtering options:

Domain:	Choose a domain to show data for in the display.
AppSpace	Choose an AppSpace to show data for in the display.
AppNode:	Choose an AppNode to show data for in the display.
AppName	Choose an AppName to show data for in the display.

Fields and Data:

Count:	The total number of processes in the AppSpace.
Process Name Filter	Enter a string to limit data shown in the display.
Clear	Clears the Application Name Filter entries from the display.
RegEx	Toggles the Application Name Filter to accept Regular Expressions for filtering.

Table:

Each row in the table is a different AppSlice. Column values are associated with the process.

Domain	The domain in which the process resides.
AppSpace	The AppSpace in which the process resides.
AppNode	The AppSpace in which the process resides.
Application Name	The name of the application in which the process is running.
Process Name	The name of the process.
Alert Level	The most critical alert state for alerts in the row: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the process.
Total Exec Time	Total execution time (in milliseconds) for all successfully completed process instances.
Delta Exec Time	Execution time accumulated during the current polling period.

Exec Time/ sec	Delta execution time per second.
Recent Exec Time	Execution time (in milliseconds) of the most recently completed process instance.
Total Elapsed Time	Total elapsed time (in milliseconds) for all successfully completed process instances.
Delta Elapsed Time	Elapsed time accumulated during the current polling period.
Elapsed Time/sec	Delta elapsed time per second.
Recent Elapsed Time	Elapsed clock time (in milliseconds) of the most recently completed process instance.
Active	The number of currently active processes
Created	The number of processes created.
Suspended	The number of process suspensions.
Failed	The number of process failures.
Completed	The number of completed processes.
Delta Active	The number of active processes since the last data update.
Active/sec	The number of active processes per second.
Delta Created	The number of created processes since the last data update.
Created/sec	The number of created processes per second.
Delta Suspended	The number of suspended processes since the last data update.
Suspended/ sec	The number of suspended processes per second.
Delta Completed	The number of completed processes since the last data update.
Completed/ sec	The number of completed processes per second.
Delta Failed	The number of failed processes since the last data update.
Failed/sec	The number of failed processes per second.
Min Exec Time	Execution time (in milliseconds) of the process instance that has completed in the shortest amount of execution time.
Max Exec Time	Execution time (in milliseconds) of the process instance that has completed in the longest amount of execution time.
Average Exec Time	Average execution time (in milliseconds) for all successfully completed process instances.
Min Elapsed Time	Elapsed clock time (in milliseconds) of the process instance that has completed in the shortest amount of elapsed time.
Max Elapsed Time	Elapsed clock time (in milliseconds) of the process instance that has completed in the longest amount of elapsed time.

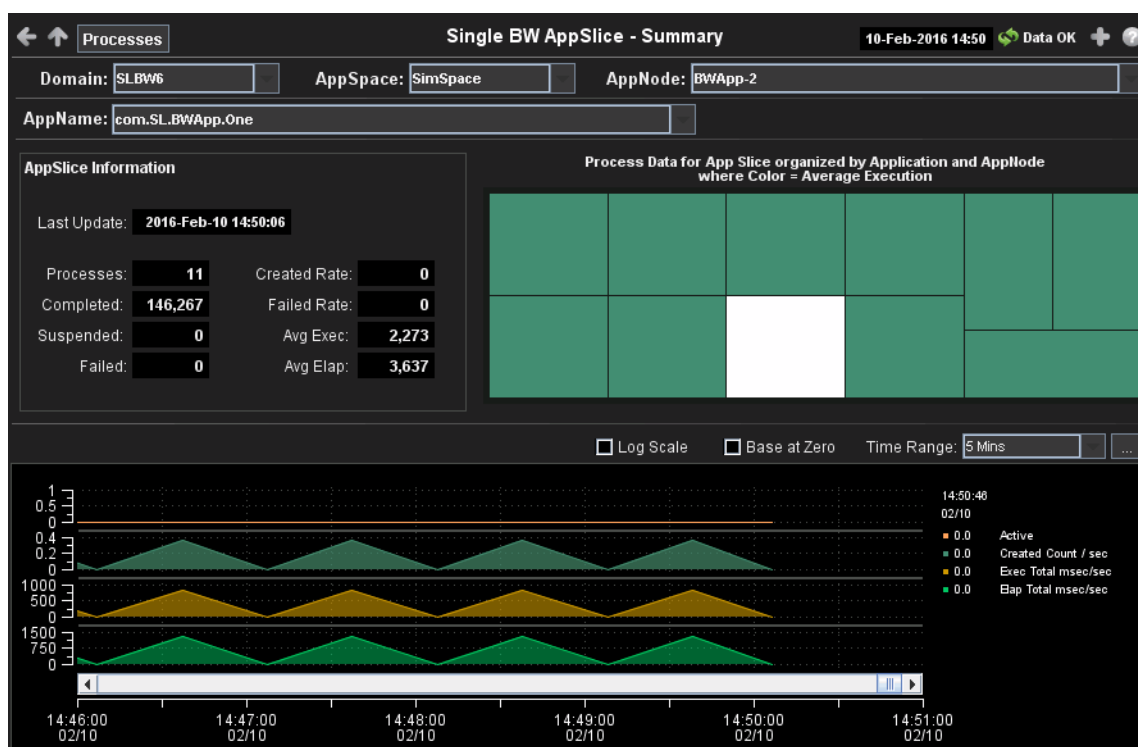
Average Elapsed Time	Average elapsed clock time (in milliseconds) for all successfully completed process instances.
Count Since Reset	The number of times the process has executed since statistics were reset.
Main Process	The name of the main process.
Application Version	The application version.
Module Name	The application module.
Module Version	The module version.
Time Stamp	The date and time the row data was last updated.
Source	Name of RTView Data Server sending this data (or localhost).


BW Single Process Summary

View current and historical execution metrics for a single BusinessWorks process. Use this display to investigate performance issues for a process.

This display includes trend graphs tracing process and activity execution counts and times.

Choose a domain, application, AppNode and AppSpace from the drop-down menus. Use the **Time-Range** to “zoom-in” or “zoom-out” on a specific time frame in the trend graph.



**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.
- AppName** Choose an application to show data for in the display.
- Process** Choose a process to show data for in the display.

Fields and Data:

- Activity Count:** The number of activities defined for the process.
- Main Process:** The name of the main process.
- Active** Number of active instances for this process definition. This number is calculated using the Hawk method named GetProcesses. This method returns information about process instances that are active at the time of update. The value here displays the current total count of all active instances discovered for this process definition. The trend below displays the same value over time.
- Active/sec** The number of currently active application processes per second.
- Created**
 - Total** The number of process instances created for this process definition.
 - Current** The number of process instances created this update cycle.
 - Rate** The number of process instances created per second.
- Completed**
 - Total** The number of process instances completed for this process definition.
 - Current** The number of process instances completed this update cycle.
 - Rate** The number of process instances completed per second.
- Errors**
 - Total** The number of errors accumulated by all process instances.
 - Current** The number of errors accumulated this update cycle.
 - Rate** The number of errors accumulated per second.
- Execution**
 - Min** The shortest execution time of any process instance, in milliseconds.
 - Max** The longest execution time of any process instance, in milliseconds.

Elapsed	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.
	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.

Trend Graphs

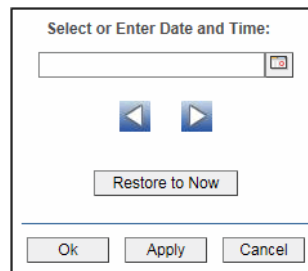
Traces application process and activity metrics for the selected process.

- **Active Count:** Traces the number of currently active processes.
- **Created Count:** Traces the number of created processes.
- **Process Elapsed Time/sec:** Traces the rate at which the application is accumulating process elapsed time, in milliseconds per second.
- **Process Exec Time/sec:** Traces the rate at which the application is accumulating process execution time, in milliseconds per second.
- **All Activities Exec Count/sec:** Traces the number of executed activities per second.
- **All Activities Exec Time/sec:** Traces the amount of execution time for executed activities per second.


Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.



Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.


BW5 Engines

These displays present performance metrics for BW5 Engines. Displays in this View are:

- [“All Engines Heatmap” on page 983](#): Performance metrics of CPU and memory utilization for all BW Engines.
- [“All Engines Table” on page 986](#): Available metrics from the Hawk microagent for each BW Engine.
- [“All Engines Grid” on page 989](#): Displays the main health metrics and a single trend graph per engine, summarizing the status of each BW Engine.
- [“Single Engine Summary” on page 991](#): Detailed performance metrics and alert status for a single BW Engine.

All Engines Heatmap

Quick view of BW5 Engines status for the selected **Filter** and **Server**. Each rectangle in the heatmap represents an engine. Rectangle size represents Max Heap Size and the color represents the most severe value in the heatmap rectangle is shown for the selected Metric. By default, the maximum **Alert Severity** is shown:

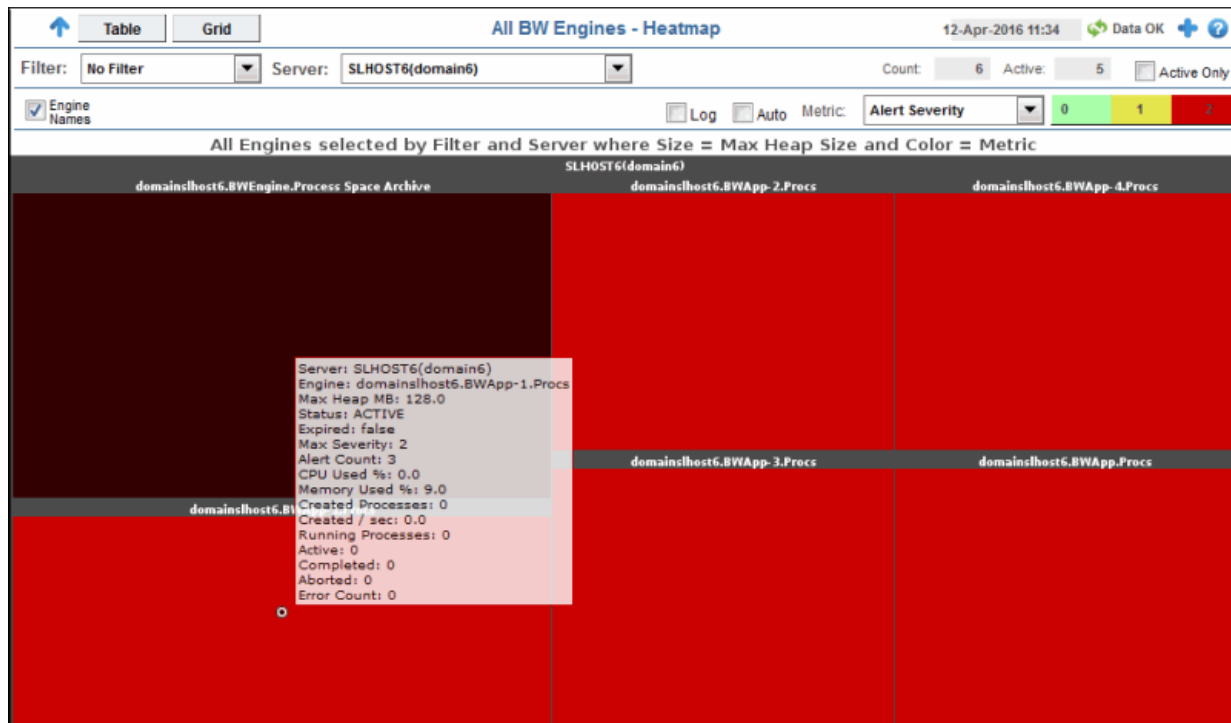
Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Mouseover to see the following performance metrics:

- **Server**: Server agent name.
- **Engine**: Engine name.
- **Max Heap MB**: Maximum heap allocated to this engine for the JVM.
- **Status**: ACTIVE, STOPPED or LIMITED.
- **Expired**: When checked, data has not been received from this host in the specified amount of time.
- **Alert Count**: Number of current alerts
- **CPU Used %**: Percent of server CPU used by engine.
- **Memory Used %**: Percentage of allocated memory currently consumed by this engine from within the JVM. Equal to the value of: (100*UsedBytes) divided by MaxBytes. NOTE: Percent used is Long.
- **Created Processes**: The total number of processes created.
- **Created / sec**: The number of processes created per second.
- **Running Processes**: The number of currently running processes.
- **Active**: The number of currently active processes.
- **Completed**: The total number of completed processes.
- **Aborted**: The total number of aborted processes.
- **Error Count**: The total number of errors.

Click on a node to drill down to the “[Single Engine Summary](#)” display to look at number of processes running, threads, history of memory utilization and other performance metrics for a specific engine. Mouse-over nodes to view details about engine performance and status.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.


Filter By:

The display might include these filtering options:


- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to show data for in the display.
- Count:** The total number of engines in the display.
- Active** Number of engines currently active.
- Active Only** If selected, only engines with a status of ACTIVE are displayed. Otherwise, if deselected, all engines for the given Filter/Server selection are displayed.
- Engine Names** Select this check box to display the names of the engines above their respective rectangles in the heatmap.

Log	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
CPU Used%	The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Memory Used%	The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Active Processes	The number of currently active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Running Processes	The number of currently running processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Created Processes	The number of created processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Created/
sec**

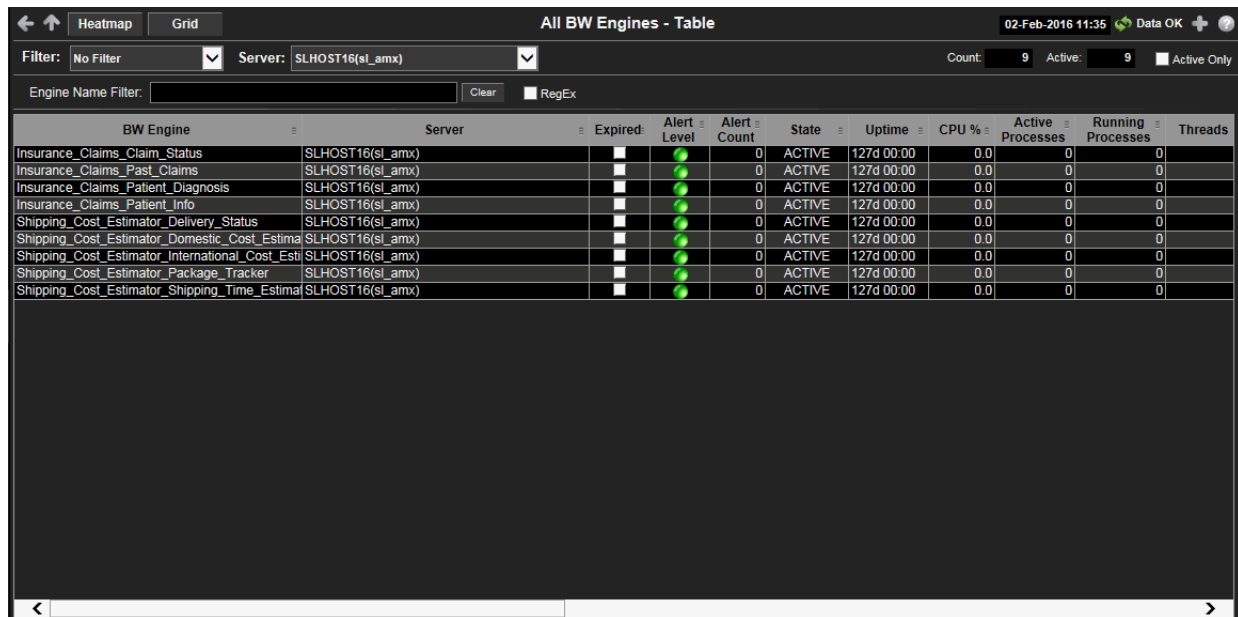
The number of created processes in the heatmap rectangle, per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Error Count

The total number of errors in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

All Engines Table

Each row in the table is an engine. Metrics are made available by the Hawk microagent for the engine (for details, refer to documentation for TIBCO BusinessWorks Administration, Appendix A: TIBCO Hawk Microagent Methods). Click on an row to drill down to the “[Single Engine Summary](#)” display.



BW Engine	Server	Expired	Alert Level	Alert Count	State	Uptime	CPU %	Active Processes	Running Processes	Threads
Insurance_Claims_Claim_Status	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Insurance_Claims_Past_Claims	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Insurance_Claims_Patient_Diagnosis	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Insurance_Claims_Patient_Info	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Delivery_Status	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Domestic_Cost_Estima	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_International_Cost_Esti	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Package_Tracker	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Shipping_Time_Estima	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:




The display might include these filtering options:

Filter:

Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.

Server:	Choose a server to show data for in the display.
Count	Number of engines currently being displayed.
Active	Number of engines currently active.
Active Only	If selected, only engines with a status of ACTIVE are displayed. Otherwise, if deselected, all engines for the given Filter/Server selection are displayed.
Engine Name Filter	Enter all or part of engine name to view specific engines. NOTE: Wild card characters are supported.
Clear	Removes Engine Name Filter and all engines for the given Filter/Server selection are displayed.
RegEx	If selected, the specified Engine Name Filter will be interpreted as a full Regular Expression rather than a simple wildcard.

Table:

BW Engine	BW Engine name.
Server	Server agent name.
Expired	When checked, data has not been received from this host in the specified amount of time.
Alert Level	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Number of current alerts
State	Engine status: ACTIVE, STOPPED, LIMITED, etc. (See All Servers Grid).
Uptime	Uptime in milliseconds since the engine was started.
CPU %	Percent of server CPU used by engine.
Active Processes	Number of active processes calculated each update period using data returned by the Hawk method GetProcesses. Note: This column will display NaN or Not Available for any engine whose status is STOPPED .
Running Processes	Number of running processes.
Threads	Number of threads used by the engine.
Memory Used%	Percentage of allocated memory currently consumed by this engine from within the JVM. Equal to the value of: (100*UsedBytes) divided by MaxBytes. NOTE: Percent used is Long.
Max Heap Size	Maximum heap allocated to this engine for the JVM.
Total Bytes	Maximum heap memory this JVM has used.
Used Bytes	Total bytes of memory within the JVM currently used by the engine. Equal to value of: MaxBytes minus FreeBytes.
Free Bytes	Amount of available memory from within the JVM.
Mem Usage KBytes	Server memory in KB used by engine.

Errors	Total number of errors since the engine was started.
Delta Errors	Current number of new errors.
Errors/sec	Error rate per second.
Created Processes	The total number of processes that were created.
Completed Processes	The total number of processes that were completed.
Aborted Processes	The total number of processes that were aborted.
Process ID	Process ID of engine as recognized by the server.
Micro Agent Instance	Unique ID of the microagent reporting the metrics.
Deployment	Name of Deployment.
Domain	Name of Domain.
BW Version	The TIBCO BusinessWorks version currently in use on the server.
Source	Name of RTView Data Server sending this data (or localhost).
Time Stamp	Time of last update.
Process Name	Name of the BW Engine process on the server. Note: This information is not displayed in the table but is present in "raw" cache data.
Host	Host name of server. Note: This information is not displayed in the table but is present in "raw" cache data.
Adapter Name	Name of adapter. Note: This information is not displayed in the table but is present in "raw" cache data.
Instance ID	Instance ID name of the engine. Note: This information is not displayed in the table but is present in "raw" cache data.
Version	Engine project version number. Note: This information is not displayed in the table but is present in "raw" cache data.

All Engines Grid

Displays the main health metrics and a single trend graph per engine, summarizing the status of each BW5 Engine. Click on an engine icon to drill down to the [“Single Engine Summary”](#) display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to show data for in the display.
- Count** Number of engines currently being displayed.
- Active** Number of engines currently active.
- Active Only** Toggle this setting to display active servers or all servers.
- Time Range** Choose a time range. Also sets range for the **Single Engine Summary** display. Options are:
All Data, 2 Mins, 5 Mins, 20 Mins, 1 Hour, 2 Hours, 4 Hours, 8 Hours, 24 Hours, 2 Days and 7 Days.

Fields and Data

Engine Name	Name of the engine.
Status	Indicates the current state of the engine: <ul style="list-style-type: none"> • ACTIVE Indicates the BW microagent is providing live data and the engine is assumed active. • SUSPENDED This state is reported by the BW microagent. • STANDBY This state is reported by the BW microagent. • STOPPING This state is reported by the BW microagent. • STOPPED This state is reported by the BW microagent. • LIMITED Live data has been received from TIBCO, but deployment data from the custom RTView microagent has not been received. • EXPIRED Indicates the server associated with the engine is unavailable or stopped sending data.
CPU Usage%	Percent of server CPU in use.
Mem Usage%	Available physical memory (MB) remaining.
Error Rate	Number of errors accumulated per second.
Total Processes	Number of process definitions for this engine.
Active Processes	Number of process instances currently active.
Trend Graphs	Traces data for the server.
CPU	Traces percent of server CPU in use.
MEM	Traces available physical memory remaining.
PROCS	Traces total number of active processes.

Single Engine Summary

Several views show historical and current performance metrics for a single engine, including the number of processes running, threads, history of memory utilization, and trend graphs of memory utilization. In this display, when an engine is **Stopped** the engine name is appended with **(X)**, the background color is light red and Uptime is zero.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Clicking the **Processes** button in the Title Bar takes you to the [“All Processes Heatmap”](#). Clicking the **JVM** button, which is automatically enabled when a JMX connection is defined for the engine, takes you to the **JVM CPU/Mem Summary** display. See **Enable Monitoring Via JMX** for more information on enabling a JMX connection.

Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to show data for in the display.
- Engine:** Choose an engine to show data for in the display. An engine is not running when the engine name is appended with **(X)**.

Fields and Data

Engine Information

Status	ACTIVE	The BW microagent is providing live data and the engine is assumed active.
	SUSPENDED	This state is reported by the BW microagent.
	STANDBY	This state is reported by the BW microagent.
	STOPPING	This state is reported by the BW microagent.
	STOPPED	This state is reported by the BW microagent.
	LIMITED	Live data has been received from TIBCO, but deployment data from the custom RTView MicroAgent has not been received.
	EXPIRED	The associated server for the engine is currently in an EXPIRED state and is unavailable or stopped sending data.
Uptime	Days hours and minutes since the engine was started.	
CPU%	Percent of server CPU used by engine.	
Memory %	Available physical memory remaining (in MB).	
Threads	Number of threads used by this engine	
Max Heap MB	Maximum heap allocated to this engine for the JVM.	

Error Counts

Total	Total errors accumulated by this engine.
Current	Number of errors accumulated this update cycle.

Process Counts

Total	<p>A BW Engine runs processes by creating instances of process definitions and making them active. A given process instance has a lifetime during which it may be suspended, swapped, queued, etc. until it is either completed or aborted.</p> <p>The Total value is calculated using the Hawk method named GetProcessDefinitions that returns statistics about the instances of each process definition including cumulative counts of completed, aborted, suspended, etc.</p>
Running	Total number of running process instances. This number is calculated using the Hawk method named GetProcessCount. It is displayed in the Monitor Engines Table as RunningProcesses. The trend below displays the same value over time as Running Procs.

Alerts

Indicates the greatest severity level and the number of open **Engine**, **Process**, and **Activity** alerts for the selected engine. Values range from **0** to **2**, where **2** is the greatest Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

Click on the alert indicator to display the **BW Engine - Tables** display, which contains a table listing the current alerts for the selected engine.

The screenshot shows the 'BW Engine - Tables' window. At the top, there's a 'Summary' tab and a status bar with '16-Jul-2018 14:51' and 'Data OK'. Below the status bar, there are dropdown menus for 'Filter' (set to 'No Filter'), 'Server' (set to 'SLHOST16(sl_qa_conn)'), and 'Engine' (set to 'sl_qa.ADB_Operations_JMS_Standard_M').

The main content area contains three tables:

- BW Engine Table:** A table with columns: State, CPU %, Running Processes, Threads, Memory Used %, Max Heap Size, Total Bytes, Used Bytes, Free Bytes, and Mem Use. The first row shows 'LIMITED', '0.0', '0', '8', '0.0', '0', '0', '0', '0', and '0'.
- Process Totals Table:** A table with columns: Aborted, Active, AverageElapsed, AverageExecution, Checkpointed, Completed, Created, and ProcessCount. The first row shows '1107306', '0', '0.0', '0.0', '0', '664628', '1771934', and '6'.
- Current Alerts for Selected BW Engine:** A table with columns: Time, Alert Name, Alert Index, and Alert Text. The table is currently empty.

- Engine** Number of engine alerts and the most critical alert state for the engine:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.
- Process** Number of process alerts and the most critical alert state for the engine:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.
- Activity** Number of activity alerts and the most critical alert state for the engine:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.

Heatmap

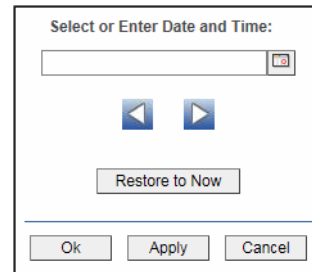
Shows processes organized by Server/Engine where Size = Creation Count and Color = Average Execution. Click on a node to drill down to a specific engine.


Trend Graphs



Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Base at Zero Select to use zero (0) as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

BW5 Processes

These displays present performance metrics for BW5 processes. Displays in this View are:

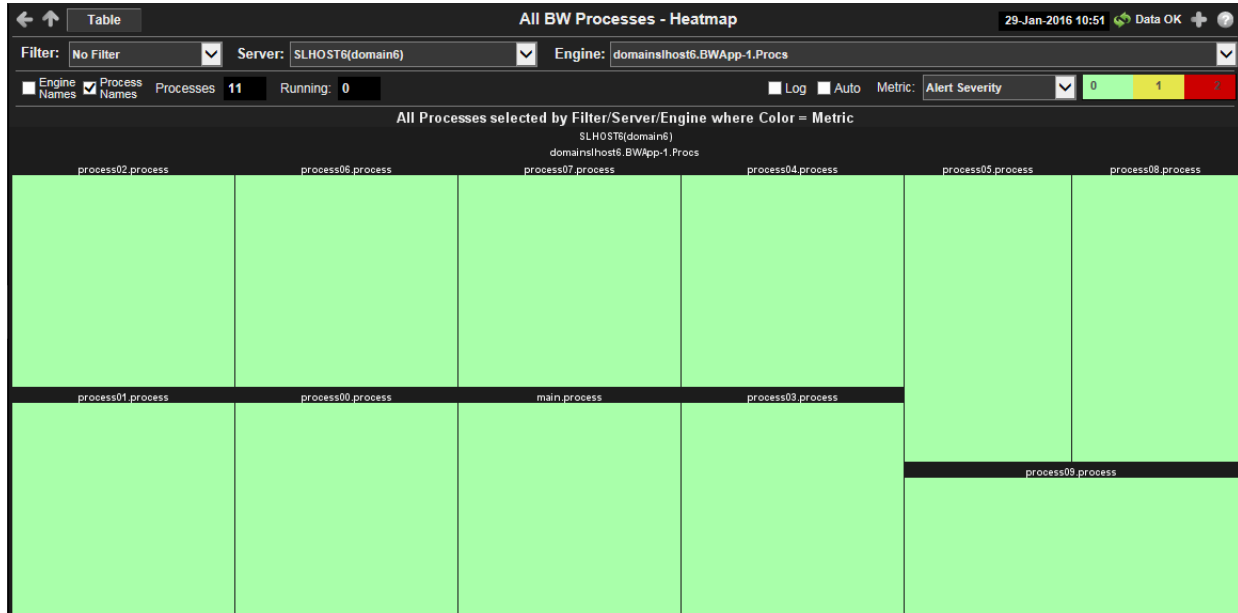
- [“All Processes Heatmap” on page 994](#): Displays process execution metrics for all BW Engines.
- [“All Processes Table” on page 998](#): Each row in the table displays all available metrics from the Hawk microagent for a process.
- [“Single Process Summary” on page 1001](#): Several views show historical and current metrics for a single process, including average execution times and execution counts.

All Processes Heatmap

Summary view of processes can show the execution times for all processes on all engines or you can filter to look at specific servers or engines. Each rectangle (node) in the heatmap represents a process. Move your mouse over a node to display current metrics. Click on a node to drill-down to the [“Single Process Summary”](#) display to view specific metrics about process behavior over a specified period of time and determine which activity in the process may be causing the bottleneck.

An engine is not running when the engine name is appended with **(X)**.

Mouse-over any node to display the current values for the metric selected from the **Metric** drop-down menu.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

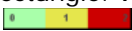



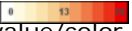






23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.




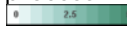


Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to show data for in the display.
- Engine:** Choose an engine to show data for in the display. An engine is not running when the engine name is appended with **(X)**.
- Engine Names** Select this check box to display the names of the engines above their respective rectangles in the heatmap.
- Process Names** Select this check box to display the names of the processes above their respective rectangles in the heatmap.
- Processes** The total number of processes in the display.
- Running** Number of processes currently running.

Log	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
Completed Count	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Active Count	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Aborted Count	The total number of aborted processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Suspended Count	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Exec Time / sec	The number of processes executed per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Created / sec	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Aborted / sec	The number of aborted processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Suspended / sec	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Most Recent Exec Time	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Exec Time	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Most Recent Elapsed Time	The elapsed time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Average Elapsed Time	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

All Processes Table

Select a server and engine from the drop-down menus. Each row in the table is a different engine. The table displays all metrics available from the Hawk microagent for an engine. (Refer to documentation for TIBCO BusinessWorks Administration, see Appendix A: TIBCO Hawk Microagent Methods).

Click on a row in the table to drill down to the “Single Engine Summary” display.

Heatmap

All BW Processes - Table

29-Jan-2016 09:14 Data OK

Filter: No Filter

Server: SLHOST6(domain6)

Engine: domainshost6.BWApp-1.Procs

BW Engine	Server	BW Process	Expired	Alert Level	Alert Count	Time Since Last Update	Total Exec Time	Exe
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main_process			2	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process00_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process01_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process02_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process03_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process04_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process05_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process06_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process07_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process08_process			0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process09_process			0	0	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.

6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Filter By:
The display might include these filtering options:

- Filter:

Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Create Customized Filters** for more information.
- Server:

Choose a server to show data for in the display.
- Engine:




Choose an engine to show data for in the display. An engine is not running when the engine name is appended with (X).

Table:

- BW Engine

BW Engine name.
- Server

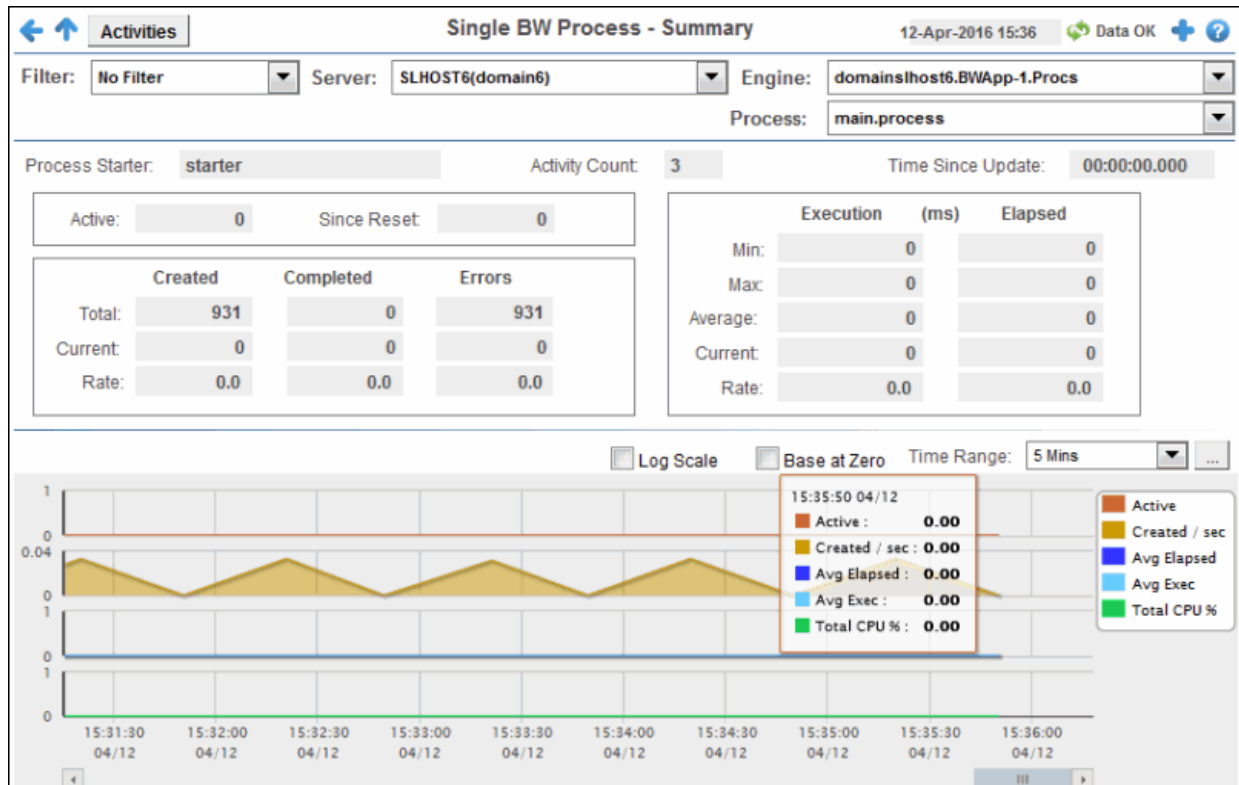
Server agent name.

BW Process	The name of the process.
Expired	When checked, data has not been received from this host in the specified amount of time.
Alert Level	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Number of current alerts
Active	Number of active processes.
Total CPU	Total CPU usage in percent.
Created/sec	Change in Created per second.
Completed/sec	Change in Completed per second.
Delta Created	Change in Created this update.
Delta Completed	Change in Completed this update.
Created	Number of process instances created for this process definition.
Completed	Number of process instances successfully completed.
Total Exec Time	Total execution time (in milliseconds) for all successfully completed process instances.
Delta Exec Time	Execution time accumulated during the current polling period.
Exec Time/sec	Delta execution time per second.
Min Exec Time	Execution time (in milliseconds) of the process instance that has completed in the shortest amount of execution time.
Max Exec Time	Execution time (in milliseconds) of the process instance that has completed in the longest amount of execution time.
Average Exec Time	Average execution time (in milliseconds) for all successfully completed process instances.
Recent Exec Time	Execution time (in milliseconds) of the most recently completed process instance.
Total Elapsed Time	Total elapsed time (in milliseconds) for all successfully completed process instances.
Delta Elapsed Time	Elapsed time accumulated during the current polling period.
Elapsed Time/sec	Delta elapsed time per second.
Min Elapsed Time	Elapsed clock time (in milliseconds) of the process instance that has completed in the shortest amount of elapsed time.
Max Elapsed Time	Elapsed clock time (in milliseconds) of the process instance that has completed in the longest amount of elapsed time.
Average Elapsed Time	Average elapsed clock time (in milliseconds) for all successfully completed process instances.

Recent Elapsed Time	Elapsed clock time (in milliseconds) of the most recently completed process instance.
Aborted	Number of times process instances have been aborted.
Delta Aborted	Change in Aborted this update.
Aborted/sec	Change in Aborted per second.
Queued	Number of times process instances have been queued for execution.
Delta Queued	Change in Queued this update.
Queued/sec	Change in Queued per second.
Suspended	Number of times process instances have been suspended.
Delta Suspended	Change in Suspended this update.
Suspended/sec	Change in Suspended per second.
Checkpointed	Number of times process instances have executed a checkpoint.
Delta Checkpointed	Change in Checkpointed this update.
Checkpointed/sec	Change in Checkpointed per second.
Swapped	Number of times process instances have been swapped to disk.
Delta Swapped	Change in Swapped this update.
Swapped/sec	Change in Swapped per second.
Time Since Last Update	Time since the last update.
Domain	Name of TIBCO Domain.
Starter	Name of the process starter for the process.
MicroAgent Instance	Unique ID of the microagent reporting the metrics.
CountSince Reset	Number of process instances that have completed since the last reset of the statistics.
Source	Name of RTView Data Server sending this data (or localhost).
Time Stamp	Time of last update.

Single Process Summary

Detailed performance metrics and alert status for a single BW process. Select a server, engine and process from the drop-down menus. The background color of the display is red when the selected engine is stopped.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to see metrics for.
- Engine:** Choose a server to see metrics for. An engine is not running when the engine name is appended with (X).
- Process:** Choose a process to see metrics for.

Process Starter	Name of the process starter for the process.
Activity Count	Number of activities defined for this process.
Time Since Update	Time since the last update to file of statistics.
Active	Number of active instances for this process definition. This number is calculated using the Hawk method named GetProcesses. This method returns information about process instances that are active at the time of update. The value here displays the current total count of all active instances discovered for this process definition. The trend below displays the same value over time.
Since Reset	Number of activity executions that have completed since the last reset of the statistics. This is the number retrieved from the Hawk method named GetProcessDefinition which returns ExecutionCountSinceReset.

Execution Counts

Most recent execution counts for this process.

Created	Total	Number of process instances created for this process definition.
	Current	Number of process instances created this update cycle.
	Rate	Number of process instances created per second.
Completed	Total	Number of process instances successfully completed.
	Current	Number of process instances successfully completed this update cycle.
	Rate	Number of process instances successfully completed per second.
Errors	Total	Number of errors accumulated by all process instances.
	Current	Number of errors accumulated this update cycle.
	Rate	Number of errors accumulated per second.


Execution (ms) Elapsed

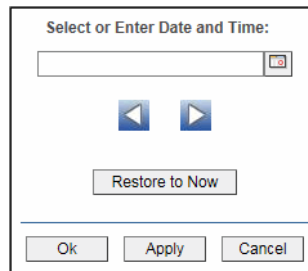
Execution and elapsed times in milliseconds for this process.

Min	Shortest time of any process instance.
Max	Longest time of any process instance.
Average	Average time across all successfully completed process instances.
Current	Time accumulated this update cycle.
Rate	Time accumulated per second.


Trend Graphs



- **Active:** Traces the number of currently active processes.
- **Created / sec:** Traces the number of created processes per second.
- **Avg Elapsed:** Traces the average number of elapsed processes. This value is calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval.
- **Avg Exec:** Traces the average number of executed processes. This value is calculated by dividing the delta executed time for the interval by the delta completed, or the number of process instances that completed in the interval.
- **Total CPU %:** Traces CPU utilization by processes, in percent.

- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

BW5 Activities

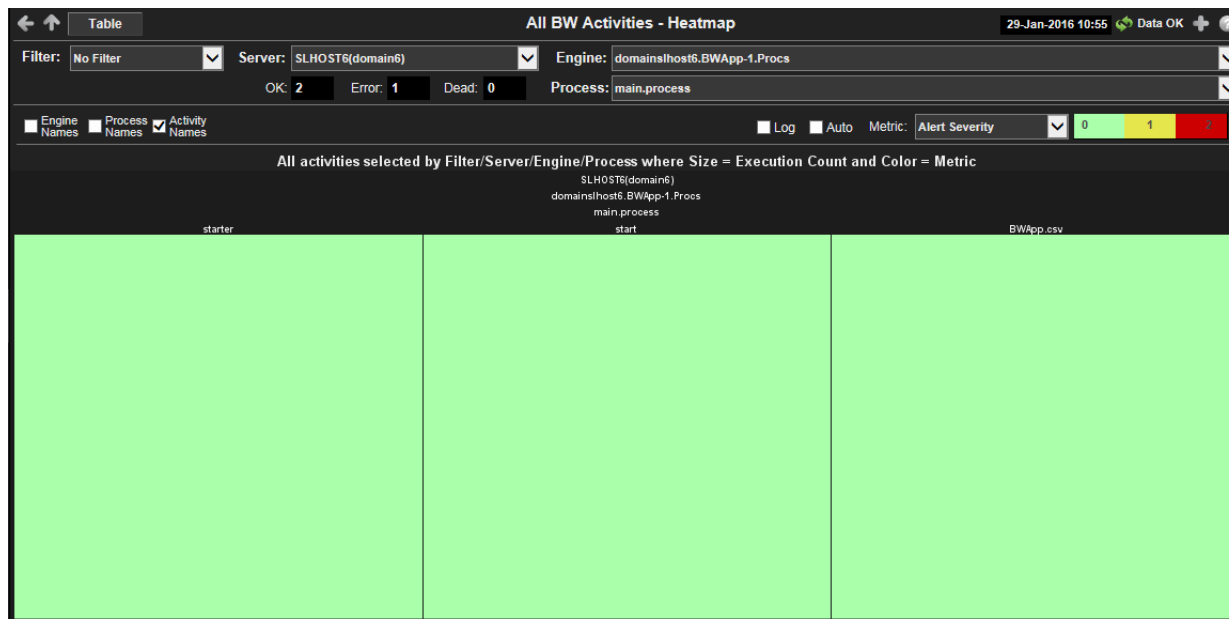
These displays present performance metrics for BW5 activities. Displays in this View are:

- [“All Activities Heatmap” on page 1003](#): Displays execution performance metrics for all BW activities.
- [“All Activities Table” on page 1006](#): Each row in the table displays all available metrics from the Hawk microagent for an activity.
- [“Single Activity Summary” on page 1009](#): Historical and current performance metrics for a single activity, including average execution times and execution counts.

All Activities Heatmap

Summary view of activities shows the execution times for all activities on all engines, or you can filter to look at specific servers, engines or processes. An engine is not running when the engine name is appended with **(X)**.

Move your mouse over a node to display current metrics. Click on a node to drill down to the “[Single Activity Summary](#)” display to view specific metrics about activity behavior over a specified period of time.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.










23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Alert Views - RTView Alerts Table display.


Filter By:

The display might include these filtering options:


Filter:	Choose a filter to show data for in the display. By default, the Filter: drop-down menu only contains the No Filter option. To create your own filtering options, see Creating Customized Filters in the User's Guide.
Server:	Choose a server to show data for in the display.
Engine:	Choose an engine to show data for in the display. An engine is not running when the engine name is appended with (X) .
Process	Select from the menu to view activities running on a specific process or all processes.
OK	Number of activities that reported their Last Return Code as OK .
Error	Number of activities that reported their Last Return Code as Error .
Dead	Number of activities that reported their Last Return Code as Dead .
Engine Names	Select this check box to display the names of the engines above their respective rectangles in the heatmap.

Process Names	Select this check box to display the names of the processes above their respective rectangles in the heatmap.
Activity Names	Select this check box to display the names of the activities above their respective rectangles in the heatmap.
Log	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Auto	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Choose a metric to view in the display.
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
Exec Count	The total number of executed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Error Count	The total number of errors in the heatmap rectangle. The color gradient  bar populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Exec Time / sec	The number of processes executed per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
Errors / sec	The number of errors per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Most Recent Exec Time

The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Max Exec Time

The maximum execution time for executed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.




All Activities Table

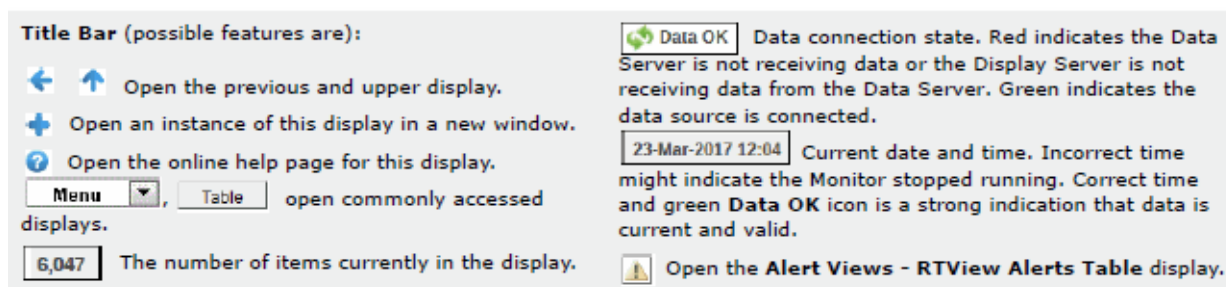
Select a server, engine and process from the drop-down menus. Each row in the table displays all metrics available from the Hawk microagent for an activity. (Refer to documentation for TIBCO BusinessWorks Administration, see Appendix A: TIBCO Hawk Microagent Methods).

Click on a row in the table to drill down to the “[Single Activity Summary](#)” display to view specific metrics about activity behavior over a specified period of time.

When the background/foreground color of a row changes color, the associated engine for the activity is currently in an EXPIRED state. An engine is EXPIRED when the associated server is unavailable or stopped sending data.

An EXPIRED activity and the associated engine are deleted from displays when the associated server exceeds its specified threshold. Processes associated with the engine are also deleted from displays.

All BW Activities - Table									
<div> Heatmap 09-Feb-2016 12:56 Data OK </div>									
<div> Filter: No Filter Server: SLHOST6(domain6) Engine: domainshost6.BWApp-1.Procs Process: main.process </div>									
BW Engine	Server	BW Process	Activity	Expired	Alert Level	Time Since Last Update	Last Ret. Code	Execution Time	Def
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main.process	starter	<input type="checkbox"/>		9,469	OK	61	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main.process	start	<input type="checkbox"/>		9,469	OK	387	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main.process	BWApp.csv	<input type="checkbox"/>		9,469	ERROR	16	

**Filter By:**

The display might include these filtering options:

Filter:	Choose a filter to show data for in the display. By default, the Filter: drop-down menu only contains the No Filter option. To create your own filtering options, see Creating Customized Filters in the User's Guide.
Server:	Choose a server to show data for in the display.
Engine:	Select from the menu to view activities running on a specific engine or all engines. An engine is not running when the engine name is appended with (X).
Process:	Select from the menu to view activities running on a specific process or all processes.

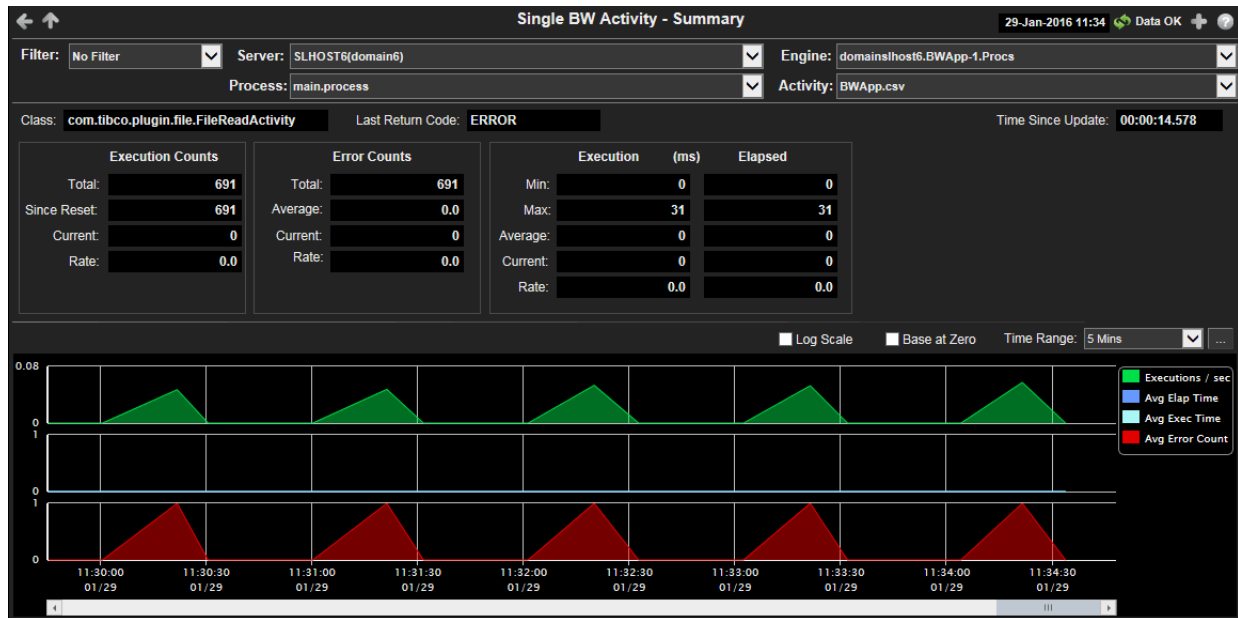
Table:

BW Engine	Name of BW Engine.
Server	Name of Server agent.
BW Process	Name of the BW engine Process on the Server.
Activity	Name of activity.
Expired	When checked, data has not been received from this host in the specified amount of time.
Alert Level	The most critical alert state for alerts in the row: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Time Since Last Update	Time since the last update.
Last Ret(urn) Code	Status code (OK DEAD ERROR) returned by most recent execution of this activity.
Execution Time	Time (in milliseconds) used by all executions of this activity. NOTE: This does not include wait time for Sleep, Call Process, and Wait For... activities.
Delta Exec(ution) Time	Execution time accumulated this update cycle.
Exec(ution) Time / sec	Execution time accumulated per second.

Min Exec(ution) Time	Time (in milliseconds) of the activity that has the shortest execution time.
Max Exec(ution) Time	Time (in milliseconds) of the activity that has the longest execution time.
Elapsed Time	Elapsed clock time (in milliseconds) used by all executions of this activity. NOTE: This does not include wait time for Sleep, Call Process, and Wait For... activities.
Delta Elapsed Time	Change in ElapsedTime this update.
Elapsed Time/sec	Change in ElapsedTime per second.
Min Elapsed Time	Elapsed clock time (in milliseconds) of the activity that has the shortest execution time.
Max Elapsed Time	Elapsed clock time (in milliseconds) of the activity that has the longest execution time.
Executions	Number of times the activity has been executed.
Delta Exec(ution)	Change in ExecutionCount this update.
Executions/sec	Change in ExecutionCount per second.
Errors	Total number of executions of the activity that have returned an error.
Delta Errors	Change in ErrorCount this update.
Errors/sec	Change in ErrorCount per second.
Domain	Name of TIBCO Domain.
ActivityClass	Name of the class that implements the activity.
CalledProcessDefs	A comma-separated list of definitions called by this activity.
Tracing	<ul style="list-style-type: none"> • true Tracing is enabled for this activity. • false Tracing is disabled for this activity.
MicroAgentInstance	Unique ID of the microagent reporting the metrics.
ExecutionCountSince Reset	Number of times the activity has been executed since the last reset of the statistics.
Source	Name of RTView Data Server sending this data (or localhost).
Time Stamp	Time of this update.

Single Activity Summary

Detailed performance metrics and alert status for a single BW activity. In this display, when an engine associated with the activity is **Stopped** the engine name is appended with **(X)** and the background color is light red.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Select from the menu to view processes running on a specific server.
- Engine:** Select from the menu to view processes running on a specific engine. An engine is not running when the engine name is appended with **(X)**.
- Process:** Select from the menu to view summary details for a specific process.
- Activity** Select from the menu to view summary details for a specific activity.
- Class** Name of the activity class.
- Last Return Code** Last return code reported from this activity.

Time Since Update Time since the last update.

Execution Counts

Most recent execution counts for this activity.

Total	Number of times the activity has been executed.
Since Reset	Number of times the activity has been executed since the last Hawk reset of the statistics.
Current	Change in ExecutionCount this update.
Rate	Change in Execution Count per second.

Error Counts

Most recent error counts for this activity.

Total	Number of errors accumulated by all activities.
Average	Average number of errors accumulated by all activities.
Current	Number of errors accumulated this update cycle.
Rate	Number of errors accumulated per second.

Execution (ms) Elapsed

Execution and elapsed times in milliseconds for this activity.

Min	Shortest time of any activity instance.
Max	Longest time of any activity instance.
Average	Average time across all successfully completed activity instance.
Current	Time accumulated this update cycle.
Rate	Time accumulated per second.


Trend Graphs

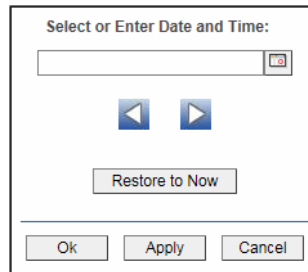
Log Scale Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Base at Zero



Select to use zero (0) as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

BW5 Servers

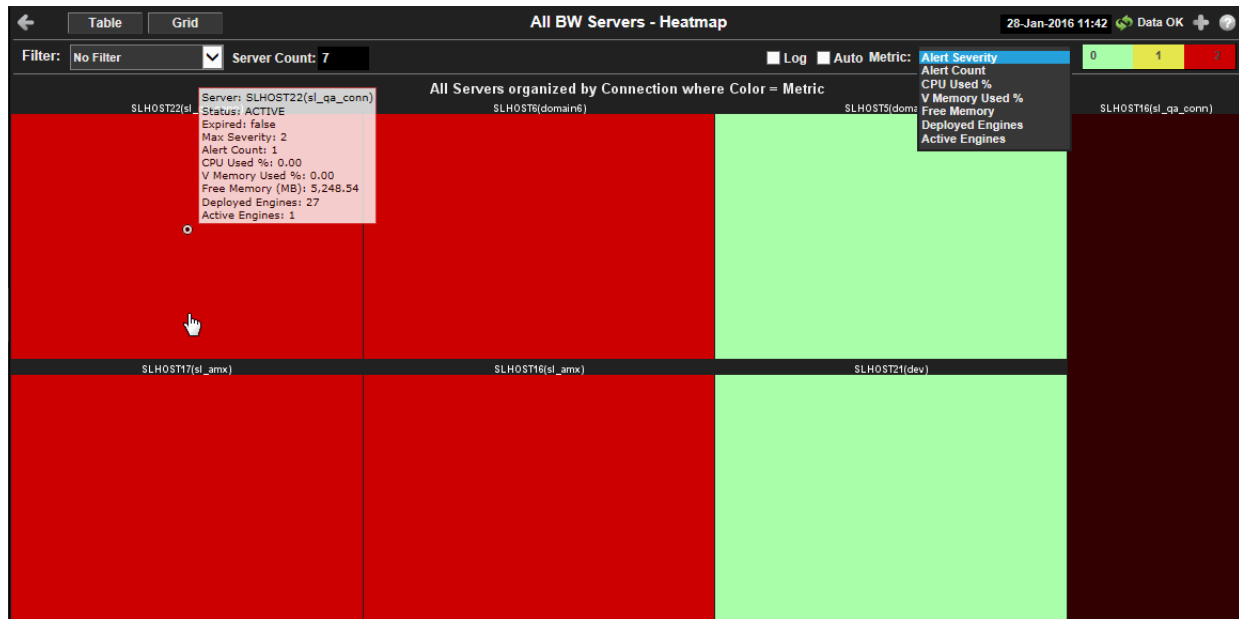
These displays present performance data for your BusinessWorks system. Displays in this View are:

- [“All Servers Heatmap” on page 1011](#)
- [“All Servers Table” on page 1014](#)
- [“All Servers Grid” on page 1015](#)
- [“Single Server Summary” on page 1017](#)
- [“Server Processes” on page 1019](#)
- [“Single Server Process - Summary” on page 1020](#)

All Servers Heatmap

Quick view of BW Servers status determined by selected Filter, organized by Connection (host) and where color equals the selected Metric. Each rectangle (node) in the heatmap represents a server.

Click on a node to drill down to the “[Single Server Summary](#)” display and view metrics for a particular server. Mouse-over any node to display the current values for the metric selected from the Metric drop-down menu.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.











23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to limit data shown in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server Count:** The total number of servers in the display.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
CPU Used%	<p>The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
V(irtual) Memory Used%	<p>The percent (%) virtual memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Free Memory	<p>The amount of free memory in the heatmap rectangle, in megabytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Deployed Engines	<p>The number of deployed engines in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Active Engines	<p>The number of active engines in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>

All Servers Table

This table provides a list view of utilization metrics for all BW servers (represented in the All Servers Heatmap). Each row in the table contains data for a particular server. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill down to the “Single Server Summary” display and view metrics for that particular server.

Server	Expired	Alert Level	State	CPU Usage (%)	Free Memory (MB)	V. Memory Usage (%)	BW Version	Deployed Engines	Active Engines	Source	Time Stamp
SLHOST16(sl_amx)	<input type="checkbox"/>		ACTIVE	5.95	926.28	18.97		9	9	localhost	01/28/16 11:48:30
SLHOST16(sl_qa_conn)	<input checked="" type="checkbox"/>		EXPIRED	10.74	916.28	19.01	v5.10	0	0	localhost	01/28/16 11:30:04
SLHOST17(sl_amx)	<input type="checkbox"/>		ACTIVE	0.69	3,323.74	2.20		9	9	localhost	01/28/16 11:48:21
SLHOST21(dev)	<input type="checkbox"/>		ACTIVE	4.00	2,446.26	20.80		0	0	localhost	01/28/16 11:48:49
SLHOST22(sl_qa_conn)	<input type="checkbox"/>		ACTIVE	0.00	5,249.51	0.00	v5.10	27	1	localhost	01/28/16 11:48:31
SLHOST5(domain5)	<input type="checkbox"/>		ACTIVE	17.33	1,763.04	0.71	v5.7	5	0	localhost	01/28/16 11:48:29
SLHOST6(domain6)	<input type="checkbox"/>		ACTIVE	3.52	915.39	1.68	v5.7	6	5	localhost	01/28/16 11:48:21

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

Filter: Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.

Table:

Server Name of Server Agent.

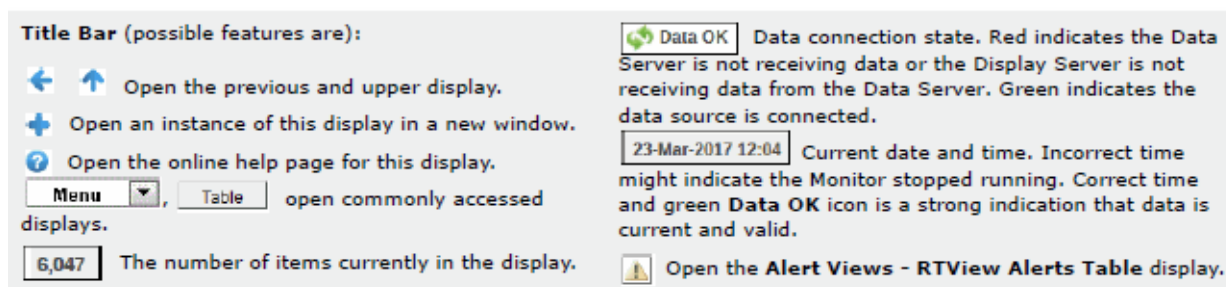
Expired When checked, data has not been received from this host in the specified amount of time.

Alert Level	<p>The most critical alert state for alerts in the row:</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
State	The current status of the application. Valid values are Running and Stopped .
CPU Usage (%)	Percent of server CPU in use.
Free Memory (MB)	Available physical memory (MB) remaining.
V. Memory Usage (%)	Percent of virtual memory used.
BW Version	The TIBCO BusinessWorks version currently in use on the server.
Deployed Engines	Total number of engines deployed on the server.
Active Engines	Number of engines currently active.
Source	Name of RTView Data Server sending this data (or localhost).
Time Stamp	Time this data was retrieved.

All Servers Grid

This grid provides a list view of utilization metrics for all BW servers (represented in the All Servers Heatmap). Track and view in parallel the general performance of all BW servers. Click on a node to drill down to the ["Single Server Summary"](#) display and view detailed metrics for that particular server.



**Filter By:**

The display might include these filtering options:

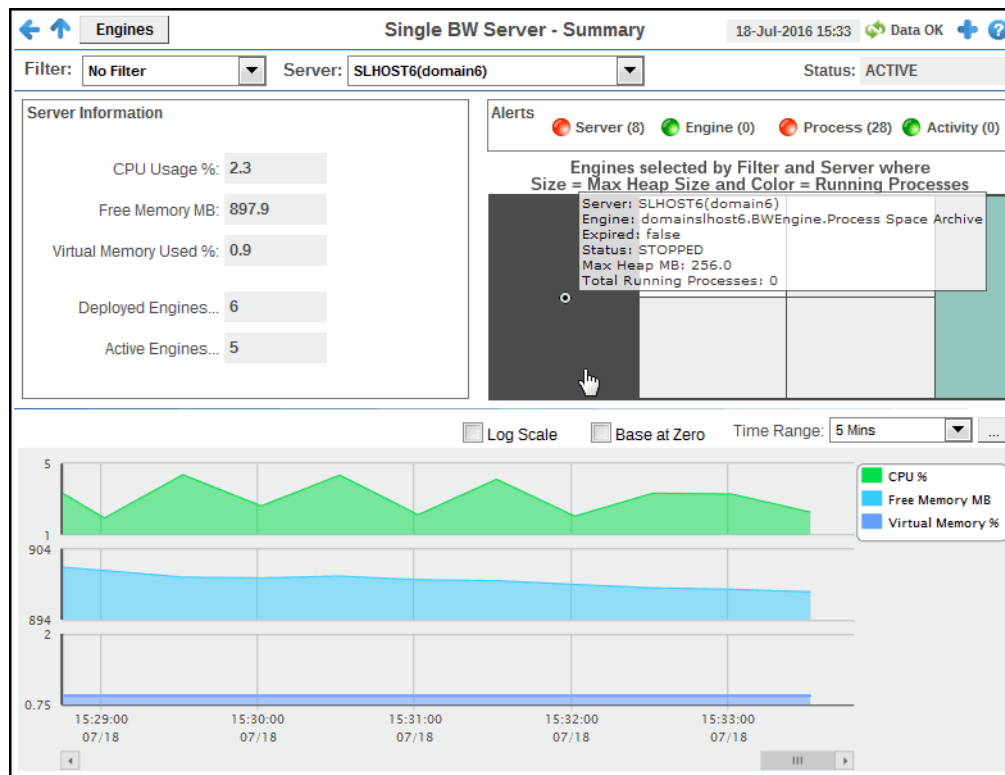
- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Time Range** Choose a time range to show data for in the display. Options are: **All Data, 2 Mins, 5 Mins, 20 Mins, 1 Hour, 2 Hours, 4 Hours, 8 Hours, 24 Hours, 2 Days and 7 Days.**

Fields and Data

- Server Name** Name of the server.
- CPU Usage%** Percent of server CPU in use.
- Free Memory** Available physical memory (MB) remaining.
- Virtual Mem Used%** Percent of virtual memory used.
- State** Server status: ACTIVE or EXPIRED.
- Deployed Engines** Total number of engines deployed on the server.
- Active Engines** Number of engines currently active.
- Trend Graphs** Shows data for the server.
- CPU** Traces percent of server CPU in use.
- MEM** Traces available physical memory remaining.
- VMEM** Traces the percent of virtual memory used.

Single Server Summary

Detailed performance metrics and alert status for a single BW server. Click on any alert indicator to drill down to the **BW Server - Tables** display to view current alerts for the selected server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to see metrics for.
- Status** Server status: ACTIVE or EXPIRED.

Server Information

- CPU Usage (%)** Percent of server CPU in use. Values are traced in trend graph (below).

Free Memory (MB)	Available physical memory remaining (in MB). Values are traced in trend graph (below).
V. Memory Usage (%)	Percent of virtual memory used. Values are traced in trend graph (below).
Deployed Engines	Number of engines currently active. Click to drill-down to details for deployed and active engines in the “All Engines Table” on page 986 display.
Active Engines	Shows data for the server. Click to drill-down to details for active engines in the “All Engines Table” on page 986 display.

Alerts

Indicates the greatest severity level and the number of open **Server**, **Engine**, **Process**, and **Activity** alerts for the selected server. Values range from **0** to **2**, where **2** is the greatest Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

Click on the alert indicator to display the **BW Server - Tables** display, which contains a table listing the current alerts for the selected engine.

The screenshot shows the 'BW Server - Tables' interface. At the top, there are tabs for 'Engines' and 'Summary'. Below them, a 'Filter' dropdown is set to 'No Filter' and a 'Server' dropdown is set to 'SLHOST16(sl_qa_conn)'. The main table, titled 'BW Server Table', has columns: Time Stamp, Server, CPU Usage (%), Free Memory (MB), V. Memory Usage (%), BW Version, Source, Deployed E, Active Engi, Expired, and State. A single row of data is visible for the selected server. Below the table, there is a section titled 'Current Alerts for Selected BW Server' with columns: Time, Alert Name, Alert Index, and Alert Text. This section is currently empty.

Heatmap

Engines selected by Filter and Server, where the heatmap rectangle size = Max Heap Size and the heatmap rectangle color = Running Processes. Dark green is the highest value for the metric shown). Click on a node to drill down to a specific engine:

- Red indicates that the engine is expired.
- Gray indicates that the engine is stopped.

Trend Graphs


Traces **CPU %**, **Free Memory MB** and **Virtual Memory %**.

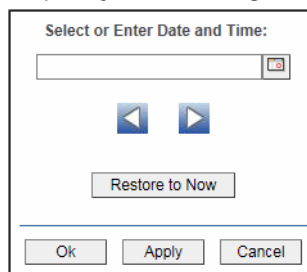
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

Base at Zero


Select to use zero (0) as the Y axis minimum for all graph traces.



Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   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.

Server Processes


Detailed information about operating system processes of a single BW Server. The heatmap shows server processes selected by Filter and Server, where the rectangle size equals memory usage and the rectangle color equals CPU percent usage.


NOTE: By default, this display is not enabled. For details, see **Enable BW Servers**.




Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

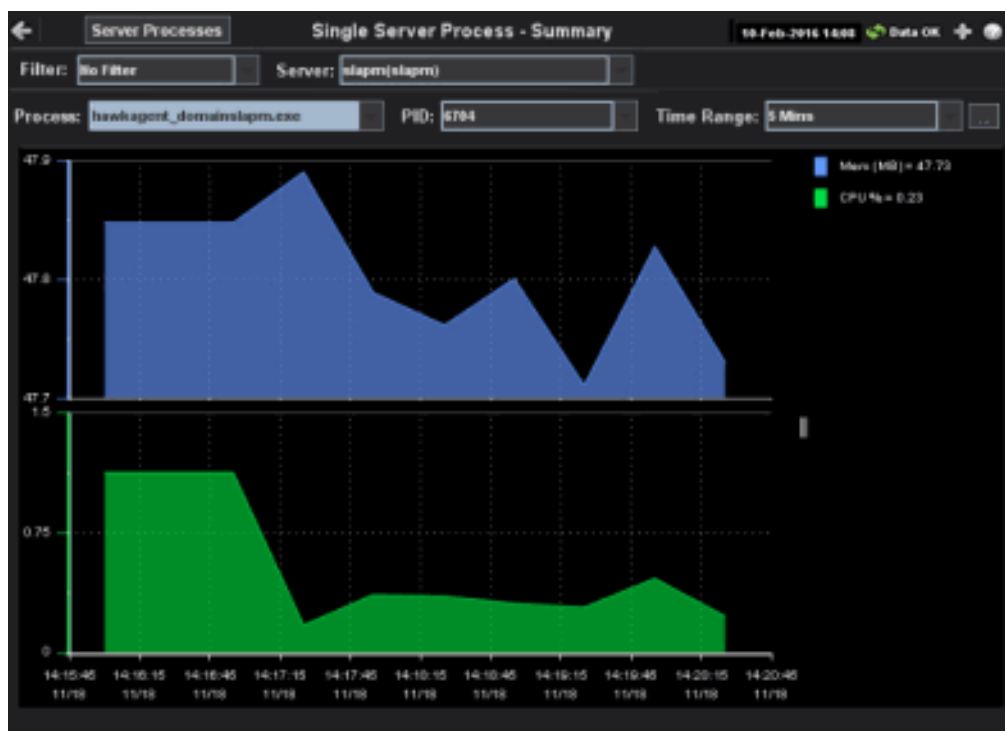
Filter: Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.


Server: Choose a server to see metrics for.

Single Server Process - Summary

Detailed information about a single operating system process running on a single BW Server.

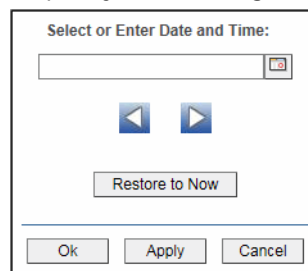
NOTE: By default, this display is not enabled. For details, see **Enable BW Servers**.



**Filter By:**

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to see metrics for.
- Process:** Choose a server process.
- PID:** Choose a server PID.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

TIBCO Enterprise Message Service

The following TIBCO Enterprise Message Service Views (and their associated displays) can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO Enterprise Message Service is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central.

- [“All EMS Servers”](#)
- [“Single EMS Server”](#)
- [“EMS Topics”](#)
- [“EMS Queues”](#)
- [“EMS Clients”](#)

All EMS Servers

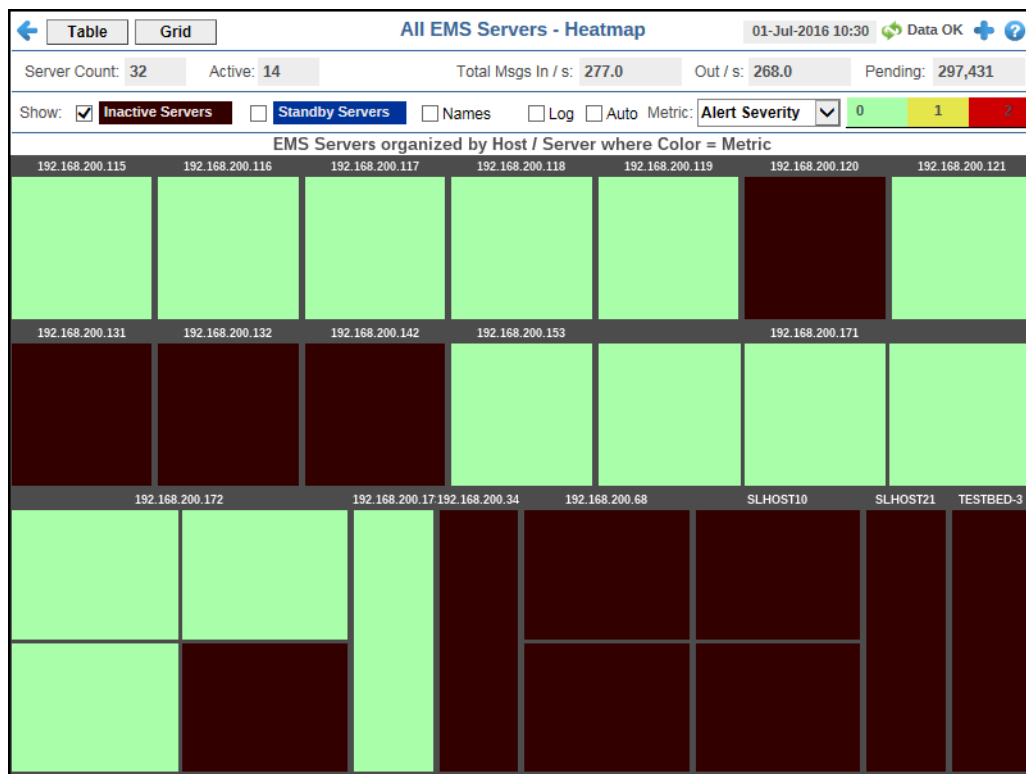
These displays present performance metrics and alert status for all EMS servers. The first three displays show different views of the same data:

- [“All Servers Heatmap”](#): Heatmap shows server and alert status for all EMS servers.
- [“All Servers Table”](#): Table shows all available utilization metrics for all EMS servers.
- [“All Servers Grid”](#): Grid enables you to see general performance of EMS servers in parallel. If you have few servers, this display is useful for verifying servers are active and generally performing as expected.
- [“All Servers Topology”](#): Topology of server routes and connections, as well as the status of active servers and standby servers that form a fault-tolerant pair.

All Servers Heatmap

View status and alerts of all EMS servers. Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **Connections**, **Pending Messages**, **Inbound Message Rate**, **Outbound Message Rate**, or **Message Memory Percent (%)**.

The heatmap is organized by host, each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the Single Server Summary display and view metrics for a particular server. Toggle between the commonly accessed **Table**, **Grid**, and **Heatmap** displays. Mouse-over rectangles to view more details about host performance and status.



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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.




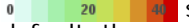

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

Server Count	The total number of active, inactive, and standby EMS servers.		
Active	The total number of currently active EMS servers.		
Total Msgs In/s	In/s	The total number of inbound messages, per second, from all producers and consumers on all EMS servers.	
	Out/s	The total number of outbound messages, per second, from all producers and consumers on all EMS servers.	
	Pending	The total number of pending messages waiting to be processed on all EMS servers. Click to open the "All Servers Table" display.	

Show	Select the type of servers for which to display data. By default, all active servers are displayed.
Inactive Servers	Select to include servers that are not currently running. Inactive Servers are represented in dark red.
Standby Servers	Select to include servers that are currently in Standby mode. Standby Servers are represented in blue.
Names	Select to display the names of servers on the hosts.
Log	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
Auto	When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
Metric	Select the metric driving the heatmap display. The default is Alert Severity. Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated "Single Server Summary" display for a detailed view of metrics for that particular server.
Alert Severity	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar , where 2 is the greatest Alert Severity.</p> <p>2 -- Metrics that have exceeded their specified ALARMLEVEL threshold and have an Alert Severity value of 2 are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p>1 -- Metrics that have exceeded their specified WARNINGLEVEL threshold and have an Alert Severity value of 1 are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p>0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of 0 and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
Alert Count	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

Connections	<p>The total number of connections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The Auto option does not impact this metric.</p>
Pend Messages	<p>The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of EmsServerPendingMsgsHigh, which is 3500. The middle value in the gradient bar indicates the middle value of the range (the default is 1750).</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
In Msg Rate	<p>The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of EmsServerInMsgRateHigh, which is 40. The middle value in the gradient bar indicates the middle value of the range (the default is 20).</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Out Msg Rate	<p>The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of EmsServerOutMsgRateHigh, which is 40. The middle value in the gradient bar indicates the middle value of the range (the default is 20).</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
Mem Msg %	<p>The percent (%) memory used by messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from 0 to the alert threshold of EmsServerMemUsedHigh, which is 40. The middle value in the gradient bar indicates the middle value of the range (the default is 20).</p> <p>When Auto is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

All Servers Table

Investigate detailed utilization metrics for all EMS servers. The **All Servers Table** contains all metrics available for servers, including the number of current client connections. Each row in the table contains data for a particular server. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill-down to the [“Single Server Summary”](#) display and view metrics for that particular server. Toggle between the commonly accessed **Table**, **Grid**, and **Heatmap** displays.

Heatmap Grid EMS All Servers Table 01-Jul-2016 10:37 Data OK

Server Count: 32 Active: 14 Total Msgs In / s: 171.0 Out / s: 211.0 Pending: 297,459

Show: ☒ Inactive Servers ☐ Standby Servers

URL	serverName	Host	Expired	Alert Level	state	versionIn
tcp://192.168.200.115:7222	EMS-SERVER	192.168.200.115	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.116:7222	EMS-SERVER	192.168.200.116	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.117:7222	EMS-SERVER	192.168.200.117	<input type="checkbox"/>		Active	6.1.0.6
tcp://192.168.200.118:7222	EMS-SERVER	192.168.200.118	<input type="checkbox"/>		Active	6.3.0.5
tcp://192.168.200.119:7222	EMS-SERVER	192.168.200.119	<input type="checkbox"/>		Active	6.3.0.5
tcp://192.168.200.120:7222	Unknown (tcp://192.168.200...	192.168.200.120	<input checked="" type="checkbox"/>			
tcp://192.168.200.121:7222	EMS-SERVER	192.168.200.121	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.131:7222	TESTBED-1 (tcp://192.168.2...	192.168.200.131	<input checked="" type="checkbox"/>			
tcp://192.168.200.132:7222	TESTBED-2 (tcp://192.168.2...	192.168.200.132	<input checked="" type="checkbox"/>			
tcp://192.168.200.142:7222	Unknown (tcp://192.168.200...	192.168.200.142	<input checked="" type="checkbox"/>			
tcp://192.168.200.153:7222	EMS-SERVER-153	192.168.200.153	<input type="checkbox"/>		Active	8.2.2.3
tcp://192.168.200.171:6010	EMS-SLDEMOS2-6010	192.168.200.171	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.171:6020	EMS-SLDEMOS2-6020	192.168.200.171	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.171:6030	EMS-SLDEMOS2-6030	192.168.200.171	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8011	EMS-SLDEMOS3-8010	192.168.200.172	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8020	EMS-SLDEMOS3-8020	192.168.200.172	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8030	EMS-SLDEMOS3-8030	192.168.200.172	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8031	Unknown (tcp://192.168.200...	192.168.200.172	<input checked="" type="checkbox"/>			
tcp://192.168.200.173:9010	EMS-SLDEMOS4-9010	192.168.200.173	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.34:7222	TESTBED-34 (tcp://192.168...	192.168.200.34	<input checked="" type="checkbox"/>			
tcp://192.168.200.68:7222	Unknown (tcp://192.168.200...	192.168.200.68	<input checked="" type="checkbox"/>			
tcp://192.168.200.68:7224	Unknown (tcp://192.168.200...	192.168.200.68	<input checked="" type="checkbox"/>			
tcp://SLHOST10:7010	Unknown (tcp://SLHOST10:7...	SLHOST10	<input checked="" type="checkbox"/>			
tcp://SLHOST10:7011	Unknown (tcp://SLHOST10:7...	SLHOST10	<input checked="" type="checkbox"/>			
tcp://SLHOST21:7222	EMS-SERVER-SLHOST21 (t...	SLHOST21	<input checked="" type="checkbox"/>			
tcp://TESTBED-3:7022	Unknown (tcp://TESTBED-3:...	TESTBED-3	<input checked="" type="checkbox"/>			

Title Bar (possible features are):




- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

- Server Count** The total number of active, inactive and standby EMS servers. **Inactive Servers** are represented in dark red. **Standby Servers** are represented in blue.
- Active** The total number of currently active EMS servers.
- Total Msgs In/s** The total number of inbound messages, per second, from all producers and consumers on all EMS servers.
- Out/s** The total number of outbound messages, per second, from all producers and consumers on all EMS servers.
- Pending** The total number of inbound and outbound messages waiting to be processed on all EMS servers.

Show	Select the type of servers to display data for. By default, all active servers are displayed.
Inactive Servers	Select to include servers that are not processing requests in the table. Inactive Servers are represented in dark red.
Standby Servers	Select to include servers that are not currently running. Standby Servers are represented in blue.
Table	This table shows information for all EMS servers. Click on a table row to drill-down to the "Single Server Summary" display and view metrics for that particular server.
URL	Select to include servers that are currently in Standby mode. Standby Servers are represented in blue.
serverName	The name of the server.
Host	The name or IP address for the host server.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Alert Level	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar, where 2 is the greatest Alert Severity.</p> <p> -- One or more alerts have exceeded their specified ALARMLEVEL threshold, have an Alert Severity value of 2, and are shown in red.</p> <p> -- One or more alerts have exceeded their specified WARNINGLEVEL threshold, have an Alert Severity value of 1, and are shown in yellow.</p> <p> -- No alerts have exceeded an alert threshold, which have an Alert Severity value of 0, and are shown in green.</p>
state	<p>The server status:</p> <p>Active -- The server is currently processing requests.</p> <p>Inactive -- The server is not currently processing requests. Inactive Servers are represented in dark red.</p> <p>Standby -- The server is functioning as a backup for a primary server. Standby Servers are represented in blue.</p>
versionInfo	The TIBCO EMS software version currently running.
faultTolerantURL	The IP address and port number for the source (application, server, and so forth) associated with the alert.
asyncDBsize	The amount of database space, in bytes, occupied by asynchronous data on the server.
backupName	The name of the backup server assigned as the backup to this server.
connectionCount	The number of clients currently connected to the server.
diskReadRate	The speed at which the server reads disk data.
diskWriteRate	The speed at which the server writes data to disk.

durableCount	The number of durables on the server.
inboundBytesRate	The rate of inbound messages in bytes per second.
inboundMessageCount	The number of inbound messages received by the server since the server was started.
inboundMessageRate	The rate of inbound messages in number of messages per second.
MaxMessageMemory	The maximum amount of memory, in bytes, allocated for use by messages on the server.
messageMemory	The amount of memory, in bytes, currently used by messages on the server.
messageMemoryPct	The amount of memory, in percent, used by messages on the server.
messageMemoryPooled	The currently allocated pool size, in bytes, for messages.
outboundBytesRate	The rate of outbound messages in bytes per second.
outboundMessageCount	The number of outbound messages sent by the server since the server was started.
outboundMessageRate	The rate of outbound messages in number of messages per second.
pendingMessageCount	The number of currently pending messages on the server.
pendingMessageSize	The amount of space, in bytes, pending messages use on the server.
processId	The process ID of the EMS server.
queueCount	The number of message queues.
startTime	The date and time that the server was started.
syncDBSize	The amount of database space, in bytes, occupied by synchronous data on the server.
topicCount	The number of currently active topics on the server.
upTime	The amount of time, in milliseconds, since the server was started.
time_stamp	The date and time this row of data was last updated.

All Servers Grid

Track and view in parallel the general performance of all EMS servers. Click on a node to drill-down to the “[Single Server Summary](#)” display and view detailed metrics for that particular server.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data


This display includes:

- Server Count** The total number of active, inactive and standby EMS servers. Inactive Servers are represented in dark red. Standby Servers are represented in blue.
- Active** The total number of currently active EMS servers.

Total Msgs	In/s	The total number of inbound messages, per second, from all producers and consumers on all EMS servers.
	Out/s	The total number of outbound messages, per second, from all producers and consumers on all EMS servers.
	Pending	The total number of inbound and outbound messages waiting to be processed on all EMS servers. Click to open the "All Servers Table" display.
Show		Select the type of servers to display data for. By default, all active servers are displayed.
	Inactive Servers	Select to include servers that are not processing requests in the table. Inactive Servers are represented in dark red.
	Standby Servers	Select to include servers that are not currently running. Standby Servers are represented in blue.
Sort By	Server Name	Select to organize the servers in the grid by server name.
	Server URL	Select to organize the servers in the grid by server URL.
Descending		When checked, lists servers in the grid in descending order.
Time Range		Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data .
Grid	Server Name	The name of the server.
	URL	The URL for the server.
	Uptime	The amount of time, in milliseconds, since the server was started.
	Pend Msgs	The number of currently pending messages on the server.
	State	The server status: Active -- The server is currently processing requests. Inactive -- The server is not currently processing requests. Inactive Servers are represented in dark red. Standby -- The server is functioning as a backup for a primary server. Standby Servers are represented in blue.
	In Rate	The rate of inbound messages in messages per second.
	Out Rate	The rate of outbound messages in messages per second.
	Trend Graphs	Shows message data for the server. Pend -- Traces the total number of pending messages on the server. In -- Traces the rate of inbound messages in messages per second. Out -- Traces the rate of outbound messages in messages per second.

All Servers Topology

View a server topology map for all EMS servers. Click on a node to drill-down to the “[Single Server Summary](#)” display and view metrics for that particular server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Clicking the **Route Table** button displays the **EMS Server Route Table** window. See “[EMS Server Route Table](#)” for more information.

Fields and Data

This display includes:

Show	The total number of active, inactive and standby EMS servers. Inactive Servers are represented in dark red. Standby Servers are represented in blue.
Inactive Servers	Select to show servers that are not processing requests in the topology. Inactive Servers are represented in dark red.
Standby Servers	Select to show servers that are not processing requests in the topology. Standby Servers are represented in blue.
Connected Routes Only	Select to show only routes that have an active connection.
Show Metrics	Available on desktop application deployments only. Shows the total input message rates, per second, on the top of each server icon and the total output message rate on the bottom of each server icon.
Topology	<p>Routes are shown between the active server and the standby server, which form a fault-tolerant pair. Either of the servers in a fault-tolerant pair can become the active server or the standby server. Show Standby Servers and Show Inactive Servers enable you to include or exclude standby and inactive servers. Inactive Servers are represented in dark red. Standby Servers are represented in blue. By default, standby servers are included in the topology and inactive servers are not.</p> <p>Typically, it takes about 30 seconds for a server to appear in the display after startup.</p> <p>The active server in a fault-tolerant pair appears in green with the suffix (A) appended to its URL. The standby server appears in blue, with the suffix (S) appended to its URL. Their link is blue and labeled FT.</p> <p>If the active server fails:</p> <ul style="list-style-type: none"> the failed server becomes inactive, its suffix changes to (X!), and the node turns red with a red outline. the standby server becomes active, its suffix changes to (A!), and the node turns green with a red outline. the link between the two servers turns red. <p>If the standby server fails:</p> <ul style="list-style-type: none"> the failed server becomes inactive, its suffix changes to (X!), and the node turns red with a red outline. the active servers' suffix changes to (A!) and it is outlined in red. the link between the two servers turns red. <p>If a failed server recovers:</p> <ul style="list-style-type: none"> the recovered server becomes the standby server, its suffix changes to (S), and the node turns blue with a grey outline. the active servers' suffix (A!) changes to (A), and the red node outline changes back to grey. the link between the two servers changes back to blue. <p>Suffix Definition</p> <p>A -- This is the active server and it is running. A! -- This is the active server and it is running but its standby has failed. S -- This is the standby server and it is running. X! -- The server is inactive.</p> <p>Node Color Definition</p> <p>-- This is the active server and it is running. Blue -- This is the standby server and it is in standby mode. -- The server is inactive.</p>

Link Color Definition

Blue -- The two servers in the pair are running.
 -- One of the servers in the pair is inactive.

Outline Color Definition

Grey -- The two servers in the pair are running.
 -- One of the servers in the pair is inactive. If the node color indicates this server is running, its pair is inactive.


EMS Server Route Table

Displays metrics for server routes on all servers. Inbound metrics, such as **inboundByteRate**, indicate an in route to the server. Outbound metrics, such as **outboundByteRate**, indicate an out route to the server.

remoteURL	remoteName	connected	stalled	inboundByteRate	inboundMessageRate
tcp://192.168.200.171:6020,tcp://192.168.200.1...	EMS-SLDEMOS2-6020		0	0.0	
tcp://192.168.200.171:6020,tcp://192.168.200.1...	EMS-SLDEMOS2-6020		0	0.0	
tcp://192.168.200.171:6030,tcp://192.168.200.1...	EMS-SLDEMOS2-6030		0	0.0	
tcp://192.168.200.171:6030,tcp://192.168.200.1...	EMS-SLDEMOS2-6030		0	0.0	
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0	
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0	
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0	
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0	
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0	
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0	
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0	
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0	
tcp://localhost:7022	EMS-SERVER2		0	0.0	
tcp://localhost:7022	EMS-SERVER2		0	0.0	
tcp://localhost:7022	EMS-SERVER2		0	0.0	
tcp://localhost:7022	EMS-SERVER2		0	0.0	
tcp://SLHOST10	EMS-SLDEMOS1-7010		0	0.0	
tcp://vmrh5-4	EMS-SLDEMOS2-6010		0	0.0	
tcp://vmrh5-4	EMS-SLDEMOS2-6010		0	0.0	

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.
- open commonly accessed displays.
- The number of items currently in the display.



Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.






Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

- remoteURL** The remote URL of the server.
- remoteName** The name of the server.
- connected** The connection state of the server route.

	 -- One or more routes for this server are disconnected.  -- All routes for this server are connected.  -- There are no routes for this server.
stalled	<p>Indicates whether the IO flow stalled on the route.</p> <p>A value of 0 (zero) = not stalled.</p> <p>A value of 1 = stalled.</p>
inboundByteRate	The rate of inbound data in bytes, per second.
inboundMessageRate	The rate of inbound messages in number of messages per second.
inboundTotalBytes	The total number of inbound bytes.
inboundTotalMessages	The total number of inbound messages.
outboundByteRate	The rate of inbound data in bytes, per second.
outboundMessageRate	The rate of outbound messages in number of messages per second.
outboundTotalBytes	The total number of outbound bytes.
outboundTotalMessages	The total number of outbound messages.
zoneName	The name of the zone for the route.
zoneType	Indicates a multi-hop or one-hop route.
active	<p>Indicates whether the server route is currently transferring data:</p> <p>1 = true (is transferring data)</p> <p>0 = false</p>
inactive	<p>Indicates whether the server route is not currently transferring data:</p> <p>1 = true (is not transferring data)</p> <p>0 = false</p>
suspended	<p>Indicates whether outbound messages to the route have been suspended:</p> <p>1 = true</p> <p>0 = false</p>
remoteURLName	The IP address and name for the remote connection.


Single EMS Server

These displays present detailed performance metrics, alert status and connection information for a single EMS server.

- **"Single Server Summary"**: Shows information for a single EMS server such as server connection details, the number of client connections, memory utilization, message performance metrics and alert status.
- **"Single Server Trends"**: Trend graphs show utilization metrics for a single EMS server, such as the number of client connections, number of pending messages and in/out rate, and memory and disk utilization.
- **"Single Server Tables"**: Tables show information about how the Monitor is connected to the EMS server, metrics queried from the server and alert details.

Single Server Summary

Track utilization and performance metrics for specific servers.






Title Bar (possible features are):










- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data


This display includes:


EMS Server		Select the EMS Server for which you want to view data. The selection made here populates this display.
Name		The name of the EMS Server selected from the EMS Server drop-down menu.
Server Information	Version	The TIBCO EMS software version currently running.
	Start Time	The date and time that the server was started.
	Backup	The name of the backup server for the server.
	FT URL	The IP address and port number, or the hostname and port number, of the fault tolerant standby server assigned to this server.
	Last Data Time	The time that a data update was last made.
	State	The server status: Active -- The server is currently processing requests. Inactive -- The server is not currently processing requests. Standby -- The server is functioning as a backup for a primary server.
	Uptime	The amount of time since the server was started. Format: dd HH:MM:SS <days> <hours>:<minutes>:<seconds> For example: 10d 08:41:38
	Persistence	
	Async Storage	The amount of database space, in bytes, used by asynchronous message persistence data on the server
	Sync Storage	The amount of database space, in bytes, used by synchronous message persistence data on the server.
	Mem Used	The amount of memory, in kilobytes, used by message persistence on the server.
	Max	The maximum amount of memory, in kilobytes, used by message persistence on the server.
	Used %	The amount of memory, in percent, used by message persistence.
	Pooled	The amount of message memory that has been pooled.
	Disk Read Rate	The speed at which the server reads message persistence disk data.
	Disk Write Rate	The speed at which the server writes message persistence disk data.
Alerts	Server	Status indicator for server-related alerts. Click to open the EMS "Single Server Tables" display and view the Server Alert Table for more detail.  -- No alerts have exceeded a specified threshold.  -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.  -- One or more alerts have exceeded their specified ALARMLEVEL threshold.


	Routes	<p>Status indicator for route-related alerts. Click to open the EMS "Single Server Tables" display and view the Server Alert Table for more detail.</p> <p> -- No alerts have exceeded a specified threshold.</p> <p> -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.</p> <p> -- One or more alerts have exceeded their specified ALARMLEVEL threshold.</p>
	Topics	<p>Status indicator for topic-related alerts. Click to open the EMS "Single Server Tables" display and view the Server Alert Table for more detail.</p> <p> -- No alerts have exceeded a specified threshold.</p> <p> -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.</p> <p> -- One or more alerts have exceeded their specified ALARMLEVEL threshold.</p>
	Queues	<p>Status indicator for queue-related alerts. Click to open the EMS "Single Server Tables" display and view the Server Alert Table for more detail.</p> <p> -- No alerts have exceeded a specified threshold.</p> <p> -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.</p> <p> -- One or more alerts have exceeded their specified ALARMLEVEL threshold.</p>
Connections / Destinations		Shows connection information for the server. The counts shown here are also visible in the "EMS Topics" and "EMS Clients" displays.
	Producers	The number of producers currently active on the server. Click to open the "EMS Clients" / "Producers" for Server display for details.
	Durables	The number of durables currently active on the server. Click to open the "EMS Clients" / "Consumer Summary" for Server display for details.
	Routes	The number of routes defined on the server.
	Connections	The number of clients currently connected to the server. Click to open the "EMS Clients" / "Connections" for Server display for details.
	Consumers	The number of consumers currently connected to the server. Click to open the "EMS Clients" / "Producer Summary" for Server display for details.
	Topics	The number of topics currently active on the server. Click to open the "EMS Topics" / "All Topics Table" display for details.
	Queues	The number of queues currently active on the server. Click to open the "EMS Topics" / "All Queues Heatmap" display for details.
Messages In	Msgs/sec	The number of inbound messages, per second, from all producers and consumers
	Bytes in/sec	The total size of inbound messages, in bytes per second, from all producers and consumers.
	Total	The total number of inbound messages, in bytes, from all producers and consumers since the server was started.



Messages Out	Msgs/sec	The number of outbound messages, per second, from all producers and consumers.
	Bytes out/sec	The total size of outbound messages, in bytes per second, from all producers and consumers.
	Total	The total of outbound messages, in bytes, from all producers and consumers since the server was started.
Pending Messages	Current	The total number of inbound and outbound messages currently waiting to be processed.
	Bytes pending	The total size of inbound and outbound messages, in bytes, currently waiting to be processed.
Trend Graphs	Shows message metrics for the selected server.	
	Pend Message	-- Traces the total number of inbound and outbound messages currently waiting to be processed.
	In Msgs / sec	-- Traces the number of inbound messages, per second, from all producers and consumers. This trend graph only displays when Use Rates is selected.
	Out Msgs / sec	-- Traces the number of outbound messages, per second, from all producers and consumers. This trend graph only displays when Use Rates is selected.
	Delta In Msgs	-- Traces the change in total inbound messages since the last update. This trend graph only displays when Use Rates is not selected.
	Delta Out Msgs	-- Traces the change in total outbound messages since the last update. This trend graph only displays when Use Rates is not selected.
	Use Rates	When this check box is selected, the inbound and outbound message rates (In Msgs/sec and Out Msgs/sec) display in the trend graph. When this check box is not selected, the delta inbound and outbound messages (Delta In Msgs and Delta Out Msgs) display in the trend graph.
	Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

Base at Zero When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.




By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Single Server Trends

View trend graphs in parallel to investigate performance issues for a specific server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

- EMS Server** Select the EMS server for which you want to view data from this drop-down menu. The selection made here populates this display.
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.
- Server Activity Trends** Specifies settings for the trend graphs.

**Trend
Graphs**

Shows metrics for the selected server.

Connections -- Traces the total number of client connections.

Msgs In/Sec -- Traces the number of inbound messages, per second, from all producers and consumers.

Msgs Out/Sec -- Traces the number of outbound messages, per second, from all producers and consumers.

Pending Msgs -- Traces the total number of messages currently waiting to be processed.

Msg Memory % -- Traces the amount of memory, in percent, used by messages.

Async Store MB -- Traces the amount of database space, in megabytes, used by asynchronous data on the server.

Disk Read KB -- Traces the amount of disk data, in kilobytes, read by the server since the server was started.

Disk Write KB -- Traces the amount of data, in kilobytes, written to disk by the server since the server was started.


Log Scale


This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero

When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Single Server Tables

View all available utilization and performance data for specific servers.

Agent	If used, the name of the RTView agent connecting to the EMS server.
User	The user name for gaining access to the server.
Password	The password associated with user name for gaining access to the server.
Subs	RTView substitutions used when connecting to this server.
ConnName	The name of the RTView connection to this server.
Active	When checked, indicates that the server is currently running.
FaultTolerantStandbyMode	When checked, indicates that the server is running as a backup server.
FaultTolerantURL	The IP address and port number for the backup server assigned to this server.
BackupName	The name of the backup server assigned as backup to this server.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Server Info Table	Select an EMS Server from the EMS Server drop-down menu. This table shows server metrics queried from the server.
time_stamp	The date and time this row of data was last updated.
Host	The name or IP address for the host server.
asyncDBSize	The amount of database space, in bytes, used by asynchronous data on the server.
backupName	The name of the backup server assigned as backup to this server.
connectionCount	The number of currently connected clients.
diskReadRate	The speed at which the server reads disk data.
diskWriteRate	The speed at which the server writes data to disk.
durableCount	The number of currently active durables.
FaultTolerantURL	The IP address and port number, or the hostname and port number, of the fault tolerant standby server assigned to this server.
inboundBytesRate	The rate of inbound messages in bytes per second.
inboundMessageCount	The number of inbound messages received by the server since the server was started.
inboundMessageRate	The rate of inbound messages in number of messages per second.
maxMessageMemory	The maximum amount of memory, in bytes, allocated for use by messages on the server.
messageMemory	The amount of memory, in bytes, currently used by messages on the server.

messageMemoryPct	The amount of memory, in percent, used by messages on the server.
messageMemoryPooled	The currently allocated pool size for messages in bytes.
outboundBytesRate	The rate of outbound messages in bytes per second.
outboundMessageCount	The number of outbound messages sent by the server since the server was started.
outboundMessageRate	The rate of outbound messages in number of messages per second
pendingMessageCount	The number of currently pending messages on the server.
pendingMessageSize	The amount of space, in bytes, pending messages use on the server.
processId	The process ID of the EMS server.
queueCount	The number of message queues.
serverName	The name of the server.
startTime	The date and time that the server was started.
state	The server status: Active -- The server is currently processing requests. Inactive --The server is not currently processing requests. Standby -- The server is functioning as a backup for a primary server.
syncDBSize	The amount of database space, in bytes, used by synchronous data on the server.
topicCount	The number of currently active topics.
upTime	The amount of time, in milliseconds, since the server was started.
versionInfo	The TIBCO EMS software version currently running.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Current Alerts Table for Selected EMS Server	Select an EMS Server from the EMS Server drop-down menu. This table lists all available data for currently active alerts. Click an alert to view details in the Alert Detail Window.
Time	The time the alert was first activated.
Alert Name	The name of the alert.
Alert Index	The EMS server that activated the alert.
Alert Text	The text that is displayed for the alert.
Package	The RTView package reporting the alert.
Category	The alert category: Server, Queue or Topic.

ID	The unique identifier for this alert instance.
Clr'd	When checked, the alert thresholds are no longer out of bounds and the alert has cleared.
Ack'd	When checked, a user has indicated that they have acknowledged the alert.
Owner	The user who has accepted ownership of this alert.
Source	The source of the alert.

**Alert
Detail
Window**

The screenshot shows the 'Alert Detail' window with the following fields and values:

- Alert Time:** 02/26/15 07:27:45
- ID:** 1000
- Name:** EmsServerMemUsedHigh
- Index:** tcp://SLHOST21:7222
- Owner:** (empty field)
- Alert Text:** High Warning Limit exceeded, current value: 9.36 limit: 5.0
- Comments:** (empty text area)
- Acknowledged:** ☐
- Cleared:** ☐
- Severity:** 1

Alert Time	The time the alert was first activated.
ID	The unique identifier for this alert instance.
Name	The name of the alert.
Index	The EMS server which activated the alert.
Owner	The user who has accepted ownership of this alert.
Alert Text	The text that is displayed for the alert.
Comments	User-supplied comments about this alert.
Acknowledged	When checked, a user has indicated that they have acknowledged the alert.
Cleared	When checked, the alert thresholds are no longer out of bounds and the alert has cleared.
Severity	Severity of the alert.

EMS Topics

These displays present several views of performance metrics for topics. You can view all topics that are defined on a specific server in the [“All Topics Table”](#) display, or you can view all servers that have a specific topic defined in the [“Single Topic Summary”](#) display. The [“Single Topic By Server”](#) display provides a list of all the servers on which those topics are defined.

- [“All Topics Heatmap”](#): A heatmap representation of a selected set of metrics from Topics organized by Server that allows you to track performance and utilization metrics and trends for all topics on a single server.
- [“All Topics Table”](#): Shows performance and utilization metrics and trends for all topics defined on a specified server, including consumer and subscriber count, memory utilization, and message performance metrics.
- [“All Topics Summary”](#): Shows performance and utilization metrics and trends for all topics defined on a specified server, including consumer and subscriber count, memory utilization, and message performance metrics.
- [“Single Topic Summary”](#): Shows detailed performance and utilization metrics and trends for a specified topic on a single server, including producer and consumer counts, and message performance metrics.
- [“Single EMS Topic-Clients”](#): View data for all consumers and producers associated with the selected topic.
- [“Single Topic By Server”](#): Table shows performance and utilization metrics for all servers that have a specified topic defined, including consumer and subscriber count, and message performance metrics.

All Topics Heatmap

A heatmap representation of a selected set of metrics from Topics organized by Server that allows you to track performance and utilization metrics and trends for all topics on a single server. View status and alerts of all topics for a server. Use the **Metric** drop-down menu to view to **Alert Severity**, **Alert Count**, **Consumers**, **Receivers**, **Pending Messages**, **Inbound Message Rate**, **Inbound Total Messages**, **Outbound Message Rate**, or **Outbound Total Messages**.

The heatmap is organized so that each rectangle represents a Topic on the selected Server. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each topic or click on a rectangle to drill-down to the [“Single Topic Summary”](#) display and view metrics for that particular Topic. You can click **Table** on this display to navigate to the [“All Topics 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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Clicking **Table** in the Title Bar takes you to the “All Topics Table” display. Clicking **Summary** in the Title Bar takes you to the “All Topics Summary” display.

Fields and Data

This display includes:

- | | |
|-----------------------|---|
| EMS Server | The EMS Server selected from this drop-down menu populates all associated Topic data in this display. |
| Filter Pattern | Enter a string to show only topics with names that contain the string. For example, if you enter the string Madrid, all topics with Madrid in the topic name are shown in the table. If no entry is made, all topic names are shown. For most use cases, you can enter a portion of the topic name. |


Filtered Topic Count

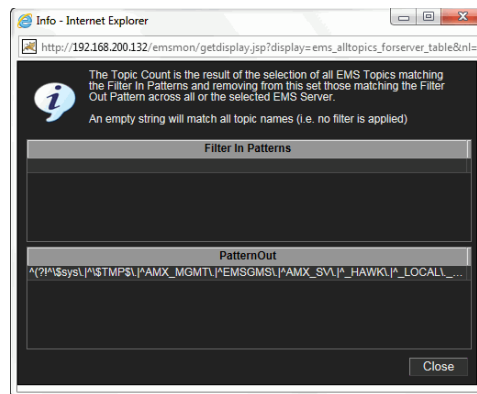
This field is broken into two different values. The first value is the total number of currently active topics on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsTopicFilterOutPattern** property in the **emsmon/conf/rtvapm.properties** file. The second value is the total number of topics on the selected server. In other words, the filtered number of topics/the total number of topics on the server.

The default value for the **\$emsTopicFilterOutPattern** property is:

```
collector.sl.rtvview.sub=$emsTopicFilterOutPattern: '^(!\\$sys\\.|^\\$TMP\\$\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\. _HAWK\\.|^_TMP\\. EMS)'
```

You can modify the filter value by editing the **\$emsTopicFilterOutPattern** property in the ["sample.properties File"](#), which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Topic Count**.

**Log**

This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

Auto

When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

Metric

Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the topics by server, where each rectangle represents a Topic. Mouse-over any rectangle to display the current values of the metrics for the Topic. Click on a rectangle to drill-down to the associated ["Single Topic Summary"](#) display for a detailed view of metrics for that particular topic.

Alert Severity

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.


2 -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

1 -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.


0 -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

Alert Count

The total number of alarm and warning alerts in a given item (index) associated with the rectangle.


The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Consumers


The total number of consumers in a given item (index) associated with the rectangle. The color gradient bar  shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of consumers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.


Durables

The total number of active durables in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of durables in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Subscribers

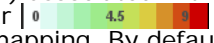
The total number of subscribers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of subscribers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Pending Msgs

The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsTopicssPendingMsgsHigh**, which is **3000**. The middle value in the gradient bar indicates the middle value of the range (the default is **1500**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.


In Msg /sec

The number of inbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsTopicsInMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.


Note: This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

In Total Msg

The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.


Out Msg/sec

The number of outbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsTopicsOutMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Note: This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

Out Total Msgs

The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

All Topics Table

Track performance and utilization metrics for all topics on a single server.

HeatmapSummary

All EMS Topics for Server - Table

06-Jul-2016 10:37Data OK

EMS Server: tcp://192.168.200.115:7222

Filter Pattern: Filtered Topic Count 11 / 15

Topic Name	URL	In Rate	In Total	Out Rate	Out Total	Pend Msgs	P
adb.custom.jmsrequest	tcp://192.1...	0	0	0	0	0	
adb.salesorder.rr	tcp://192.1...	0	0	0	0	0	
adb.salesorder.sub	tcp://192.1...	0	0	0	0	0	
adb.standard.jmsrequest	tcp://192.1...	0	582,832	0	583,934	2,233,301	1,9
MessageSelector	tcp://192.1...	0	0	0	0	0	
rtv.amx.governance.internal.stats	tcp://192.1...	0	0	0	0	0	
rtv.amx.governance.stats	tcp://192.1...	0	0	0	0	0	
sample	tcp://192.1...	0	0	0	0	0	
topic.sample	tcp://192.1...	0	0	0	0	0	
topic.sample.exported	tcp://192.1...	0	0	0	0	0	
topic.sample.imported	tcp://192.1...	0	0	0	0	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.

6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Clicking **Heatmap** in the Title Bar takes you to the "All Topics Heatmap" display. Clicking **Summary** in the Title Bar takes you to the "All Topics Summary" display.

Fields and Data

This display includes:

- EMS Server

The EMS Server selected from this drop-down menu populates all associated Topic data in this display.

Filter Pattern

Enter a string to show only topics with names that contain the string. For example, if you enter the string Madrid, all topics with Madrid in the topic name are shown in the table. If no entry is made, all topic names are shown. For most use cases, you can enter a portion of the topic name.


Filtered Topic Count

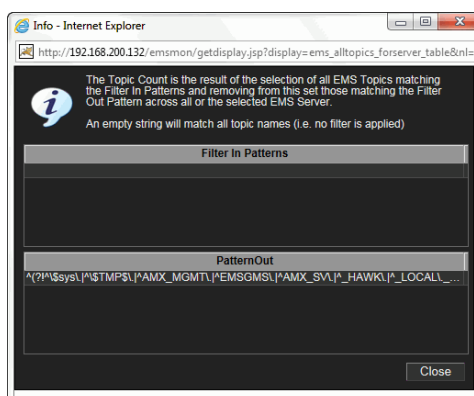
This field is broken into two different values. The first value is the total number of currently active topics on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsTopicFilterOutPattern** property in the **emsmon/conf/rtvapm.properties** file. The second value is the total number of topics on the selected server. In other words, the filtered number of topics/the total number of topics on the server.

The default value for the **\$emsTopicFilterOutPattern** property is:

```
collector.sl.rtvew.sub=$emsTopicFilterOutPattern: '^(!^\\$sys\\.|^\\$TMP\\$\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\. _HAWK\\.|^TMP\\. EMS)'
```

You can modify the filter value by editing the **\$emsTopicFilterOutPattern** property in the ["sample.properties File"](#), which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Topic Count**.

**Table**

This table describes all topics on the selected server. Click a row to view metrics for a single topic in the ["Single Topic Summary"](#) display.

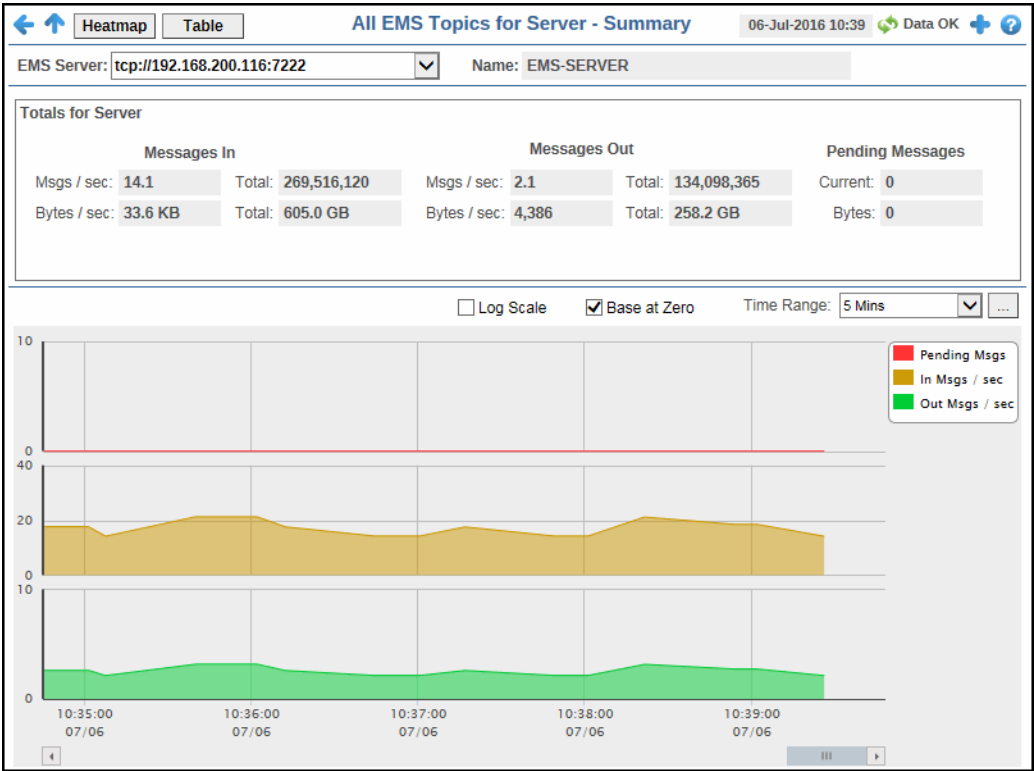
Topic Name	The name of the topic.
URL	The IP address and port number for the server.
In Rate	The number of inbound messages for the topic, per second. Note: This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.
In Total	The total number of inbound messages for the topic.
Out Rate	The number of outbound messages for the topic, per second. Note: This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.
Out Total	The total number of outbound messages for the topic.
Pend Msgs	The number of currently pending messages for the topic.
Pend Size	The amount of space, in bytes, used by pending messages for the topic.
activeDurableCount	The number of currently active durables or the topic.


consumerCount	The number of consumers for the topic.
durableCount	The number of durables for the topic.
failSafe	When checked, the message is marked as failsafe delivery.
fcMaxBytes	The maximum number of bytes allocated for use by flow control.
global	When checked, the message is global and is routed to other servers.
inboundByteRate	The amount of inbound messages for the topic, in bytes per second.
inboundTotalBytes	The total amount of inbound messages for the topic, in bytes, since the server started.
maxBytes	The maximum size, in bytes, that the topic can store for delivery to each durable or non-durable online subscriber on that topic.
maxMsgs	The maximum number of messages before the server indicates an error and overflow policies are activated.
outboundByteRate	The amount of outbound messages for the topic, in bytes per second.
outboundTotalBytes	The total amount of outbound messages for the topic, in bytes.
overflowPolicy	Indicates whether an overflow policy is set for the topic: 0 = No policy is set. 1 = A policy is set.
secure	When checked, the topic is designated as secure and enforces permission policies.
static	When checked, the topic has a static destination.
subscriberCount	The number of subscribers for the topic.
description	Descriptive text to help the administrator identify this resource.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
time_stamp	The date and time this row of data was last updated.
DeltainboundTotalMessages	Displays the change (delta) in inboundTotalMessages from the previous cache refresh to the current cache refresh.
DeltainboundTotalBytes	Displays the change (delta) in inboundTotalBytes from the previous cache refresh to the current cache refresh.
DeltaoutboundTotalMessages	Displays the change (delta) in outboundTotalMessages from the previous cache refresh to the current cache refresh.

DeltaoutboundTotalBytes	Displays the change (delta) in outboundTotalBytes from the previous cache refresh to the current cache refresh.
prefetch	Lists the maximum number of messages consumers can fetch.
expiryOverride	If set to a non-zero value for a destination and the server delivers a message to the destination, the server replaces the producer's expiration value with this value.
store	Provides the store for this destination where persistent messages are stored.
URLTopic	The topic's URL.

All Topics Summary

Track performance and utilization metrics and trends for all topics on a single server.





Note: Clicking **Heatmap** in the Title Bar takes you to the “All Topics Heatmap” display. Clicking **Table** in the Title Bar takes you to the “All Topics Table” display.

Fields and Data

This display includes:

EMS Server	The EMS Server selected from this drop-down menu populates all associated Topic data in this display.
Name	The name of the server selected in the EMS Server drop down list.
Totals for Server	Shows metrics for all topics on the selected server.
Messages In	<p>Msgs/sec -- The number of inbound messages for all topics on the server, per second.</p> <p>Total -- The total number of inbound messages for all topics on the server since the server was started.</p> <p>Bytes/sec -- The size of inbound messages, in bytes per second, for all topics on the server.</p> <p>Total -- The total size of inbound messages, in kilobytes, for all topics on the server since the server was started.</p>
Messages Out	<p>Msgs/sec -- The number of outbound messages for all topics on the server, per second.</p> <p>Total -- The total number of outbound messages for all topics on the server since the server was started.</p> <p>Bytes/sec -- The size of outbound messages, in bytes per second, for all topics on the server.</p> <p>Total -- The total size of outbound messages for all topics on the server, in kilobytes, since the server was started.</p>
Pending Messages	<p>Current -- The total number of messages for all topics on the server currently waiting to be processed.</p> <p>Bytes -- The total size of messages, in bytes, for all topics on the server currently waiting to be processed.</p>
Trend Graphs	Shows metrics for all topics on the selected server. <p>Pend Msgs -- Traces the total number of messages for all topics on the server currently waiting to be processed.</p> <p>In Msgs / sec -- Traces the number of inbound messages for all topics, per second.</p> <p>Out Msgs / sec -- Traces the number of outbound messages for all topics, per second.</p>


Log Scale


This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero



When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



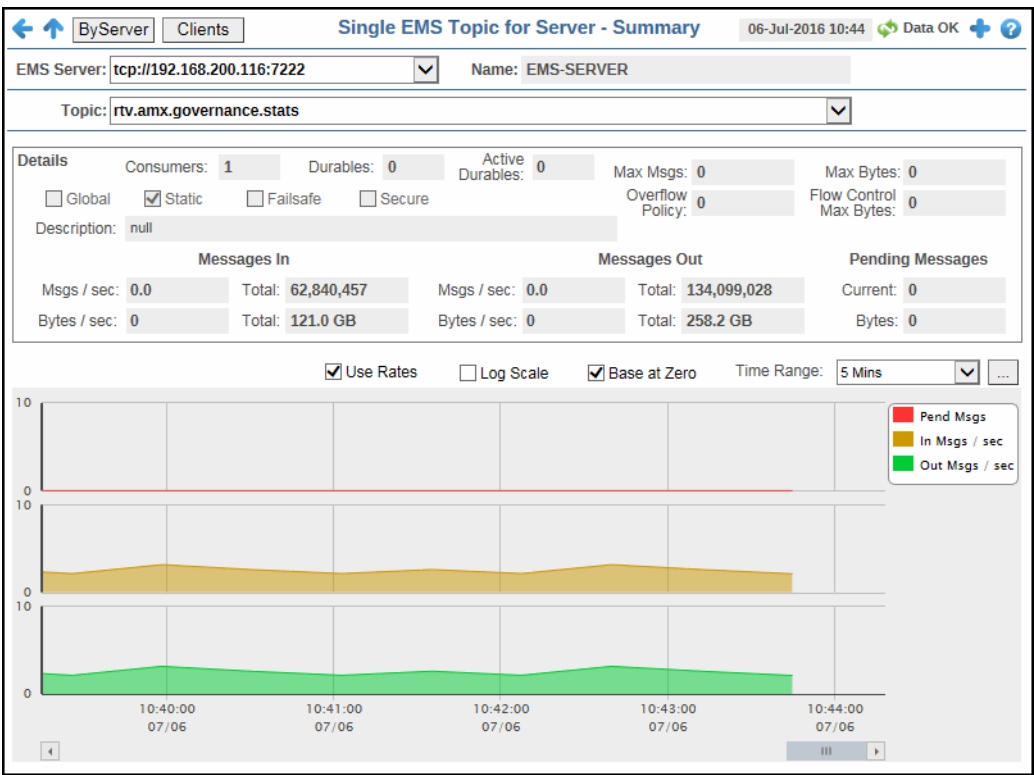
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Single Topic Summary

Track performance and utilization metrics for a single topic on a single server.



Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu

Table

 open commonly accessed displays.

6,047

 The number of items currently in the display.

Data OK

 Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Clicking **Clients** in the Title Bar takes you to the “Single EMS Topic-Clients” display for the selected topic.

Fields and Data

This display includes:

- EMS Server**

The EMS Server selected from this drop-down menu populates the Topics drop-down menu with the Topics belonging to this EMS Server.
- Name**

The name of the EMS server selected from the EMS Server drop-down menu.

Topic Select a Topic from the drop-down menu to view details for the selected Topic.

Browse Click to browse the contents of the selected topic in a separate window. The topic browser table displays up to 100,000 rows of messages.

Topic	Message ID	Time	Status
Topic: MyTopic	0		
Topic: MyTopic	1	5/10/2001 12:30	BROKEN
Topic: MyTopic	2	5/10/2001 12:30	BROKEN
Topic: MyTopic	3	5/10/2001 12:30	BROKEN
Topic: MyTopic	4	5/10/2001 12:30	BROKEN
Topic: MyTopic	5	5/10/2001 12:30	BROKEN
Topic: MyTopic	6	5/10/2001 12:30	BROKEN
Topic: MyTopic	7	5/10/2001 12:30	BROKEN
Topic: MyTopic	8	5/10/2001 12:30	BROKEN
Topic: MyTopic	9	5/10/2001 12:30	BROKEN
Topic: MyTopic	10	5/10/2001 12:30	BROKEN
Topic: MyTopic	11	5/10/2001 12:30	BROKEN
Topic: MyTopic	12	5/10/2001 12:30	BROKEN
Topic: MyTopic	13	5/10/2001 12:30	BROKEN
Topic: MyTopic	14	5/10/2001 12:30	BROKEN
Topic: MyTopic	15	5/10/2001 12:30	BROKEN
Topic: MyTopic	16	5/10/2001 12:30	BROKEN
Topic: MyTopic	17	5/10/2001 12:30	BROKEN
Topic: MyTopic	18	5/10/2001 12:30	BROKEN
Topic: MyTopic	19	5/10/2001 12:30	BROKEN
Topic: MyTopic	20	5/10/2001 12:30	BROKEN

By default, this button is disabled due to the fact that use of this option could significantly impact performance. To enable it, add the following substitution to the properties file with which you execute the Display Server and/or Viewer:

sl.rtview.sub=\$emsDestBrowseButtonVisFlag:1


Details Shows metrics for the topic selected from the Topic drop-down menu.


Consumers	The current number of consumers for the topic.
Durables	The number of durable subscribers (active and inactive) to the topic.
Active Durables	The number of active durable subscribers to the topic.
Max Msgs	The maximum number of messages allocated for the topic.
Max Bytes	The maximum of memory, in bytes, allocated for use by the topic.
Global	When checked, the message is global and is routed to other servers.
Static	When checked, the topic has a static destination.
Failsafe	When checked, the message is marked as failsafe delivery.
Secure	When checked, the topic is designated as secure and enforces permission policies.
Overflow Policy	Indicates whether an overflow policy is set for the topic: 0 = No policy is set. 1 = A policy is set.
Flow Control Max Bytes	The maximum amount of memory, in bytes, allocated for flow control use by the topic.
Description	Description of the Topic.
Messages In	
Msgs/sec	The number of inbound messages, per second, for the selected topic.
Total	The total number of inbound messages for the selected topic since the server was started.
Bytes/sec	The size of inbound messages, in bytes per second, for the selected topic.
Total	The total size of inbound messages, in bytes, for the selected topic since the server was started.
Messages Out	
Msgs/sec	The number of outbound messages, per second, for the selected topic.
Total	The total number of outbound messages for the selected topic since the server was started.
Bytes/sec	The size of outbound messages, in bytes per second, for the selected topic.

Pending Messages	Total	The total size of outbound messages, in bytes, for the selected topic since the server was started.
	Current	The number of messages for the selected topic currently waiting to be processed.
	Bytes	The size of the messages for the selected topic, in bytes, currently waiting to be processed.
Trend Graphs	Shows message data for the selected topic.	
	Pend Msgs	-- Traces the number of messages currently waiting to be processed.
	In Msgs / sec	-- Traces the number of inbound messages, per second. This trend graph only displays when Use Rates is selected.
	Out Msgs / sec	-- Traces the number of outbound messages, per second. This trend graph only displays when Use Rates is selected.
	Delta In Msgs	-- Traces the change in total inbound messages since the last update. This trend graph only displays when Use Rates is not selected.
	Delta Out Msgs	-- Traces the change in total inbound messages since the last update. This trend graph only displays when Use Rates is not selected.
	Use Rates	When this check box is selected, the inbound and outbound message rates (In Msgs/sec and Out Msgs/sec) display in the trend graph. When this check box is not selected, the delta inbound and outbound messages (Delta In Msgs and Delta Out Msgs) display in the trend graph.
	Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

Base at Zero When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Single EMS Topic-Clients

View data for all consumers and producers associated with the selected topic.

← ↑

ByServer

Summary

Single EMS Topic - Clients

06-Jul-2016 10:51

Data OK

+

?

EMS Server: tcp://192.168.200.116:7222

Name: EMS-SERVER

Topic: rtv.amx.governance.stats

Show Active Only

Producers

Count: 0

ID	clientID	Msgs / sec	Msgs Total	Bytes / sec	Total Bytes	userName	hos
----	----------	------------	------------	-------------	-------------	----------	-----

Consumers

Count: 1

ID	clientID	Msgs/sec	Msgs Total	Bytes/sec	Total Bytes	userName	hos
5614606		0.0	1,739,409	0.0	3,626,981,...	admin	SLHOST21

Title Bar (possible features are):

← ↑

Open the previous and upper display.

+

Open an instance of this display in a new window.

?

Open the online help page for this display.

Menu

Table

open commonly accessed displays.

6,047

The number of items currently in the display.

Data OK

Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Clicking **Summary** in the Title Bar takes you to the “Single Topic Summary” display. Clicking **ByServer** in the Title Bar takes you to the “Single Topic By Server” display.

Fields and Data

This display includes:

- EMS Server

The EMS Server selected from this drop-down menu populates the Topics drop-down menu with the Topics belonging to this EMS Server.
- Name

The name of the EMS Server selected from the EMS Server drop-down menu.

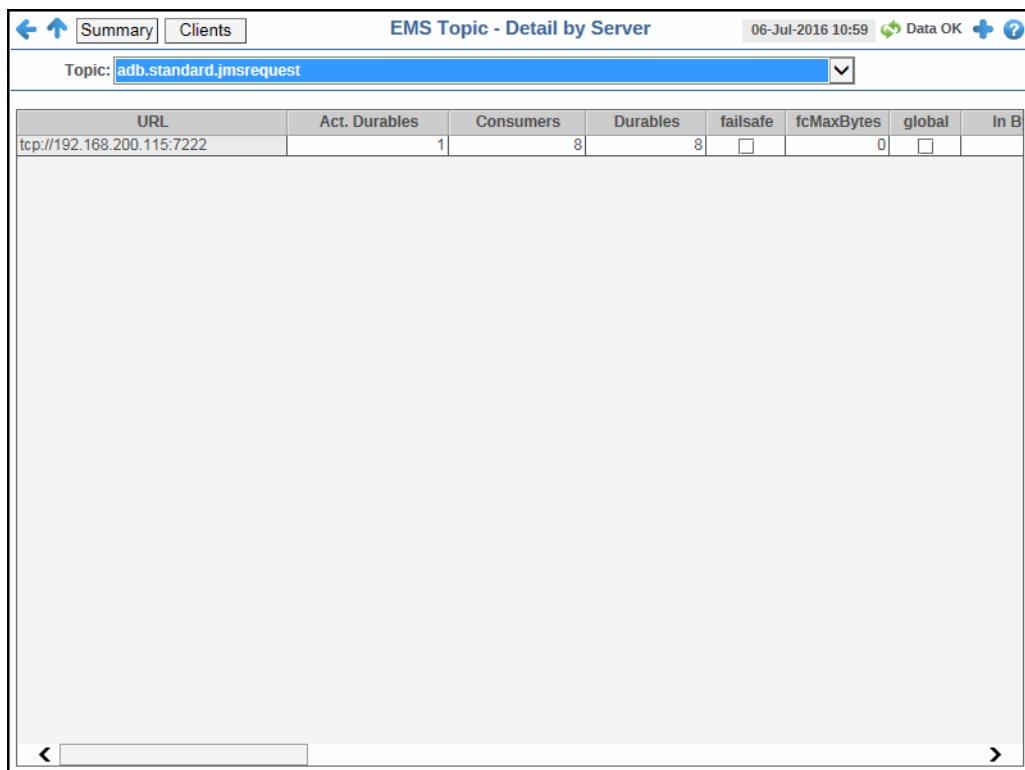
Topic	Select a Topic from the drop-down menu to view details for the selected Topic.
Show Active Only	Select this check box to view only the active producers and consumers for the selected Server/ Topic combination.
Producers	Shows data for all producers for the selected topic.
ID	A unique string identifier assigned to each producer.
clientID	A unique string identifier assigned to each client.
Msgs / sec	The number of messages, per second, emitted by the producer.
Msgs Total	The total number of messages emitted by the producer since the server was started.
Bytes / sec	The size of messages, in bytes per second, emitted by the producer.
Total Bytes	The total size of messages, in bytes, emitted by the producer since the server was started.
userName	The user name.
host	The name of the host.
sessionID	A unique string identifier assigned to each session.
connection ID	A unique string identifier assigned to each connection.
createTime	The amount of time, in milliseconds, since the producer was created.
time_stamp	The date and time this row of data was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Consumers	Shows data for all consumers of messages for the selected topic.
ID	A unique string identifier assigned to each consumer.
clientID	A unique string identifier assigned to each client.
Msgs / sec	The number of messages, per second, processed by the consumer.
Msgs Total	The total number of messages processed by the consumer.
Bytes / sec	The size of messages, in bytes per second, processed by the consumer.
Total Bytes	The total size of messages, in bytes, processed by the consumer since the server was started.
userName	The user name.
host	The name of the host machine.
Msgs Sent	The number of messages sent to the consumer that were not yet acknowledged by the consumer's session. The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.

Size Msg Sent	<p>The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Ack Msgs	<p>The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Sent Msgs	<p>The total number of messages sent to the consumer since the consumer was created.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Elap. Since Last Ack	<p>The amount of time (in milliseconds) that has elapsed since the last time a message sent to the consumer was acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Elap. Since Last Sent	<p>The amount of time (in milliseconds) that has elapsed since the last time the server sent a message to the consumer.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
destination Prefetch	<p>The actual destination prefetch value used by the server at runtime.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
prefetch Delivered Count	<p>The number of prefetch messages delivered to the consumer by the server. For consumers receiving messages on any destination with positive prefetch value, this value is never more than the prefetch value of the destination. This value cannot be used to identify the status of the consumer, but it can be used in conjunction with other consumer information values to identify consumers who stopped receiving messages due to application-specific problems.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
durable Name	<p>The name of the durable.</p>
routeName	<p>The queue owner server name if the consumer's destination is a routed queue.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
isActive	<p>When checked, the consumer is active and can receive messages from the server.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
isSystem	<p>This check box is checked if the consumer was automatically created by the system.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
sessionAck Mode	<p>Lists the consumer's session acknowledge mode as a constant defined in TibjmsAdmin.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
session ID	<p>A unique string identifier assigned to each session.</p>

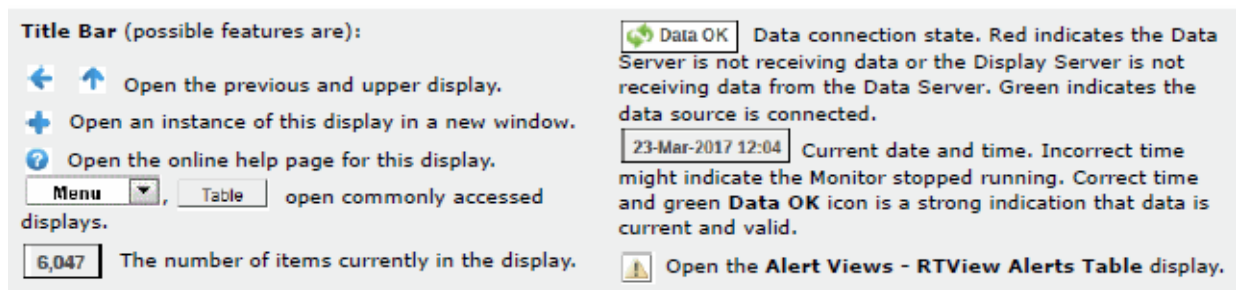
connection ID	A unique string identifier assigned to each connection.
createTime	The amount of time, in milliseconds, since the consumer was created.
time_stamp	The date and time this row of data was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

Single Topic By Server

Track performance and utilization metrics of a single topic across all servers that have the topic defined on it. Compare topic activity among servers.



EMS Topic - Detail by Server							
Topic: adb.standard.jmsrequest							
URL	Act. Durables	Consumers	Durables	failsafe	fcMaxBytes	global	In B
tcp://192.168.200.115:7222	1	8	8	<input type="checkbox"/>	0	<input type="checkbox"/>	



Note: Clicking **Clients** in the Title Bar takes you to the “Single EMS Topic-Clients” display for the selected topic. Clicking **Summary** in the Title Bar takes you to the “Single Topic Summary” display.

Fields and Data

This display includes:

Topic	The Topic selected from this drop-down menu populates this display.
Table	Shows details about the selected Topic for each server that has the Topic defined. Select a server from the list to view details in the “Single Topic Summary” display.
URL	The IP address and port number for the server.
Act. Durables	The number of currently active durables.
Consumers	The current number of consumers.
Durables	The number of active and inactive durables.
failsafe	When checked, the message is marked as failsafe delivery.
fcMaxBytes	The maximum number of bytes allocated for use by flow control.
global	When checked, the message is global and is routed to other servers.
In Byte Rate	The amount of inbound messages for the topic, in bytes per second.
In Msgs Rate	The amount of inbound messages for the topic, in number of messages per second.
In Total Bytes	The total number of inbound bytes for the topic.
In Total Msgs	The total number of inbound messages for the topic.
maxBytes	The maximum size, in bytes, that the topic can store for delivery to each durable or non-durable online subscriber on the topic.
maxMsgs	The maximum number of messages allocated for use by the topic.
Out Byte Rate	The amount of outbound messages (in bytes) per second.
Out Msg Rate	The number of outbound messages per second.
Out Total Bytes	The total amount of outbound messages for the topic, in bytes, since the server was started.

Out Total Msgs	The total number of outbound messages for the topic since the server was started.
overflowPolicy	Policy Indicates whether an overflow policy is set for the topic: 0 = No policy is set. 1 = A policy is set.
Pending Msgs	The number of currently pending messages for the topic.
Pending Msgs Size	The amount of space, in bytes, pending messages use for the topic.
secure	When checked, the topic is designated as secure and enforces permission policies.
static	When checked, the topic has a static destination.
Subscribers	The number of subscribers for the topic.
time_stamp	The date and time this row of data was last updated.
description	Descriptive text to help the administrator identify this resource.

EMS Queues

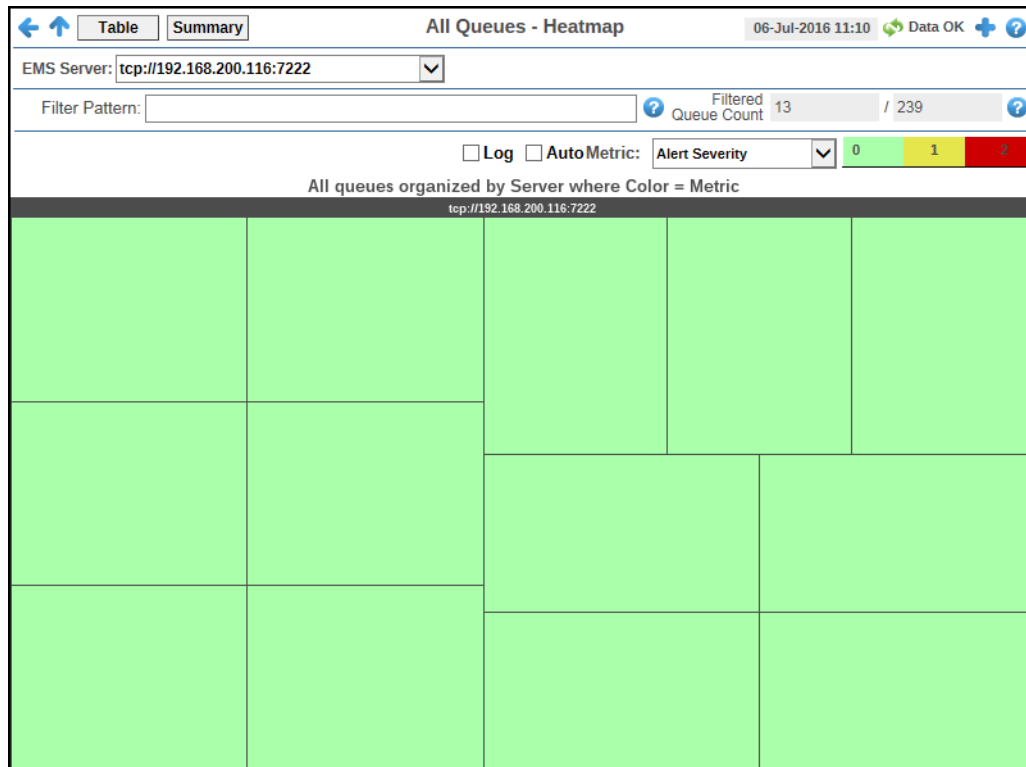
These displays present several views of performance metrics for queues. You can view all queues that are defined on a specific server in the [“All Queues Heatmap”](#) display, or you can view all servers that have a specific queue defined in the [“Single Queue Summary”](#) display. The [“Single EMS Queue-Clients”](#) display provides a list of all the servers on which those queues are defined.

- [“All Queues Heatmap”](#): A heatmap representation of a selected set of metrics that shows performance and utilization metrics and trends for all queues defined on a specified server, including message performance metrics.
- [“All Queues Table”](#): Shows performance and utilization metrics for all queues defined on a specified server.
- [“All Queues Summary”](#): Shows performance and utilization metrics and trends for all queues defined on a specified server, including message performance metrics.
- [“Single Queue Summary”](#): Shows detailed performance and utilization metrics and trends for a specified queue on a single server, including producer and consumer counts, and message performance metrics.
- [“Single EMS Queue-Clients”](#): View data for all consumers and producers associated with the selected queue.
- [“Single Queue By Server”](#): Table shows performance and utilization metrics for all servers that have a specified queue defined, including consumer and receiver count, and message performance metrics.

All Queues Heatmap

A heatmap representation of the [“All Queues Table”](#) display that allows you to track performance and utilization metrics and trends for all queues on a single server. View status and alerts of all queues for a server. Use the **Metric** drop-down menu to view to **Alert Severity**, **Alert Count**, **Consumers**, **Receivers**, **Pending Messages**, **Inbound Message Rate**, **Inbound Total Messages**, **Outbound Message Rate**, or **Outbound Total Messages**.

The heatmap is organized so that each rectangle represents a queue on the selected server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the “[Single Queue Summary](#)” display and view metrics for a particular queue. Toggle between the commonly accessed **Table** (link to the “[All Queues Table](#)” display) and **Heatmap** displays. Mouse-over rectangles to view more details about the performance and status of each queue.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Clicking **Table** in the Title Bar takes you to the “[All Queues Table](#)” display. Clicking **Summary** in the Title Bar takes you to the “[All Queues Summary](#)” display.

Fields and Data

This display includes:

EMS Server The EMS Server selected from this drop-down menu populates all the associated Queue data in this display.


Filter Pattern Enter a string to show only queues with names that contain the string. For example, if you enter the string Madrid, all queues with Madrid in the queue name are shown in the table. If no entry is made, all queue names are shown. For most use cases, you can enter a portion of the queue name.

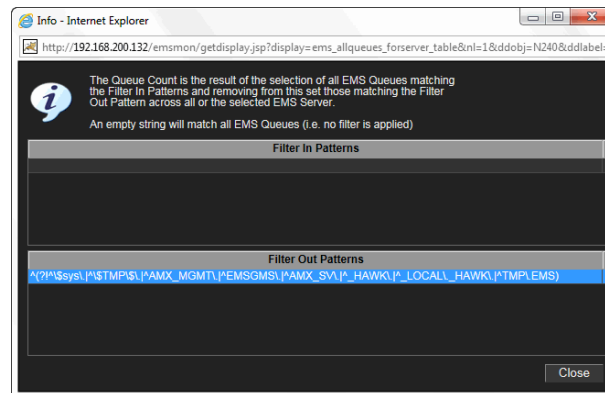
Filtered Queue Count This field is broken into two different values. The first value is the total number of currently active queues on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsQueueFilterOutPattern** property in the **emsmon/conf/rtvapm.properties** file. The second value is the total number of queues on the selected server. In other words, the filtered number of queues/the total number of queues on the server.

The default value for the **\$emsQueueFilterOutPattern** property is:

```
collector.sl.rtvew.sub=$emsQueueFilterOutPattern: '^(!^\\$sys\\.|^\\$TMP\\$\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\._HAWK\\.|^TMP\\.EMS)'
```

You can modify the filter value by editing the **\$emsQueueFilterOutPattern** property in the **“sample.properties File”**, which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Queue Count**.



Log This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

Auto When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

Metric Select the metric driving the heatmap display. The default is **Alert Severity**. Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the topics by server, where each rectangle represents a Queue. Mouse-over any rectangle to display the current values of the metrics for the Queue. Click on a rectangle to drill-down to the associated **“Single Queue Summary”** display for a detailed view of metrics for that particular queue.

Alert Severity

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.


-- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

1 -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.


-- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

Alert Count

The total number of alarm and warning alerts in a given item (index) associated with the rectangle.


The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Consumers

The total number of consumers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

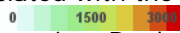
The **Auto** option does not impact this metric.

Receivers

The total number of receivers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

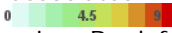
The **Auto** option does not impact this metric.

Pending Msgs

The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsQueuesPendingMsgsHigh**, which is **3000**. The middle value in the gradient bar indicates the middle value of the range (the default is **1500**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.


In Msgs /sec

The number of inbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsQueuesInMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.


Note: This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

In Total Msg

The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.


Out Msgs/sec

The number of outbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsQueuesOutMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Note: This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

Out Total Msgs

The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

All Queues Table

Track performance and utilization metrics for all queues on a single server.

Queue Name	URL	In Rate	In Total	Out Rate	Out Total	Pend Msgs	P
amx.governance.internal.stats	tcp://192.1...	0	206,708,124	0	206,708,124	103	
amx.governance.stats	tcp://192.1...	0	62,845,363	0	62,845,367	0	
cl_logservice_queue	tcp://192.1...	0	0	0	0	0	
cl_payload_queue	tcp://192.1...	0	0	0	0	0	
com.tibco.amf.admin.deploymentServerQueue.inst...	tcp://192.1...	0	4	0	4	0	
com.tibco.amf.admin.deploymentServerQueue.inst...	tcp://192.1...	0	9	0	9	0	
com.tibco.amf.admin.deploymentServerQueue.SL...	tcp://192.1...	0	0	0	0	0	
com.tibco.amf.admin.deploymentServerQueue.SL...	tcp://192.1...	0	11	0	11	0	
com.tibco.amf.admin.deploymentServerQueue.Sys...	tcp://192.1...	0	0	0	0	0	
com.tibco.amf.admin.deploymentServerQueue.Sys...	tcp://192.1...	0	8	0	8	0	
com.tibco.amf.admin.deploymentServerQueue.Sys...	tcp://192.1...	0	0	0	0	0	
queue.sample	tcp://192.1...	0	0	0	0	0	
sample	tcp://192.1...	0	0	0	0	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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Clicking **Heatmap** in the Title Bar takes you to the "All Queues Heatmap" display. Clicking **Summary** in the Title Bar takes you to the "All Queues Summary" display.

Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Queue data in this display.

Filter Pattern

Enter a string to show only queues with names that contain the string. For example, if you enter the string Madrid, all queues with Madrid in the queue name are shown in the table. If no entry is made, all queue names are shown. For most use cases, you can enter a portion of the queue name.


Filtered Queue Count

This field is broken into two different values. The first value is the total number of currently active queues on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsQueueFilterOutPattern** property in the **emsmon/conf/rtvapm.properties** file. The second value is the total number of queues on the selected server. In other words, the filtered number of queues/the total number of queues on the server.

The default value for the **\$emsQueueFilterOutPattern** property is:

```
collector.sl.rtvview.sub=$emsQueueFilterOutPattern: '^(?!^\\$sys\\.|^\\$TMP\\$\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\.|^_HAWK\\.|^_TMP\\.EMS)'
```

You can modify the filter value by editing the **\$emsQueueFilterOutPattern** property in the **"sample.properties File"**, which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Queue Count**.

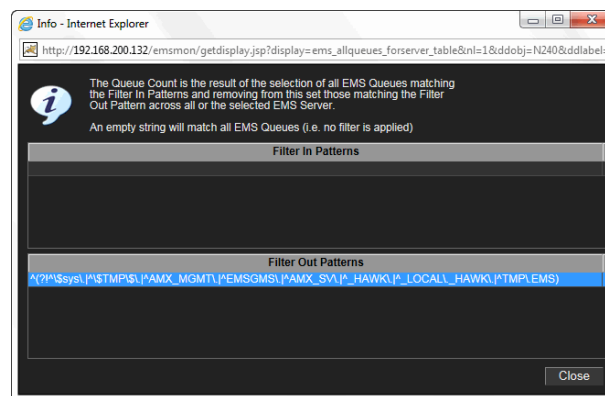


Table This table describes all queues on the selected server. Click a row to view metrics for a single queue in the **"Single Queue Summary"** display.

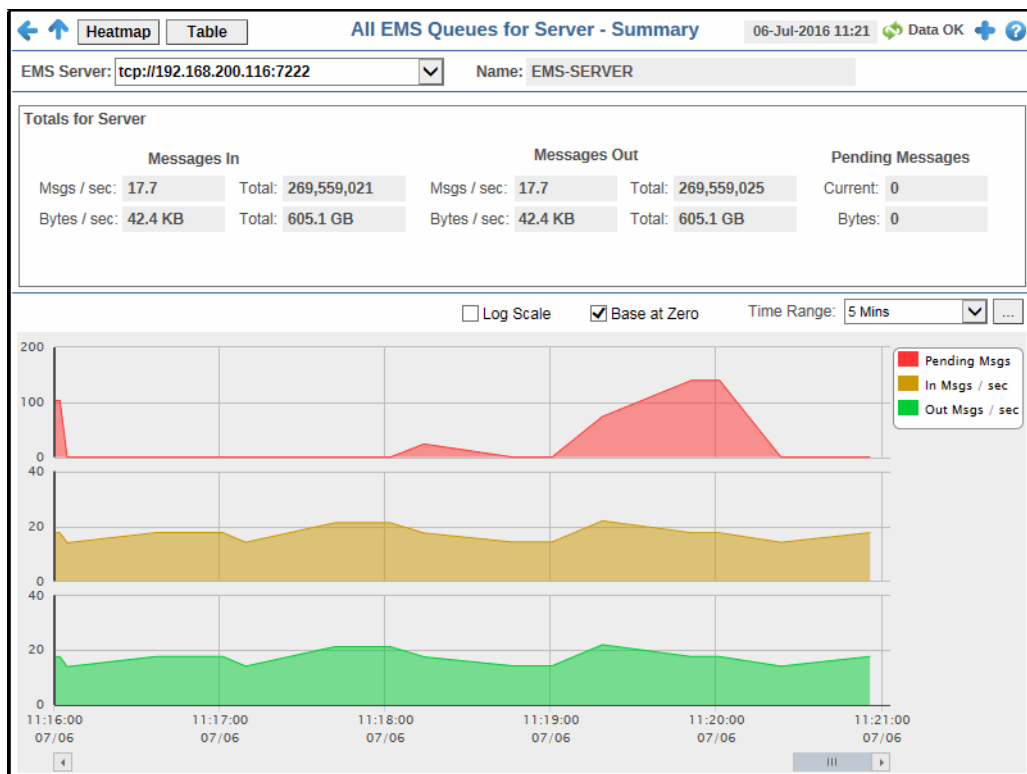
Queue Name	The name of the queue.
URL	The IP address and port number for the server.
In Rate	The number of inbound messages for the queue, per second. Note: This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.
In Total	The total number of inbound messages for the queue.
Out Rate	The number of outbound messages for the queue, per second. Note: This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.
Out Total	The total number of outbound messages for the queue.
Pend Msgs	The number of currently pending messages for the queue.

Pend Size	The amount of space, in bytes, used by pending messages for the queue.
activeDurableCount	The current number of active durables.
consumerCount	The number of active and inactive consumers.
durableCount	The number of active and inactive durables.
failSafe	When checked, the message is marked as failsafe delivery.
fcMaxBytes	The maximum number of bytes allocated for use by flow control.
global	When checked, the message is global and is routed to other servers.
inboundByteRate	The amount of inbound messages for the queue, in bytes per second.
inboundTotalBytes	The total amount of inbound messages for the queue, in bytes.
maxBytes	The maximum amount of bytes allocated for use by the queue.
maxMsgs	The maximum number of messages allocated for use by the queue.
outboundByteRate	The amount of outbound messages for the queue, in bytes per second.
outboundTotalBytes	The total amount of outbound messages for the queue, in bytes.
overflowPolicy	Indicates whether an overflow policy is set for the queue: 0 = No policy is set. 1 = A policy is set.
secure	When checked, the queue is designated as secure and enforces permission policies.
static	When checked, the queue has a static destination.
subscriberCount	The number of subscribers that receive queue message.
description	Descriptive text to help the administrator identify this resource.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

time_stamp	The date and time this row of data was last updated.
DeltaInboundTotalMessages	The change in total inbound messages since the last update.
DeltaInboundTotalBytes	The change in total inbound message bytes since the last update.
DeltaOutboundTotalMessages	The change in total outbound messages since the last update.
DeltaOutboundTotalBytes	The change in total outbound message bytes since the last update.
prefetch	Lists the maximum number of messages consumers can fetch.
expiryOverride	If set to a non-zero value for a destination and the server delivers a message to the destination, the server replaces the producer's expiration value with this value.
store	Provides the store for this destination where persistent messages are stored.
deliveredMessageCount	Indicates the total number of messages that have been delivered and acknowledged.
URLQueue	The IP address and port for the queue.
exclusive	When checked, the server sends all messages on this queue to one consumer.
maxRedelivery	The maximum number of attempts for attempting redelivery of a message.
receiverCount	The number of receivers that receive queue message.

All Queues Summary

Track performance and utilization metrics and trends for all queues on a single server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.


Note: Clicking **Heatmap** in the Title Bar takes you to the "All Queues Heatmap" display. Clicking **Table** in the Title Bar takes you to the "All Queues Table" display.


Fields and Data


This display includes:


- EMS Server** The EMS Server selected from this drop-down menu populates all associated queue data in this display.
- Name** The name of the server selected in the **EMS Server** drop down list.
- Totals For Server** Shows metrics for all queues on the selected server.

Messages In	Msgs/sec -- The total number of inbound messages for all queues on the server, per second.
	Total -- The total number of inbound messages for all queues on the server since the server was started.
	Bytes/sec -- The amount of inbound messages, in bytes per second, for all queues on the server.
	Total -- The amount of inbound messages, in kilobytes, for all queues on the server since the server was started.
Messages Out	Msgs/sec -- The total number of outbound messages for all queues on the server, per second.
	Total -- The total number of outbound messages for all queues on the server since the server was started.
	Bytes/sec -- The amount of outbound messages, in bytes per second, for all queues on the server.
	Total -- The amount of outbound messages for all queues on the server, in kilobytes, since the server was started.
Pending Messages	Current -- The total number of messages currently waiting to be processed.
	Bytes -- The amount of messages, in bytes, currently waiting to be processed.
Trend Graphs	<p>Shows metrics for all queues on the selected server.</p> <p>Pending Msgs -- Traces the number of messages currently waiting to be processed.</p> <p>In Msgs / sec -- Traces the number of inbound messages for all queues, per second.</p> <p>Out Msgs / sec -- Traces the number of outbound messages for all queues, per second.</p>
Log Scale	<p>This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.</p>

- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



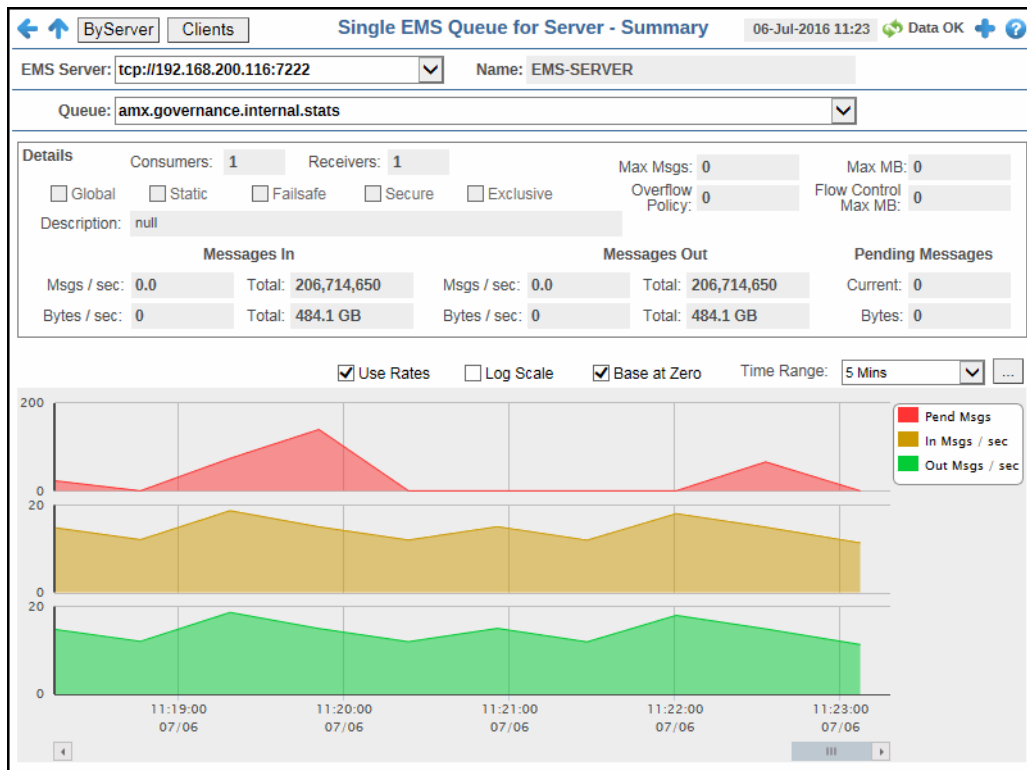
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Single Queue Summary

Track performance and utilization metrics for a single queue on a single server.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Clicking **Clients** in the Title Bar takes you to the "Single EMS Queue-Clients" display. Clicking **By Server** in the Title Bar takes you to the "Single Queue By Server".

Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates the **Queues** drop-down menu with the queues belonging to this EMS Server.
- Name** The name of the EMS Server selected from the **EMS Server** drop-down menu.

- Queue** Select a queue from the drop-down menu. The selection made here populates this display.
- Browse** Click to browse the contents of the selected queue in a separate window. The queue browser table displays up to 100,000 rows of messages.

Message ID	Queue	Destination
Msg 1748	008	075
Msg 1749	024	052
Msg 1750	044	034
Msg 1751	014	001
Msg 1752	012	018
Msg 1753	002	007
Msg 1754	012	018
Msg 1755	002	007
Msg 1756	012	018
Msg 1757	002	007
Msg 1758	012	018
Msg 1759	002	007
Msg 1760	012	018
Msg 1761	002	007
Msg 1762	012	018
Msg 1763	002	007
Msg 1764	012	018
Msg 1765	002	007
Msg 1766	012	018
Msg 1767	002	007
Msg 1768	012	018
Msg 1769	002	007
Msg 1770	012	018

By default, this button is disabled due to the fact that use of this option could significantly impact performance. To enable it, add the following substitution to the properties file with which you execute the Display Server and/or Viewer:

sl.rtview.sub=\$emsDestBrowseButtonVisFlag:1

Details Shows metrics for the queue selected from the **Queue** drop-down menu.

Consumers The number of consumers currently interacting with the queue.

Receivers The number of consumers currently receiving messages from the queue.

Max Msgs The maximum number of messages allocated for the queue.

Max MB The maximum amount of memory, in megabytes, allocated for use by the queue.

Global When checked, the message is global and is routed to other servers.

Static When checked, the queue has a static destination.

Failsafe When checked, the message is marked as failsafe delivery.

Secure When checked, the queue is designated as secure and enforces permission policies.

Exclusive When checked, the server sends all messages on this queue to one consumer.

Overflow Policy Indicates whether an overflow policy is set for the queue:
0 = No policy is set.
1 = A policy is set.

Flow Control Max MB The maximum amount of memory, in megabytes, allocated for flow control use by the queue.

Description Description of the Queue.

Messages In **Msgs/sec** The number of inbound messages, per second, for the selected queue.


Total The total number of inbound messages for the selected queue since the server was started.


Bytes/sec The size of the inbound messages, in bytes per second, for the selected queue.


Total The total size of inbound messages, in bytes, for the selected queue since the server was started.

Messages Out	Msgs/sec	The number of outbound messages, per second, for the selected queue.
	Total	The total number of outbound messages for the selected queue since the server was started.
	Bytes/sec	The size of outbound messages, in bytes per second, for the selected queue.
	Total	The total size of outbound messages, in bytes, for the selected queue since the server was started.
Pending Messages	Current	The total number of messages for the selected queue currently waiting to be processed.
	Bytes	The size, in bytes, of messages for the selected queue currently waiting to be processed.
Trend Graphs	Shows metrics for the selected queue on the specified server.	
	Pending Msgs	-- Traces the number of messages currently waiting to be processed.
	In Msgs / sec	-- Traces the number of inbound messages, per second. This trend graph only displays when Use Rates is selected.
	Out Msgs / sec	-- Traces the number of outbound messages, per second. This trend graph only displays when Use Rates is selected.
	Delta In Msgs	-- Traces the change in total inbound messages since the last update. This trend graph only displays when Use Rates is not selected.
	Delta Out Msgs	-- Traces the change in total inbound messages since the last update. This trend graph only displays when Use Rates is selected.
	Use Rates	When this check box is selected, the inbound and outbound message rates (In Msgs/sec and Out Msgs/sec) display in the trend graph. When this check box is not selected, the delta inbound and outbound messages (Delta In Msgs and Delta Out Msgs) display in the trend graph.
	Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

Base at Zero When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

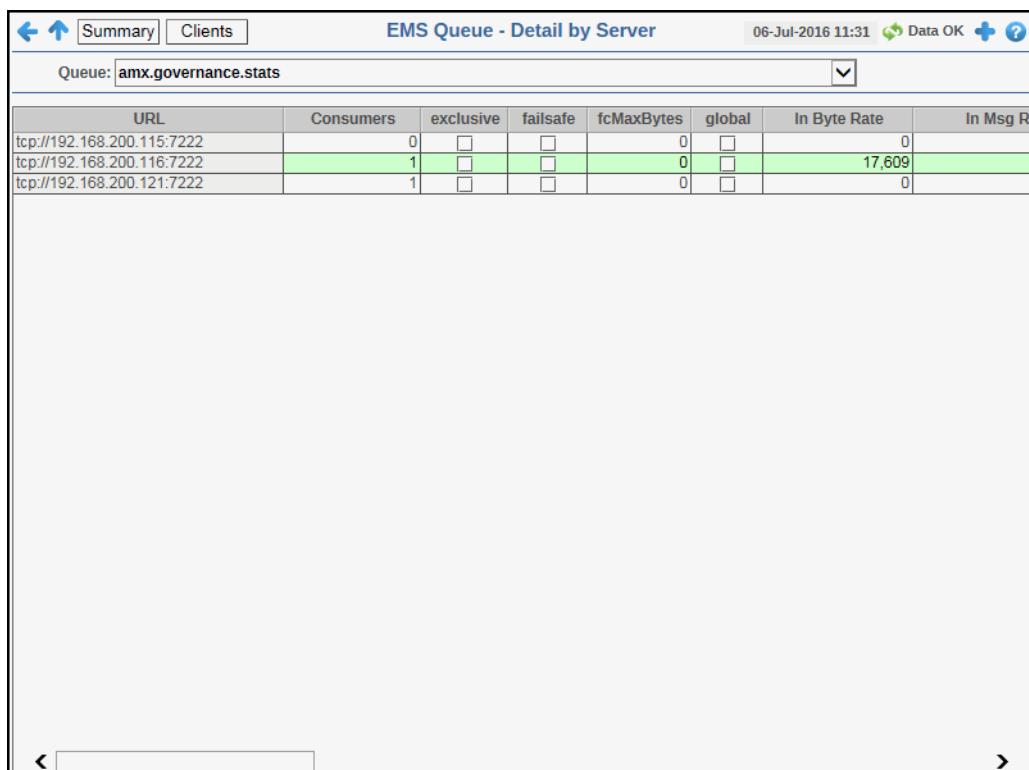
Queue	Select a Queue from the drop-down menu to view details for the selected Queue.
Show Active Only	Select this check box to view only the active producers and consumers for the selected EMS Queue.
Producers	Shows data for all producers for the selected queue.
ID	A unique string identifier assigned to each producer.
clientID	A unique string identifier assigned to each client.
Msgs / sec	The number of messages, per second, that are emitted by the producer.
Msgs Total	The total number of messages emitted by the producer since the server was started.
Bytes / sec	The size of messages, in bytes per second, that are emitted by the producer.
Total Bytes	The total size of messages, in bytes, emitted by the producer since the server was started.
userName	The user name.
host	The name of the host.
sessionID	A unique string identifier assigned to each session.
connection ID	A unique string identifier assigned to each connection.
createTime	The amount of time, in milliseconds, since the producer was created.
time_stamp	The date and time this row of data was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Consumers	Shows data for all consumers associated with the selected queue.
ID	A unique string identifier assigned to each consumer.
clientID	A unique string identifier assigned to each client.
Msgs / sec	The number of messages, per second, that are processed by the consumer.
Msgs Total	The total number of messages that have been processed by the consumer.
Bytes / sec	The size of messages, in bytes per second, that are processed by the consumer.
Total Bytes	The total size of messages, in bytes, processed by the consumer since the server was started.
userName	The user name.
host	The name of the host machine.

Msgs Sentt	<p>The number of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Size Msg Sent	<p>The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Ack Msgs	<p>The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Sent Msgs	<p>The total number of messages sent to the consumer since the consumer was created.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Elap. Since Last Ack	<p>The amount of time (in milliseconds) that has elapsed since the last time a message sent to the consumer was acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Elap. Since Last Sent	<p>The amount of time (in milliseconds) that has elapsed since the last time the server sent a message to the consumer.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
destination Prefetch	<p>The actual destination prefetch value used by the server at runtime.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
prefetch Delivered Count	<p>The number of prefetch messages delivered to the consumer by the server. For consumers receiving messages on any destination with positive prefetch value, this value is never more than the prefetch value of the destination. This value cannot be used to identify the status of the consumer, but it can be used in conjunction with other consumer information values to identify consumers who stopped receiving messages due to application-specific problems.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
durable Name	<p>The name of the durable.</p>
routeName	<p>The queue owner server name if the consumer's destination is a routed queue.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
isActive	<p>When checked, the consumer is active and can receive messages from the server.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
isSystem	<p>This check box is checked if the consumer was automatically created by the system.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>


sessionAck Mode	Lists the consumer's session acknowledge mode as a constant defined in TibjmsAdmin . The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.
session ID	A unique string identifier assigned to each session.
connection ID	A unique string identifier assigned to each connection.
createTime	The amount of time, in milliseconds, since the consumer was created.
time_stamp	The date and time this row of data was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

Single Queue By Server

Track performance and utilization metrics of a single queue across all servers. Compare queue activity among servers.



URL	Consumers	exclusive	failsafe	fcMaxBytes	global	In Byte Rate	In Msg R
tcp://192.168.200.115:7222	0	<input type="checkbox"/>	<input type="checkbox"/>	0	<input type="checkbox"/>	0	
tcp://192.168.200.116:7222	1	<input type="checkbox"/>	<input type="checkbox"/>	0	<input type="checkbox"/>	17,609	
tcp://192.168.200.121:7222	1	<input type="checkbox"/>	<input type="checkbox"/>	0	<input type="checkbox"/>	0	



Note: Clicking **Summary** in the Title Bar takes you to the [“Single Queue Summary”](#). Clicking **Clients** in the Title Bar takes you to the [“Single EMS Queue-Clients”](#) display.

Fields and Data

This display includes:

Queue	The Queue selected from this drop-down menu populates this display.
Table	Shows details about the selected Queue for each server that has the queue defined. Select a server to view details in the “Single Queue Summary” display.
URL	The URL of the server.
Consumers	The number of active and inactive consumers.
exclusive	When checked, the server sends all messages on this queue to one consumer.
failSafe	When checked, the message is marked as failsafe delivery.
fcMaxBytes	The maximum number of bytes allocated for use by flow control.
global	When checked, the message is global and is routed to other servers.
In Byte Rate	The amount of inbound messages for the queue, in bytes per second.
In Msg Rate	The amount of inbound messages for the queue, in number of messages per second.
In Total Bytes	The total number of inbound bytes for the queue.
In Total Msgs	The total number of inbound messages for the queue.
maxBytes	The maximum amount of bytes allocated for use by the queue.
maxMsgs	The maximum number of messages allocated for use by the queue.
maxRedelivery	The maximum number of attempts for attempting redelivery of a message.

Out Byte Rate	The amount of outbound messages (in bytes) per second.
Out Msg Rate	The number of outbound messages per second.
Out Total Bytes	The total amount of outbound messages, in bytes, since the server was started.
Out Total Msgs	The total number of outbound messages since the server was started.
overflowPolicy	Indicates whether an overflow policy is set for the queue: 0 = No policy is set. 1 = A policy is set.
Pending Msgs	The number of currently pending messages.
Pending Msgs Size	The amount of space, in bytes, pending messages use for the queue.
Receivers	The number of receivers of queue messages.
secure	When checked, the topic is designated as secure and enforces permission policies.
static	When checked, the topic has a static destination.
time_stamp	The date and time this row of data was last updated.
description	Descriptive text to help the administrator identify this resource.

EMS Clients

These displays present performance metrics for all server connections, including users, routes between servers, producers, consumers and durables connected to a specific EMS server.

- ["Connections"](#): Shows connection information on a single server.
- ["Bridges, Users, Ports"](#): Shows utilization metrics for bridges, users and ports on a single server.
- ["Routes"](#): Shows bridges for server routes on a single server.
- ["Producers"](#): Shows utilization metrics for producers on a single server.
- ["Producer Summary"](#): Shows utilization metrics for producers on a single server.
- ["Consumers"](#): Shows utilization metrics for consumers on a single server.
- ["Consumer Summary"](#): Shows utilization metrics for consumers on a single server.
- ["Durables"](#): Shows utilization metrics for durables on a single server.

Connections

View connections on a single server.

← ↑ **EMS Connections for Server** 06-Jul-2016 11:40 Data OK + ?

EMS Server: **tcp://192.168.200.116:7222** Name: **EMS-SERVER** ☒ Show Active Only

Client ID Filter: ? Filtered Connection Count 85 / 85 ?

User Filter: ?

Conn ID	Client ID	Conn URL	User	host	type	consumerCount	producerCount
1181372		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181373	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181375		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181376	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181378		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181379	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181388		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181396		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181398		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181399	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181402		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181403		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181405	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181407		[admin@SLHOST16]	admin	SLHOST16	QUEUE	0	
1181410		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181415		[admin@SLHOST16]	admin	SLHOST16	QUEUE	0	
1181424	AMX_MGMT.DevEn...	[admin@SLHOST16]	admin	SLHOST16	CONN...	4	
1181427	AMX_SV:8bf3d299-...	[admin@SLHOST16]	admin	SLHOST16	CONN...	37	
1181590		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181591	EMSGMS._amxadm...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181593		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181594	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181596		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181597	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181599		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181600	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181603		[admin@SLHOST16]	admin	SLHOST16	CONN...	7	
1181609	AMX_MGMT.Syste...	[admin@SLHOST16]	admin	SLHOST16	CONN...	5	

Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Connections data in this display.
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.
- Show Active Only** Select this check box to display only active connections.
- Client ID Filter** Filter field that allows you to filter the list of connections by client ID.


Filtered Connection Count


This field is broken into two different values. The first value is the total number of currently active connections on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsConnectionFilterOutPattern** property in the **emsmon/conf/rtvapm.properties** file. The second value is the total number of connections on the selected server. In other words, the filtered number of connections/the total number of connections on the server.

The default value for the **\$emsConnectionFilterOutPattern** property is:

```
collector.sl.rtvew.sub=$emsConnectionFilterOutPattern: '^(?!^[admin\\@])'
```

You can modify the filter value by editing the **\$emsConnectionFilterOutPattern** property in the ["sample.properties File"](#), which will override the default value.


Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Connection Count**.



User Filter Filter field that allows you to filter the list of connections by user name.

Connections This table describes the current connections on the selected server.

Conn ID	The unique numeric ID assigned to this connection that can be used for deletion.
Client ID	The unique string identifier assigned to the client.
Conn URL	The connection URL.
User	The user name.
host	The name of the host to which the server is connected.
type	The type of connection: Queue, Topic or System.
consumerCount	The total number of consumers currently connected.
producerCount	The total number of producers currently connected.
sessionCount	The total number of sessions currently connected.
startTime	The date and time the server was started
upTime	The amount of time, in milliseconds, since the server was started.



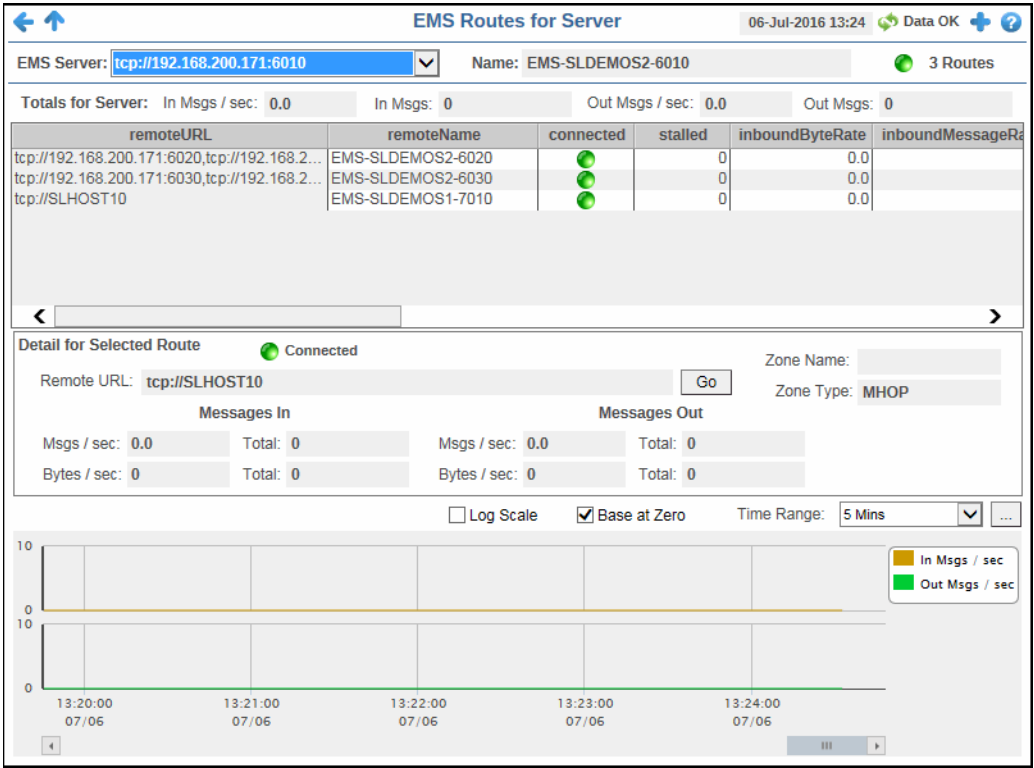
Fields and Data

This display includes:

EMS Server	The EMS Server selected from this drop-down menu populates all associated Bridges, Users, and Ports data in this display.								
Name	The name of the EMS Server selected from the EMS Server drop-down menu.								
Bridges	This table describes the bridges for the selected server. <table> <tr> <td>source</td><td>The topic or queue which is the source of the bridge.</td></tr> <tr> <td>target</td><td>The topic or queue which is the target of the bridge.</td></tr> <tr> <td>selector</td><td>The message selector string or blank if none has been set.</td></tr> <tr> <td>Expired</td><td>When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</td></tr> </table>	source	The topic or queue which is the source of the bridge.	target	The topic or queue which is the target of the bridge.	selector	The message selector string or blank if none has been set.	Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
source	The topic or queue which is the source of the bridge.								
target	The topic or queue which is the target of the bridge.								
selector	The message selector string or blank if none has been set.								
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.								
Users	This table describes the users on the selected server. <table> <tr> <td>name</td><td>The name of the connected user.</td></tr> <tr> <td>external</td><td>When checked, the user is defined in an external system.</td></tr> <tr> <td>description</td><td>Textual description of the user.</td></tr> </table>	name	The name of the connected user.	external	When checked, the user is defined in an external system.	description	Textual description of the user.		
name	The name of the connected user.								
external	When checked, the user is defined in an external system.								
description	Textual description of the user.								
Listen Ports	This table describes the connections the selected server is to listen for. <table> <tr> <td>port</td><td>The IP address and port number on which the server is to listen for connections.</td></tr> <tr> <td>URL</td><td>The URL on which the server is to listen for connections.</td></tr> </table>	port	The IP address and port number on which the server is to listen for connections.	URL	The URL on which the server is to listen for connections.				
port	The IP address and port number on which the server is to listen for connections.								
URL	The URL on which the server is to listen for connections.								

Routes

Track utilization metrics for server routes on a single server. Inbound metrics, such as **inboundByteRate**, indicate an in route to the server. Outbound metrics, such as **outboundByteRate**, indicate an out route to the server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.




23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Routes data in this display.
- Name** The name of the EMS server selected from the **EMS Server** drop-down menu.
- Routes** The number of server routes and the connection state.

-  -- One or more routes for this server are disconnected.
-  -- All routes for this server are connected.
-  -- There are no routes for this server.

Totals For Server




Shows metrics for all server routes on the selected server.

In Msgs / sec	The number of inbound messages, per second.
In Msgs	The total number of inbound messages.
Out Msgs / sec	The number of outbound messages, per second.
Out Msgs	The total number of outbound messages.

Table

This table shows metrics for each server route on the selected server. Select a route to view details.


remoteURL	The URL of the remote server.
remoteName	The name of the remote server.
connected	When checked, the server route is connected.
stalled	Indicates whether the IO flow stalled on the route. A value of 0 (zero) = not stalled. A value of 1 = stalled.
inboundByteRate	The rate of inbound data in bytes, per second.
inboundMessageRate	The rate of inbound messages in number of messages per second.
inboundTotalBytes	The total number of inbound bytes.
inboundTotalMessages	The total number of inbound messages.
outboundByteRate	The rate of outbound data in bytes per second.
outboundMessageRate	The rate of outbound messages in number of messages per second.
outboundTotalBytes	The total number of outbound bytes.
outboundTotalMessages	The total number of outbound messages.
zoneName	The name of the zone for the route.
zoneType	Indicates a multi-hop or one-hop zone.
active	Indicates whether the server route is currently transferring data: 1 = true 0 = false


Detail for Selected Route	inactive	Indicates whether the server route is currently transferring data: 1 = true 0 = false
	suspended	Indicates whether outbound messages to the route have been suspended: 1 = true 0 = false
	remoteURLName	The IP address and name for the remote connection.
	Shows metrics for the server route selected from the table.	
	Connected	The server route connection state.  -- The server route is disconnected  -- The server route is connected.
	Zone Name	The name of the zone for the route.
	Remote URL	The IP address and port number for the server route connection. Click the  button to open the selected server in the EMS Single Server Summary display.
	Zone Type	Indicates a multi-hop or one-hop zone.
	Messages In	Msgs/sec -- The number of inbound messages, per second. Total -- The total number of inbound messages since the connection was established. Bytes/sec -- The amount of inbound messages, in bytes per second, for this server route. Total -- The amount of inbound messages, in kilobytes, for this server route since the connection was established.
	Messages Out	Msgs/sec -- The number of outbound messages, per second. Total -- The total number of outbound messages since the connection was established. Bytes/sec -- The amount of outbound messages, in bytes per second. Total -- The amount of outbound messages, in kilobytes, since the connection was established.
Trend Graphs	Shows message data for the selected route.	
	In Msgs / sec -- Traces the number of inbound messages, per second.	
	Out Msgs / sec -- Traces the number of outbound messages, per second.	
	Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero


When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Producers

Track utilization metrics for producers on a single server.

EMS Producers for Server					
06-Jul-2016 13:27 Data OK					
EMS Server: tcp://192.168.200.171:6010		Name: EMS-SLDEMOS2-6010		8 Producers	
Client ID Filter:		DestName Filter:		<input checked="" type="checkbox"/> Show Active Only	
Count: 8	Msgs / sec: 4.0	Total Msgs: 6,213,157	Bytes / sec: 800	Total Bytes: 1,222.9 MB	
ID	Client ID	Destination	Msgs / sec	Total Msgs	By
2007		\$sys.admin	0.0	3,302,620	
2049		\$sys.admin	0.0	2,196,187	
2055		\$sys.admin	0.0	276,789	
2121		\$sys.admin	0.0	348,891	
2136		\$sys.admin	4.0	43,490	
2138		\$sys.admin	0.0	22,003	
2139		\$sys.admin	0.0	11,948	
2140		\$sys.admin	0.0	11,229	

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Note: Clicking on a row in the Producers table and then clicking the Dest. button in the Title Bar takes you to the ["Single Queue Summary"](#) display for the selected producer.

Fields and Data


This display includes:

- EMS Server** The EMS Server selected from this drop-down list displays a list of the currently connected Producers.
- Name** The name of the EMS server selected from the **EMS Server** drop-down menu.

Producers	The number of currently connected producers on the server.
Client ID Filter	Filter field that allows you to filter the list of producers by client ID.
DestName Filter	Filter field that allows you to filter the list of producers by destination name.
Show Active Only	Select this check box to display only active producers.
Count	The number of currently connected producers on the server.
Msgs / sec	The number of messages, per second, for the producer.
Total Msgs	The total number of messages for the producer.
Bytes / sec	The amount of messages, in bytes per second, for the producer.
Total Bytes	The total size of messages, in bytes, for the producer.
Table	This table shows metrics for each producer on the selected server. Double-clicking on a row in the Producers table displays details for the producer in the "Producer Summary" drill-down display.
ID	A unique string identifier assigned to each producer.
Client ID	A unique string identifier assigned to each client.
Destination	The name of the destination.
Msgs / sec	The number of messages, per second, for the producer.
Total Msgs	The total number of messages for the producer.
Bytes / sec	The size of messages, in bytes per second, for the producer.
Total Bytes	The total size of messages, in bytes, for the producer.
User	The user name.
Host	The name of the host.
sessionID	A unique string identifier assigned to each session.
ConnID	A unique string identifier assigned to each connection.
createTime	The amount of time, in milliseconds, since the producer was created.
time_stamp	The date and time this row of data was last updated.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
destinationType	The configured destination type.

Producer Summary

Displays details for an individual producer. You can access this display by double-clicking on a producer in the “Producers” 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.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

- EMS Server** The selected EMS Server populates the Producer ID/ Client ID drop-down menu with associated Producer IDs/Client IDs. This drop down list defaults to the EMS Server that was selected on the previous display.
- Name** The name of the EMS server selected from the **EMS Server** drop-down menu.
- Producer ID/ Client ID** Drop-down menu containing the Producer IDs/Client IDs. This drop down list defaults to the Producer ID/Client ID that was selected on the previous display.

**Detail for
Selected
Producer**

Shows metrics for the producer selected from the table.

Dest Name	The name of the destination.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Dest Type	The configured destination type.
Producer ID	A unique string identifier assigned to each producer.
Conn ID	A unique string identifier assigned to each connection.
Session ID	A unique string identifier assigned to each session.
Creation Time	The amount of time, in milliseconds, since the producer was created.
User	The user name.
Client ID	A unique string identifier assigned to each client.
Msgs/sec	The number of messages, per second, for the producer.
Total Msgs	The total number of messages for the producer.
Bytes/sec	The size of messages, in bytes per second, for the producer.
Total Bytes	The total size of messages, in bytes, for the producer.

Trend Graphs

Shows message data for the selected producer.

Msgs / sec -- Traces the number of messages for the producer, per second.


Bytes / sec -- Traces the size of messages for the producer, in bytes.


Log Scale This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero



When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Consumers

Track utilization metrics for consumers on a single server.

← ↑ Dest.

EMS Consumers for Server06-Jul-2016 13:57Data OK

EMS Server: tcp://192.168.200.171:6010Name: EMS-SLDEMOS2-601012 Consumers

Client ID Filter: DestName Filter: Show Active Only

Count: 12Msgs / sec: 8.0Bytes / sec: 45.0 KBTotals Msgs: 6,224,639Total Bytes: 28.1 GB

ID	Client ID	Dest Name	Msgs / sec	Total Msgs	By
7		\$TMP\$.EMS-SLDEMOS2-6020.>	0.0	0	
9		\$TMP\$.EMS-SLDEMOS2-6030.>	0.0	0	
2583		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11EF3...	4.0	3,304,272	
2653		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11F2B...	0.0	2,197,813	
2667		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11F35...	0.0	276,789	
2785		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11F91...	0.0	350,505	
2808		\$TMP\$.EMS-SLDEMOS1-7010.>	0.0	0	
2809		\$TMP\$.EMS-SLDEMOS1-7020.>	0.0	0	
2810		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA5.6	4.0	45,130	
2812		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA7.4	0.0	23,655	
2813		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA8.2	0.0	13,594	
2814		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA9...	0.0	12,881	

Title Bar (possible features are):

← ↑ Open the previous and upper display.

+ Open an instance of this display in a new window.

? Open the online help page for this display.

Menu Table open commonly accessed displays.

6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

Note: Clicking on a row in the Consumers table and then clicking the **Dest.** button in the Title Bar takes you to the “Single Topic Summary” display for the selected consumer.

Fields and Data

This display includes:

- EMS Server

The EMS Server selected from this drop-down list displays a list of the currently connected Consumers.
- Name

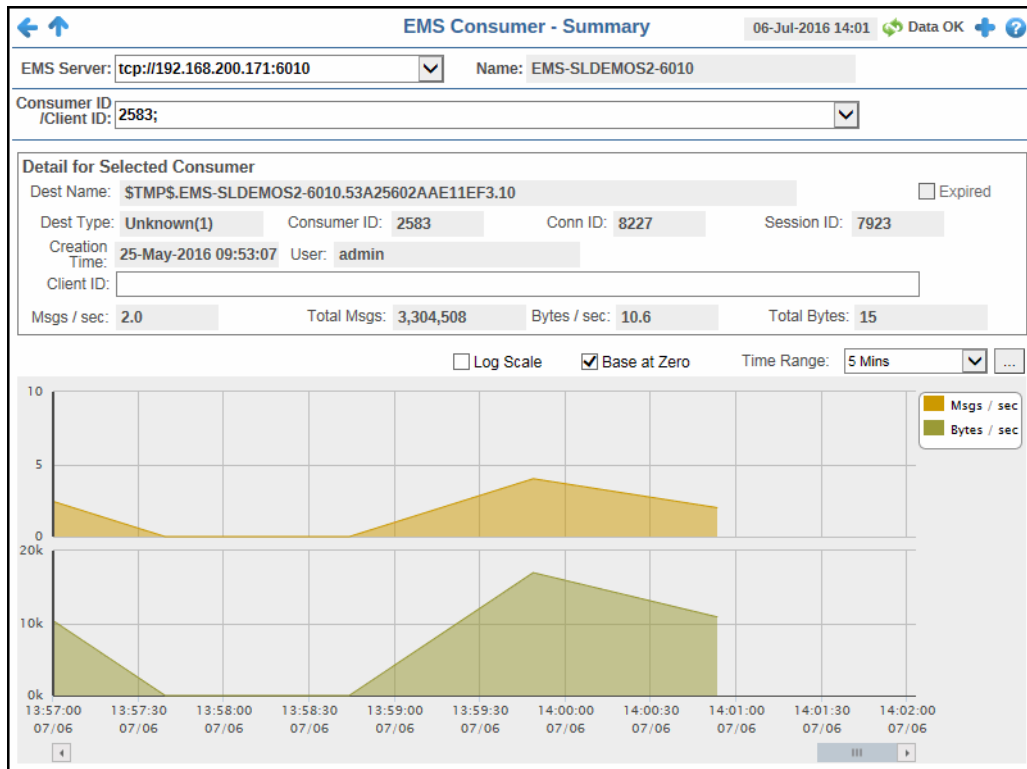
The name of the EMS Server selected from the EMS Server drop-down menu.

Consumers	The number of currently connected consumers on the server.
Client ID Filter	Filter field that allows you to filter the list of consumers by client ID. This filter works in conjunction with the DestName Filter to display the list of consumers.
DestName Filter	Filter field that allows you to filter the list of consumers by destination name. This filter works in conjunction with the Client ID Filter to display the list of consumers.
Show Active Only	Select this check box to display only active consumers.
Count	The number of currently connected consumers on the server.
Msgs / sec	The number of messages, per second, for the consumer.
Bytes / sec	The amount of messages, in bytes per second, for the consumer.
Total Msgs	The total number of messages for the consumer.
Total Bytes	The total size of messages, in bytes, for the consumer.
Table	This table shows metrics for each consumer on the selected server. Double-clicking on a row in the Consumers table displays details for the consumer in the "Consumer Summary" drill-down display.
ID	A unique string identifier assigned to each consumer.
Client ID	A unique string identifier assigned to each client.
Dest Name	The name of the destination.
Msgs / sec	The number of messages, per second, for the consumer.
Total Msgs	The total number of messages for the consumer.
Bytes / sec	The size of messages, in bytes per second, for the consumer.
Total Bytes	The total size of messages, in bytes, for the consumer.
User	The user name.
Host	The name of the host machine.
Session ID	A unique string identifier assigned to each session.
Conn ID	A unique string identifier assigned to each connection.
Curr Msg Sent Count	<p>The number of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p>The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Curr Msg Sent Size	<p>The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p>Note: The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Total Msg Ack Count	<p>The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session.</p> <p>Note: The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Total Msg Sent Count	<p>The total number of messages sent to the consumer since the consumer was created.</p> <p>Note: The sl.rtvew.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>

Elapsed Since Last Ack	<p>The amount of time (in milliseconds) that has elapsed since the last time a message sent to the consumer was acknowledged by the consumer's session.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Elapsed Since Last Sent	<p>The amount of time (in milliseconds) that has elapsed since the last time the server sent a message to the consumer.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Destination Prefetch	<p>The actual destination prefetch value used by the server at runtime.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Prefetch Deliv Count	<p>The number of prefetch messages delivered to the consumer by the server. For consumers receiving messages on any destination with positive prefetch value, this value is never more than the prefetch value of the destination. This value cannot be used to identify the status of the consumer, but it can be used in conjunction with other consumer information values to identify consumers who stopped receiving messages due to application-specific problems.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Durable Name	The name of the durable.
Route Name	<p>The queue owner server name if the consumer if the consumer's destination is a routed queue.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Is Active	<p>When checked, the consumer is active and can receive messages from the server.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Is System	<p>This check box is checked if the consumer was automatically created by the system.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Session Ack Mode	<p>Lists the consumer's session acknowledge mode as a constant defined in TibjmsAdmin.</p> <p>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</p>
Create Time	The amount of time, in milliseconds, since the consumer was created.
time_stamp	The date and time this row of data was last updated.
Expired	<p>When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p>
Dest Type	The configured destination type.

Consumer Summary

Displays details for an individual consumer. You can access this display by double-clicking on a producer in the “Consumers” 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.
- open commonly accessed displays.
- The number of items currently in the display.

- 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.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:

- EMS Server** The selected EMS Server populates the Consumer ID/ Client ID drop-down menu with Consumer IDs/Client IDs belonging to this EMS Server. This drop down list defaults to the EMS Server that was selected on the previous display.
- Name** The name of the EMS Server selected from the **EMS Server** drop-down menu.
- Consumer ID/ Client ID** Drop-down menu containing the Consumer IDs/Client IDs. This drop down list defaults to the Consumer ID/Client ID that was selected on the previous display.

**Detail for
Selected
Consumer**

Shows metrics for the consumer selected from the table.

Dest Name	The name of the destination.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the Expire Time field in the Duration region in the RTView Configuration Application > (Project Name) > Solution Package Configuration > TIBCO Enterprise Message Service > DATA STORAGE tab. The Delete Time field (also in the Duration region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Dest Type	The configured destination type.
Consumer ID	A unique string identifier assigned to each consumer.
Conn ID	A unique string identifier assigned to each connection.
Session ID	A unique string identifier assigned to each session.
Creation Time	The amount of time, in milliseconds, since the consumer was created.
User	The user name.
Client ID	A unique string identifier assigned to each client.
Msgs/sec	The number of messages, per second, for the consumer.
Total Msgs	The total number of messages for the consumer.
Bytes/sec	The size of messages, in bytes per second, for the consumer.
Total Bytes	The total size of messages, in bytes, for the consumer.

Trend Graphs

Shows message data for the selected producer.

Msgs / sec -- Traces the number of messages for the consumer, per second.


Bytes / sec -- Traces the size of messages for the consumer, in bytes.


Log Scale This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero



When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



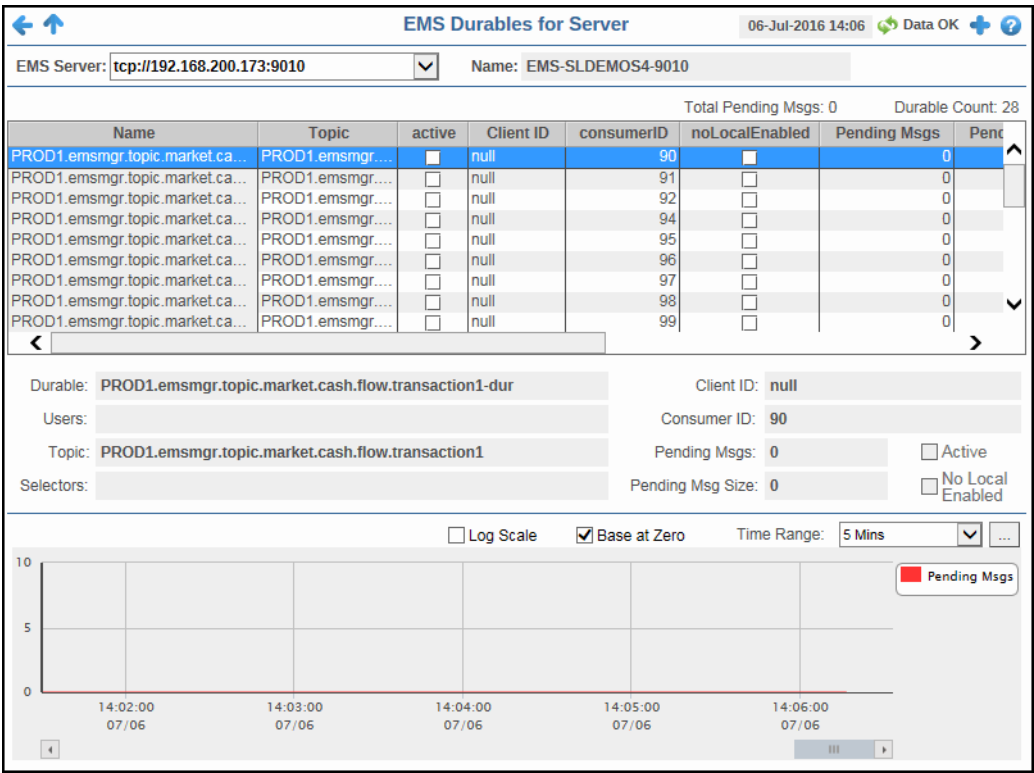
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Durables

Track utilization metrics for durables on a single server.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data

This display includes:


- EMS Server** The EMS Server selected from this drop-down menu populates all associated Durables data in this display.
- Name** The name of the EMS Server selected from the **EMS Server** drop-down menu.
- Total Pending Msgs** The total number of pending messages for the durable.
- Durable Count** The number of currently connected durables on the server.
- Table** This table shows metrics for each durable on the selected server.


Name	The name of the durable.
Topic	The name of the topic.
Active	Indicates whether the durable is active.
Client ID	A unique string identifier assigned to each client.
consumerID	A unique string identifier assigned to each consumer.
NoLocalEnabled	Indicates whether the subscriber receives messages from all connections its local connection. Enabled -- The subscriber does not receive messages sent from its local connection. Disabled -- The subscriber receives messages from all connections.
Pending Msgs	The total number of pending messages for the selected durable.
Pending Size	The total amount of pending messages, in bytes, for the selected durable.
Selector	Indicates that the subscriber only receives messages that match this selector.
userName	The name of the user of this durable subscriber.
time_stamp	The date and time this row of data was last updated.
Durable	The name of the durable selected from the table.
Users	The names of the users of this durable subscriber.
Topic	The name of the topic.
Selectors	Indicates that the subscriber only receives messages that match this selector.
Client ID	A unique string identifier assigned to each client.
Consumer ID	A unique string identifier assigned to each consumer.
Pending Msgs	The total number of pending messages for the selected durable.
Pending Msg Size	The total size of pending messages, in bytes, for the selected durable.
Active	Indicates whether the durable is active.
No Local	Indicates whether the subscriber receives messages from all connections its local connection. Enabled The subscriber does not receive messages sent from its local connection. Disabled The subscriber receives messages from all connections.
Trend Graphs	Shows message data for the selected consumer. Pending Msgs -- Traces the number of pending messages for the durable. Log Scale This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.


Base at Zero



When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

TIBCO FTL

The Solution Package for TIBCO FTL Views (and their associated displays) can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO FTL is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central. The available Views are:

- "FTL Servers"
- "FTL Clients"
- "FTL Events"


FTL Servers


These displays present performance metrics and alert status for all FTL servers. Displays are:


- **"All Servers Heatmap"**: Heatmap shows server and alert status for all FTL servers in all realms.
- **"All Servers Table"**: Table shows all available utilization metrics for all FTL servers.
- **"All Group Servers Table"**: Table shows the status and ID of all FTL Group Servers.
- **"All Satellites Table"**: Table shows the status and ID of all satellites.
- **"Server Summary"**: Current and historical metrics for a single FTL server.

All Servers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your servers for each available metric. You can view the servers in the heatmap based on the following metrics: the current alert severity, the current alert count, the total number of clients, the current amount of CPU being used, the current amount of memory being used for processing, the current amount of virtual memory being used for processing, and the number of inbox send faults. By default, this display shows the heatmap based on the **Alert Severity** metric.



Each heatmap rectangle represents a server. The rectangle color indicates the most critical alert state. You can click on a rectangle to drill-down to the **"Server Summary"** display and view metrics for that server. Clicking on the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays. You can also mouse-over the rectangles to view more details about host performance and status.

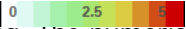
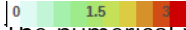
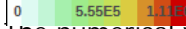
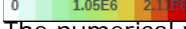
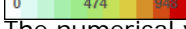





Fields and Data

This display includes:

- Names** Select to display the names of servers on the hosts.
- Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
- Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
- Metric** Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated "Server Summary" display for a detailed view of metrics for that particular server.
- Alert Severity** The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.
- 2** Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.
- 1** Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.
- 0** Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.
- Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle.
- The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

# Clients	The total number of clients in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlServerClientCountHigh . The middle value in the gradient bar indicates the middle value of the range.
CPU Usage	The total amount of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlServerCpuUsageHigh . The middle value in the gradient bar indicates the middle value of the range.
Memory	The current memory being used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlServerMemoryHigh . The middle value in the gradient bar indicates the middle value of the range.
V(irtual) Memory	The current virtual memory being used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlServerVirtualMemoryHigh . The middle value in the gradient bar indicates the middle value of the range.
#Inbox Faults	The total number of inbox faults. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlServerInboxSendFaultsHigh . The middle value in the gradient bar indicates the middle value of the range.

All Servers Table

Investigate detailed utilization metrics for all FTL servers. The **All Servers Table** contains all metrics available for servers, including the number of current client connections. Each row in the table describes a different server, and clicking on a table row drills-down to the “[Server Summary](#)” display allowing you to view metrics for that particular server. Clicking the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays. You can click a column header to sort column data in numerical or alphabetical order.

All FTL Servers - Table

26-Oct-2017 16:40 Data OK

Count: 3

All Servers Table

Connection	Current Realm Server	Alert Level	Alert Count	On Backup	Backup Server Status	% CPU	Client C
Realm1	192.168.200.227:8050		0	<input type="checkbox"/>	OK	0.70	
Realm1_Sat1	192.168.200.228:8050		0	<input type="checkbox"/>	OK	0.40	
Realm1_Sat2	192.168.200.229:8050		0	<input type="checkbox"/>	OK	2.00	

<

>

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

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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.




Fields and Data

This display includes:

Count The total number of active, inactive, and standby FTL servers. Column values are for the server except where noted. **Inactive Servers** are shown in dark red. **Standby Servers** are shown in blue. **Inactive Servers** are shown in dark gray. **Active Backup Servers** are shown in yellow.

All Servers Table

Connection The name of the connection.

Current Realm Server	The IP address and port number for the currently connected realm server.
Alert Level	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar, where 2 is the greatest Alert Severity.</p> <ul style="list-style-type: none">  One or more alerts have exceeded their specified ALARMLEVEL threshold, have an Alert Severity value of 2, and are shown in red.  One or more alerts have exceeded their specified WARNINGLEVEL threshold, have an Alert Severity value of 1, and are shown in yellow.  No alerts have exceeded an alert threshold, which have an Alert Severity value of 0, and are shown in green.
Alert Count	The number of alerts currently on the server.
On Backup	When checked, the server is in backup server mode.*
Backup Server Status	The current backup server status.*
%CPU	The percent CPU used on the server.*
Client Count	<p>The number of clients currently connected to the server.*</p> <p>Note that the client count might not match the number of clients found in the "All Clients Table", possibly due to the following:</p> <ul style="list-style-type: none"> One client might have one or more group joins resulting in a higher client count. For example, if a client has two group joins, CLIENT_COUNT equals 3, but will only be listed as a single client in the "All Clients Table". A TIBCO bridge could have one or more logical bridges running inside the bridge process, which could result in an increased CLIENT_COUNT even though there is actually only one client. Other FTL services could get a Client ID and, hence, be included in the CLIENT_COUNT even though they are not necessarily clients.
Clients Running	The number of connected clients on the server that currently have a status of RUNNING (which can be less than or equal to the client count).*
Cumulative Client Connects	The total number of clients the server has connected since the server was last started.*
Process Peak RSS (KB)	The maximum RSS memory used, in kilobytes.*
Process RSS (KB)	The current RSS memory used, in kilobytes.*
Process VM (KB)	The current virtual memory used, in kilobytes.*
#Bridge Servers	The number of bridge servers connected.*
#EFTL Clusters	The number of EFTL clusters connected.*
#Group Clients	The number of group clients connected.*
#Group Servers	The number of group servers connected.*

#Persistence Servers	The number of persistence servers connected.*
#Satellites	The number of satellites connected.*
Inbox Send Faults	The total number of faults when sending messages to inbox subscribers.*
Delta Inbox Send Faults	The number of faults when sending messages to inbox subscribers since the last data update.
Rate Inbox Send Faults	The rate of faults when sending messages to inbox subscribers
User CPU Time	Total amount of time the CPU spent, in microseconds, processing object code for users.*
System CPU Time	Total amount of time the CPU spent, in microseconds, processing operating system calls.*
Clients Destroyed	The total number of destroyed clients since the server was last started.*
Client Exceptions	The total number of client exceptions since the server was last started.*
Clients Needing Restart	The total number of clients that had to reconnect since the server was last started.*
Clients Off-line	The number of clients currently offline.*
Clients Out of Sync	The number of clients currently out of sync.*
Client Reconnects	The total number of clients that had to reconnect since the server was last started.*
Client Time-outs	The total number of clients connections that time out due to inactivity since the server was last started.*
Primary Realm Server	The configured primary realm server.*
Backup Realm Server	The configured secondary realm server.*
Server ID	The unique server ID.*
Uptime	The number of days, hours and minutes since the server was last started.*
Version	The FTL version on the server.*
Source	The source of the incoming data.

Expired

When checked (Expired=true), monitoring data for the FTL Server row has not been received within the time specified for expiration, which is defined (in seconds) using the **\$ftlServerRowExpirationTime** substitution located in the **conf\rtvpm_tftlmon.properties** file. If the row has been expired for an extended period of time, the **\$ftlRowExpirationTimeForDelete** substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the FTL Server rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.

To edit the default/current values, copy the lines below from **rtvpm_tftlmon.properties** file, paste them into the **sample.properties** file, and modify the lines in the **sample.properties** file:

```
#####
# CACHE / HISTORIAN SETTINGS

sl.rtvview.sub=$ftlServerRowExpirationTime:10
sl.rtvview.sub=$ftlRowExpirationTimeForDelete:0
```

Data Timestamp

The date and time this row of data occurred in FTL.*

Timestamp

The date and time this row of data was last updated in RTView.

All Group Servers Table

This display allows you to view the status of all FTL group servers, see their parent realm, and see which are primary and secondary backup servers. Each row in the table is a different parent realm, and clicking on a table row drills-down to the [“Clients by Group”](#) display allowing you to view detailed metrics for that group. You can click a column header to sort column data in ascending or descending order.

Parent Realm	server_mode	server_state	Expired	Timestamp
Realm1	Primary	RUNNING	<input type="checkbox"/>	24-Oct-2017 17:06:02

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.
- open commonly accessed displays.

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.

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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

Filter By:

Server: Select the server for which you want to view data, or select **All Group Servers** to view data for all servers.

Count The total number of group servers. Inactive Servers are shown with a dark gray background and Backup Servers are shown with a yellow background.

Group Server Status Table

server_mode	<p>The server mode. The server could be running as a primary server, acting as a backup server or running as a satellite. Valid values are: *</p> <p>Primary -- the server is running as a primary server.</p> <p>Secondary -- the server is acting as a secondary server.</p>
server_state	<p>The current server state. Valid values are: *</p> <p>Running -- the server is up and running.</p> <p>Stopped -- the server is stopped.</p>
Expired	<p>When checked (Expired=true), monitoring data for the FTL Server row has not been received within the time specified for expiration, which is defined (in seconds) using the \$ftlServerRowExpirationTime substitution located in the conf\rtvapm_tftlmon.properties file. If the row has been expired for an extended period of time, the \$ftlRowExpirationTimeForDelete substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the FTL Server rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.</p> <p>To edit the default/current values, copy the lines below from rtvapm_tftlmon.properties file, paste them into the sample.properties file, and modify the lines in the sample.properties file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS sl.rtvview.sub=\$ftlServerRowExpirationTime:10 sl.rtvview.sub=\$ftlRowExpirationTimeForDelete:0</pre>
Timestamp	The date and time this row of data was last updated in RTView.

All Satellites Table

View the status, ID, and parent realm of all satellites. Each row in the table is a different satellite, and clicking on a table row drills-down to the [“Server Summary”](#) display allowing you to view performance metrics for the server hosting the satellite. You can click a column header to sort column data in ascending or descending order.

← ↑ 🏠 All FTL Satellites - Table 24-Oct-2017 17:07 🟢 Data OK + ?

Server: Realm1

Count: 2

All Satellites Table				
Parent Realm	Satellite	Satellite Label	Satellite Status	Satellite UUID
Realm1	Realm1_Sat1	192.168.200.229:8050	Running	d28a50aa-a0a8-44f0-b3ea-8da9359
Realm1	Realm1_Sat2	192.168.200.229:8050	Running	d28a50aa-a0a8-44f0-b3ea-8da9359

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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

Filter By:

Server Select the server for which you want to view data.

Count The total number of table of satellite servers associated with the selected server. If the satellite server is not running, the row background is shown in yellow. When the satellite server is inactive (**Expired=true**) the row is shown in dark gray.


All Satellites Table

Parent Realm The name of the host server for the satellite.*

Satellite	The satellite IP address and port number.*
Satellite Status	<p>The server state. Valid values are: *</p> <p>Running -- the satellite is up and running.</p> <p>Stopped -- the satellite is stopped.</p>
Satellite UUID	The unique ID for the satellite.*
Expired	<p>When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the \$tftlServerRowExpirationTime substitution located in the conf\rtvapm_tftlmon.properties file. If the row has been expired for an extended period of time, the \$tftlRowExpirationTimeForDelete substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.</p> <p>To edit the default/current values, copy the lines below from rtvapm_tftlmon.properties file, paste them into the sample.properties file, and modify the lines in the sample.properties file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS sl.rtvview.sub=\$tftlServerRowExpirationTime:10 sl.rtvview.sub=\$tftlRowExpirationTimeForDelete:0</pre>
Timestamp	The date and time this row of data was last updated in RTView.

Server Summary

This display allows you to investigate performance issues for the selected server. You can track current and historical performance metrics for a single FTL server and view how many components (satellites, EFTL clusters; bridge, group and persistence servers) the server hosts.



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.




Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

Filter By:

Server Select the FTL Server for which you want to view data.

Server Metrics

Alert State The maximum alert level on the server:

-  One or more alerts have exceeded their specified **ALARMLEVEL** threshold.
-  One or more alerts have exceeded their specified **WARNINGLEVEL** threshold.
-  No alerts have exceeded an alert threshold.

Expired When checked (Expired=true), monitoring data for the FTL Server row has not been received within the time specified for expiration, which is defined (in seconds) using the **\$ftlServerRowExpirationTime** substitution located in the **conf\rtvapm_tftlmon.properties** file. If the row has been expired for an extended period of time, the **\$ftlRowExpirationTimeForDelete** substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the FTL Server rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.

To edit the default/current values, copy the lines below from **rtvapm_tftlmon.properties** file, paste them into the **sample.properties** file, and modify the lines in the **sample.properties** file:

```
#####
# CACHE / HISTORIAN SETTINGS

sl.rtvapm.sub=$ftlServerRowExpirationTime:10
sl.rtvapm.sub=$ftlRowExpirationTimeForDelete:0
```

on Backup When checked, this server has a backup server.*

Realm Server The server IP address or host name.*

Label The server label.*

FTL Version The FTL version on the server.*

Backup Server The IP address and port of the backup server.*

Backup Status The current backup server status.*

Up Time The number of days, hours and minutes since the server was last started.*

Inbox Faults The total number of faults when sending messages to inbox subscribers.*

Satellites The number of satellites.*

Clients

Current Clients The number of clients currently on the server.*

Running Clients The number of clients currently active on the server.*

Cumulative Client Connections	The total number of clients the server has connected since it the server was last started.*
Destroyed	The total number of destroyed clients since the server was last started.*
Exceptions	The total number of client exceptions since the server was last started.*
Needs Restart	The total number of clients that had to reconnect since the server was last started.*
Offline	The number of clients currently offline.*
Out of Sync	The number of clients currently out of sync.*
Timed out	The total number of clients connections that timed out due to inactivity since the server was last started.*
Reconnects	The total number of clients that had to reconnect since the server was last started.*

Process

% CPU	The amount of CPU used, in percent.*
Peak RSS	The maximum RSS memory used, in kilobytes.*
RSS	The current RSS memory being used, in megabytes.*
Virtual Memory	The current virtual memory being used, in megabytes.*

Components

Bridge Servers	The number of bridge servers connected.*
EFTL Clusters	The number of EFTL clusters connected.*
Group Clients	The number of group clients connected.*
Group Servers	The number of group servers connected.*
Persistence Servers	The number of persistence servers connected.*
Satellites	The number of satellites connected.*

Performance Metrics Trends

Traces the following for the selected server:


% CPU -- The percent CPU used.


Peak RSS -- The maximum amount of RSS memory used.


RSS -- The RSS memory used.

Process VM -- The current virtual memory used.

Inbox Send Faults/sec -- The rate of faults when sending messages to inbox subscribers (per second).

- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.


FTL Clients

These displays present performance metrics and alert status for FTL clients. Available displays are:

- **"All Clients Heatmap"**: Shows alert status for all FTL clients in all realms in a heatmap format.
- **"All Clients Table"**: Shows all available utilization metrics for all FTL clients on a selected server in a tabular format.
- **"Clients by Group"**: Lists all FTL clients by group with member details and CPU utilization in a tabular format.
- **"Client Summary"**: Displays current and historical metrics for a single FTL client.
- **"Client Metrics"**: Provides details about metric calculations for each FTL client.







All Clients Heatmap


This heatmap allows you to view the status and alerts of all FTL clients. You can view the clients in the heatmap based on the following metrics: the current alert severity, the current alert count, the current CPU usage, the rate of messages received, and the rate of messages sent.


Each heatmap rectangle represents a client, and clicking on a client drills-down to the “[Client Summary](#)” display allowing you to view metrics for that client. The rectangle color indicates the most critical alert state. Clicking on the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays. You can also mouse-over the rectangles to view more details about host performance and status.



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.
-   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.



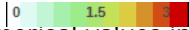

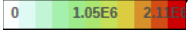
 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Filter By:


Server Choose the server for which you want to view data.

Fields and Data


This display includes:

Names	Select this check box to display the names of clients.
Log	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
Auto	When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
Metric	Select the metric driving the heatmap display. The default is Alert Severity. Each Metric has a color gradient bar that maps values to colors. Each rectangle represents a client. Mouse-over any rectangle to display the current values of the metrics for the client. Click on a rectangle to drill-down to the associated "Client Summary" display for a detailed view of metrics for that particular client.
Alert Severity	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar , where 2 is the greatest Alert Severity.</p> <p>2 Metrics that have exceeded their specified ALARMLEVEL threshold and have an Alert Severity value of 2 are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p>1 Metrics that have exceeded their specified WARNINGLEVEL threshold and have an Alert Severity value of 1 are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p>0 Metrics that have not exceeded either specified threshold have an Alert Severity value of 0 and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
Alert Count	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
CPU Usage	<p>The total amount of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlClientCpuUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Memory	<p>The current amount of memory used for processing. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlClientMemoryHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
V(irtual) Memory	<p>The current amount of virtual memory being used for processing. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of TftlClientVirtualMemoryHigh. The middle value in the gradient bar indicates the middle value of the range.</p>


**Msgs
Rcvd/sec**

The rate of messages received (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlClientMsgsRcvdRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Msgs
Sent/sec**

The rate of messages sent (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlClientMsgsSentRateHigh**. The middle value in the gradient bar indicates the middle value of the range.


All Clients Table

Investigate detailed utilization metrics for all FTL clients. The **All Clients Table** contains all metrics available for clients, including the number of current client connections. Each row in the table contains data for a particular client, and clicking on a table row drills-down to the “[Client Summary](#)” display allowing you to view metrics for that particular client. You can click a column header to sort column data in ascending or descending order. Clicking the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays.

Note that the number of clients found in this table might not match the client count found in the “[All Servers Table](#)”, possibly due to the following:

- One client might have one or more group joins resulting in a higher client count. For example, if a client has two group joins, CLIENT_COUNT equals 3, but will only be listed as a single client in this table.
- A TIBCO bridge could have one or more logical bridges running inside the bridge process, which could result in an increased CLIENT_COUNT even though there is actually only one client.
- Other FTL services could get a Client ID and, hence, be included in the CLIENT_COUNT even though they are not necessarily clients.

All FTL Clients - Table							
Server: Realm1		Count: 4					
Show: <input type="checkbox"/> Expired		All Clients Table					
Realm Server	Client	Alert Level	Alert Count	Client Status	% CPU	Process VM (KB)	Proc
Realm1	consumer		0	RUNNING	0.01	1,234,436	
Realm1	_GroupServer.PrimaryGroupServer		0	RUNNING	0.99	1,921,396	
Realm1	_GroupServer.BackupGroupServer		0	RUNNING	0.42	1,211,916	
Realm1	.tiblatsend		0	RUNNING	0.56	1,209,216	



Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

Filter By:

- Server** Select the server for which you want to view data.
- Show: Expired** Select this check box to display those rows containing expired data. Leave unchecked to display only those rows that are not expired.
- Count** The total number of active, inactive, and standby FTL clients. **Inactive Clients** are shown in dark red. **Standby Clients** are shown in blue.

All Clients Table

- Realm Server** The name of the server.
- Client** The name of the client.
- Alert Level** The maximum alert level. Values range from 0 to 2, as indicated in the color gradient bar, where **2** is the greatest Alert Severity.
- One or more alerts have exceeded their specified **ALARMLEVEL** threshold, have an Alert Severity value of **2**, and are shown in red.
 - One or more alerts have exceeded their specified **WARNINGLEVEL** threshold, have an Alert Severity value of **1**, and are shown in yellow.
 - No alerts have exceeded an alert threshold, which have an Alert Severity value of **0**, and are shown in green.
- Alert Count** The current number of alerts.
- Client Status** The status of the client. For example, RUNNING.*
- % CPU** The amount of CPU used, in percent.*
- Process VM (KB)** The amount of virtual memory used for processing, in kilobytes. *
- Process RSS (KB)** The current RSS memory being used, in kilobytes.*
- Process Peak RSS (KB)** The maximum RSS memory used, in kilobytes.*
- Msgs Rcvd/sec** The number of messages received per second.

Msgs Sent/sec	The number of messages sent per second.
Delta Msgs Rcvd	The total number of messages received since the last data update.
Delta Msgs Sent	The total number of messages sent since the last data update.
Total Msgs Rcvd	The total number of messages received since the client started.*
Total Msgs Sent	The total number of messages sent since the client started.*
Store Mismatch Msgs	Any non-zero value indicates a store mismatch misconfiguration, which occurs when a direct path transport connects two endpoints that are associated with two different persistence stores.*
Dynamic Formats	The number of distinct dynamic formats that the client creates within the sample interval.*
User CPU	The amount of time the user has used the CPU, in microseconds.*
Application	The application name of the client.*
Application Instance	The application instance of the client.*
Client ID	The unique ID for the client.*
Process ID	The unique ID for the process.*
FTL User	The FTL user that is being used by the client.*
Effective User	The UID of the client (which is used for most access checks).*
Host	The host name.*
Host IP	The host IP address.*
FTL Version	The FTL version on the host.*
Expired	<p>When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the \$tftlRowExpirationTime substitution located in the conf\rtvapm_tftlmon.properties file. If the row has been expired for an extended period of time, the \$tftlRowExpirationTimeForDelete substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 120 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 120 seconds of inactivity and that expired rows will be removed from the cache after 3600 seconds (one hour) of inactivity.</p> <p>To edit the default/current values, copy the lines below from rtvapm_tftlmon.properties file, paste them into the sample.properties file, and modify the lines in the sample.properties file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS sl.rtview.sub=\$tftlServerRowExpirationTime:120 sl.rtview.sub=\$tftlRowExpirationTimeForDelete:0</pre>

Data Timestamp The date and time this row of data occurred in FTL.*

Local Timestamp The date and time this row of data was last updated in RTView.

Clients by Group

This display lists all clients and their associated FTL groups for a particular server. Each row in the table is a different client, and clicking on a table row drills-down to the “Client Summary” display allowing you to view metrics for that particular client. You can click a column header to sort column data in ascending or descending order.

← ↑

All FTL Groups - Table

27-Oct-2017 15:29 Data OK + ?

Server: Realm1

Group Server Mode: Primary

State: RUNNING

Count: 9

Clients by Group

Group Name	Client Label	Ordinal	Member Type	Client Status	% CPU	Client ID
testgroup1		1	FULL_MEMBER			
testgroup1		0	OBSERVER_MEMBER			
testgroup1		2	FULL_MEMBER			
testgroup3		0	OBSERVER_MEMBER			
testgroup3		1	FULL_MEMBER			
testgroup3		2	FULL_MEMBER			
testgroup4		1	FULL_MEMBER			
testgroup4		2	FULL_MEMBER			
testgroup4		0	OBSERVER_MEMBER			

<

>

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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

Filter By:

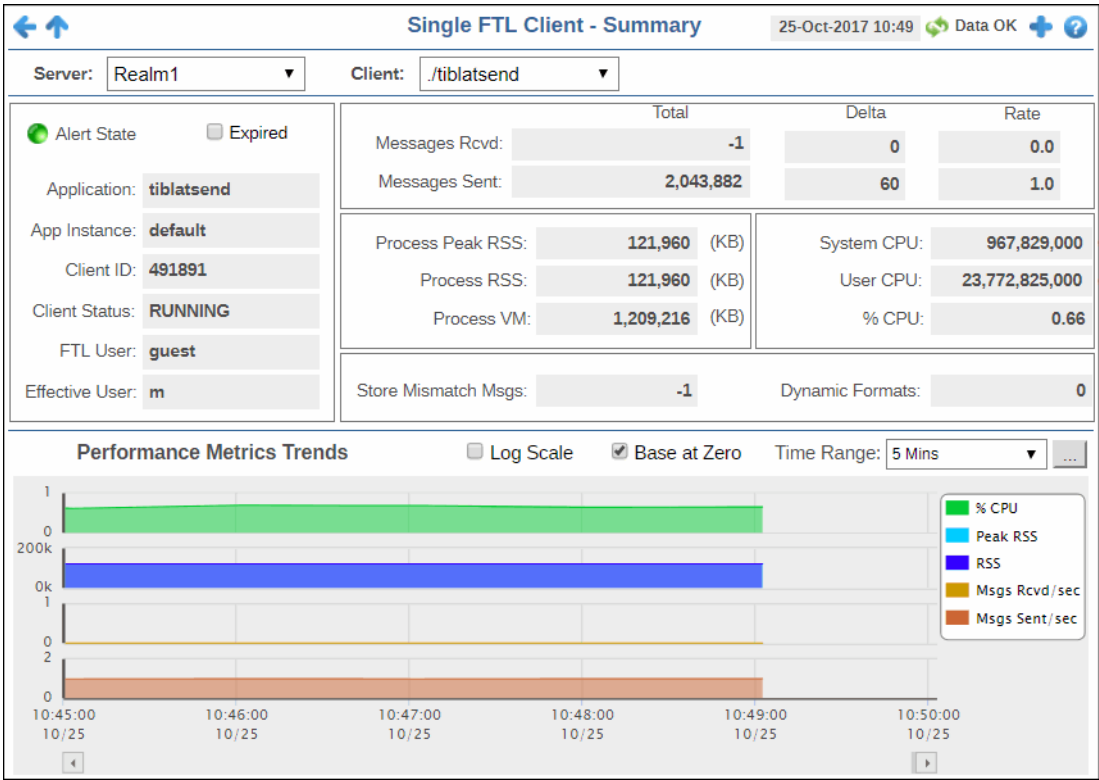
Server	Select the server for which you would like to view data.
Group Server Mode	The server mode. Valid values are: Primary -- The server is acting as a primary server. Secondary --The the server is acting as a secondary server.
State	Running -- The client is running and its local realm definition is up-to-date. Needs Restart -- The client needs to be restarted to update its local realm definition. Timed Out - The server has lost the heartbeat signal from the client. Either the client has stopped or a network issue is obstructing the signal. Exception - The client is running, but its realm definition is out-of-date. Out-of-Sync -- The realm definition revision on the client and server are different.
Count	The total number of FTL clients. Inactive clients are shown in dark gray and when the client is associated with an active backup server the client is are shown in yellow.

Clients by Group Table

Group Name	The name of the group.*
Client Label	The client's label.*
Ordinal	The number representing the client's position within the group. A value of -1 indicates that the client has been disconnected from the group server.*
Member Type	The client's member type.*
Client Status	The status of the client. For example, RUNNING.*
% CPU	The amount of CPU used, in percent.*
Client ID	The ID of the client.*
Process ID	The process ID.*
Expired	When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the \$fttlRowExpirationTime substitution located in the conf\rtvapm_tftlmon.properties file. If the row has been expired for an extended period of time, the \$fttlRowExpirationTimeForDelete substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 120 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 120 seconds of inactivity and that expired rows will be removed from the cache after 3600 seconds (one hour) of inactivity. To edit the default/current values, copy the lines below from rtvapm_tftlmon.properties file, paste them into the sample.properties file, and modify the lines in the sample.properties file: <pre>##### # CACHE / HISTORIAN SETTINGS sl.rtvapm.sub=\$fttlServerRowExpirationTime:120 sl.rtvapm.sub=\$fttlRowExpirationTimeForDelete:0</pre>
Timestamp	The date and time this row of data was last updated in RTView.

Client Summary

Track current and historical performance metrics for a single FTL client. Use this display to investigate performance issues of a client.



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.

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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

Filter By:

- Server** Choose an FTL server to display.
- Client** Choose a client to display.

Fields and Data

Alert State	<p>The maximum alert level has been exceeded. Values range from 0 to 2, as indicated in the color gradient bar, where 2 is the greatest Alert Severity.</p> <ul style="list-style-type: none">  One or more alerts have exceeded their specified ALARMLEVEL threshold, have an Alert Severity value of 2, and are shown in red.  One or more alerts have exceeded their specified WARNINGLEVEL threshold, have an Alert Severity value of 1, and are shown in yellow.  No alerts have exceeded an alert threshold, which have an Alert Severity value of 0, and are shown in green. 	
Expired	<p>When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the \$fttlRowExpirationTime substitution located in the conf\rtvapm_tftlmon.properties file. If the row has been expired for an extended period of time, the \$fttlRowExpirationTimeForDelete substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 120 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 120 seconds of inactivity and that expired rows will be removed from the cache after 3600 seconds (one hour) of inactivity.</p> <p>To edit the default/current values, copy the lines below from rtvapm_tftlmon.properties file, paste them into the sample.properties file, and modify the lines in the sample.properties file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS sl.rtvview.sub=\$fttlServerRowExpirationTime:120 sl.rtvview.sub=\$fttlRowExpirationTimeForDelete:0</pre>	
Application	The application name of the client.*	
App Instance	The application instance of the client.*	
Client ID	The unique identifier for the client.*	
Client Status	The status of the client. For example, RUNNING.*	
FTL User	The FTL user that is being used by the client.*	
Effective User	The UID of the client (which is used for most access checks).*	
Messages Rcvd	Total	The total number of messages received since the client started.*
	Delta	The number of messages received since the last data update.
	Rate	The number of messages received per second.
Messages Sent	Total	The total number of messages sent since the client started.*
	Delta	The number of messages sent since the last data update.
	Rate	The number of messages sent per second.
Process Peak RSS	The maximum RSS memory used, in kilobytes.*	
Process RSS	The current RSS memory being used, in kilobytes.*	
Process VM	The current virtual memory being used, in kilobytes.*	

System CPU	The amount of CPU used by the system, in kilobytes.*
User CPU	The amount of CPU used by the client, in kilobytes.*
% CPU	The percent of CPU used by the client.*
Store Mismatch Msgs	Any non-zero value indicates a store mismatch misconfiguration, which occurs when a direct path transport connects two endpoints that are associated with two different persistence stores.*
Dynamic Formats	The number of distinct dynamic formats that the client creates within the sample interval.*

Performance Metrics Trends

Traces the following for the selected client:

% CPU -- The amount of CPU used, in percent.

Peak RSS -- The maximum RSS memory used, in kilobytes.

RSS -- The current RSS memory being used, in kilobytes.


Msgs Rcvd/sec -- The number of messages received per second.


Msgs Sent/sec -- The number of messages sent per second.



Log Scale This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

Base at Zero When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



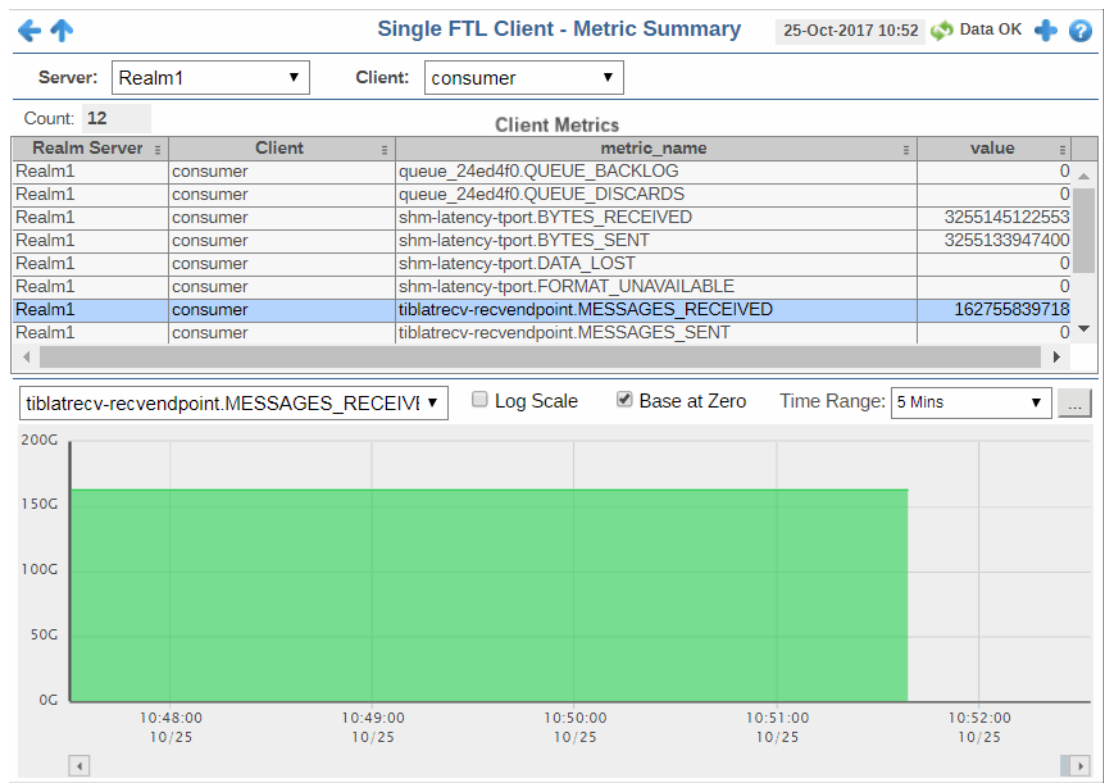
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Client Metrics

Track detailed performance and utilization metrics for a single FTL client.



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.


Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.



- Filter By:**
- Server** Select the server containing the client for which you want to see data.
 - Client** Select the client for which you want to see data.
 - Count** The number of rows/metrics in the table.

Client Metrics Table

Realm Server	The name of the server.
Client	The name of the client.
metric_name	The name of the metric.*
value	The current value of the metric.*
Trend Graph	Select a metric from the drop-down menu to trace in the trend graph for the selected client.
Log Scale	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
Base at Zero	When this option is checked, zero is set as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days , or display All Data . To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

FTL Events

This display presents FTL events for a particular client. The available display is:

- **"Events"**: Table that lists all FTL events for a particular client.

Events

This display allows you to view FTL events on one or all clients. You can view the event type, the client status, the associated application, and the FTL user name among other details.

← ↑ All FTL Events - Table 25-Oct-2017 10:54 Data OK + ?


Server: All Servers Client: All Clients

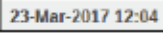
Count: 5 Time Range: 5 Mins ...

Timestamp	Server	Client	Event Type	Client Status	Client
03-Oct-2017 06:32:16.801	Realm1	client41	DISCONNECTED	DESTROYED	1361
03-Oct-2017 06:32:16.801	Realm1	client41	DISCONNECTED	DESTROYED	1355
03-Oct-2017 06:32:16.801	Realm1	client42	DISCONNECTED	DESTROYED	1351
03-Oct-2017 06:32:16.801	Realm1	client43	DISCONNECTED	DESTROYED	1353
03-Oct-2017 06:32:16.801	Realm1	client43	DISCONNECTED	DESTROYED	1357

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.

Note: Fields/columns with an asterisk (*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

Filter By:


Server Select an FTL server containing the client for which you want to view data.


Client Select an FTL client for which you want to view data.

Count The number of rows in the table.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Fields and Data

Timestamp	The date and time this row of data was last updated in RTView.
Server	The name of the server.
Client	The name of the client.
Event Type	The type of event.*
Client Status	The status of the client.*
Client ID	The ID of the client.*
Process ID	The process ID.*
Application	The application name of the client.*
Application Instance	The name of the application instance.*
FTL User	The FTL user that is being used by the client.*
Realm Server	The IP address and port of the realm server.*
Host	The name of the host.*
IP Address	The IP address of the client.*
Version	The version of the client running on the server.*

TIBCO Hawk

With the Solution Package for TIBCO Hawk™ you can centralize alert events triggered by Hawk alert rule bases and performance data relevant to monitored hosts. This enables RTView to be the event correlation engine and event management system for alerts that are generated by TIBCO Hawk, RTView Enterprise Monitor and other monitoring tools such as Oracle Enterprise Manager.

The TIBCO Hawk Views displays can be found under **Components** tab > **Connectors/TIBCO Hawk**. The displays within the Views will be populated with data once the Solution Package for TIBCO Hawk is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central. The following Views are available:

- "Hawk Agents Table"
- "Hawk Alerts Table"
- "RTView Agent Administration"

Hawk Agents Table

This table provides a list of agents as well as network connectivity details about each agent.

All Hawk Agents - Table View					
Agent Count: 11		Hawk Agents			
Agent	Status	Last Alert Level	Cluster	IP Address	Platform
agentW46	Alive	ALERT_LOW	192.168.200.0	192.168.200.146	amd64:Windows Server 2008 R2
SLHOST93	Alive	ALERT_HIGH	192.168.200.0	192.168.200.93	amd64:Linux:2.6.32-358.11.1.el6
WIN44	Alive	ALERT_HIGH	192.168.200.0	192.168.200.144	amd64:Windows Server 2008 R2
SLHOST21(dev)	Alive	ALERT_MEDIUM	192.168.200.0	192.168.200.121	amd64:Windows 7-6.1
SLHOST5(domain5)	Alive	ALERT_MEDIUM	192.168.200.0	192.168.200.105	x86:Windows XP:5.1
SLHOST6(domain6)	Alive	ALERT_MEDIUM	192.168.200.0	192.168.200.106	x86:Windows XP:5.1
SLHOST15(sl_amx)	Alive	ALERT_HIGH	192.168.200.0	192.168.200.115	amd64:Windows 7-6.1
SLHOST17(sl_amx)	Alive	ALERT_HIGH	192.168.200.0	192.168.200.117	amd64:Windows 7-6.1
SLHOST15(sl_qa_conn)	Alive	NO_ALERT	127.0.0.0	192.168.200.115	amd64:Windows 7-6.1
SLHOST16(sl_qa_conn)	Alive	ALERT_HIGH	192.168.200.0	192.168.200.116	amd64:Windows 7-6.1
SLHOST18(sl_qa_conn)	Alive	NO_ALERT	192.168.200.0	192.168.200.118	amd64:Windows 7-6.1

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.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

Fields and Data:

Agent Count: The total number of agents in the table.

Table:

Each row in the table is a different agent.

Agent	The name for the agent which is composed of the hostname and Hawk domain (in parenthesis). Agent names which do not contain an explicit Hawk domain are members of the "default" domain.
Status	The agent status, either Alive or Expired .
Last Alert Level	The most recent and most critical alert level.
Cluster	The IP address of the cluster to which this agent belongs.
IP Address	The IP subnet address for the group of machines to which this agent belongs.
Platform	The physical CPU class and operating system version.
Last Update	The date and time the row data was last updated.

Hawk Alerts Table




Use this display to view all Hawk alerts that have occurred in the system.

Each row in the table is a different active alert. Use the drop-down menus to filter the alerts listed. Click a column heading to sort the table on that column data.

The row color indicates the following:

Row Color Code:

Tables with colored rows indicate the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

Hawk Alerts Table

26-Oct-2016 10:27

Data OK

Agent Filter:

All Agents

Alert Text Filter

Clear

Rulebase Filter

Clear

☐ Show Cleared Alerts

Alert Count:10

Time	Agent	Alert ID	Alert Level	RuleBase	
26-Oct-2016 10:27:33	SLHOST6(domain6)	10	ALERT_MEDIUM	TibRV_Alerts	Received from
26-Oct-2016 10:27:19	SLHOST5(domain5)	10	ALERT_MEDIUM	TibRV_Alerts	Received from
26-Oct-2016 10:26:59	SLHOST5(domain5)	11	ALERT_LOW	System_Alerts	Server Process
26-Oct-2016 10:17:32	SLHOST17(sl_amx)	7	ALERT_HIGH	test	Current Proces
26-Oct-2016 10:06:05	SLHOST5(domain5)	13	ALERT_LOW	System_Alerts	System Uptime
26-Oct-2016 10:02:46	SLHOST16(sl_qa_conn)	4	ALERT_HIGH	generate_Alerts	Current Proces
26-Oct-2016 10:01:26	SLHOST6(domain6)	13	ALERT_LOW	System_Alerts	System Uptime
26-Oct-2016 00:26:52	SLHOST6(domain6)	11	ALERT_LOW	System_Alerts	Server Process
26-Oct-2016 00:20:33	SLHOST5(domain5)	12	ALERT_LOW	System_Alerts	Service Print S
26-Oct-2016 00:16:21	SLHOST6(domain6)	12	ALERT_LOW	System_Alerts	Service Print S

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.

- Agent Filter

Choose one or **All Agents**.
- Alert Text Filter

Enter a string to filter alerts listed.
Clear to remove this filter.
- Rulebase Filter

Enter a rule to filter alerts listed.
Clear to remove this filter.
- Show Cleared Alerts

When checked, cleared alerts are included in the table.
- Alert Count

The number of alerts in the table.
- Time

The date and time the alert occurred.

Agent	The name of the agent associated with the alert.
Alert ID	The unique string identifier for the alert.
Alert Level	<p>● ALERT_HIGH indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.</p> <p>● ALERT_MEDIUM indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.</p> <p>● ALERT_LOW indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.</p>
RuleBase	The alert system type (e.g.System_Alerts).
Alert Text	Descriptive text about the alert.
Cleared	When checked, the alert has been cleared.

RTView Agent Administration

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

RTView Agent Metrics Administration						
10-Nov-2014 16:31 Data OK						
Data Received from Remote Agents						
AgentName	AgentClass	Client ID	Total Rows Rcvd	Delta Rows rcvd	Rows Rcvd / sec	Last Receive Time
slapm	SL-RTVMGR-Agent	30002	43,412	0	0.0	10-Nov-2014 16:31:42
slapm	SL-HOSTMON-Agent	30017	53,750	35	8.6	10-Nov-2014 16:31:43
slapm	SL-BWMON-Agent	30018	423,741	8	4.0	10-Nov-2014 16:31:43
slcl4-64	SL-HOSTMON-Agent	30005	68,536	0	0.0	10-Nov-2014 16:31:37
slcl4-64	SL-BWMON-Agent	30006	91,694	0	0.0	10-Nov-2014 16:31:35
slcl4-64	SL-RTVMGR-Agent	30003	41,913	4	1.9	10-Nov-2014 16:31:43
slhost6	SL-HOSTMON-Agent	30026	23,418	0	0.0	10-Nov-2014 16:31:40
slhost6	SL-RTVMGR-Agent	30027	26,933	4	2.0	10-Nov-2014 16:31:42
slhost6	SL-BWMON-Agent	30032	26,321	14	2.3	10-Nov-2014 16:31:44
slhpux11	SL-BWMON-Agent	30012	34,363	0	0.0	10-Nov-2014 16:31:42
slhpux11	SL-HOSTMON-Agent	30010	64,394	0	0.0	10-Nov-2014 16:31:42
slhpux11	SL-RTVMGR-Agent	30011	41,820	64	15.4	10-Nov-2014 16:31:44
slvmrh2	SL-BWMON-Agent	30004	7,874	0	0.0	10-Nov-2014 16:31:38
slvmrh2	SL-RTVMGR-Agent	30001	45,352	0	0.0	10-Nov-2014 16:31:40
slvmrh2	SL-HOSTMON-Agent	30009	46,787	1	0.2	10-Nov-2014 16:31:44
slvmware	SL-BWMON-Agent	30013	6,085	0	0.0	10-Nov-2014 16:31:31
slvmware	SL-RTVMGR-Agent	30016	43,399	2	1.0	10-Nov-2014 16:31:43
slvmware	SL-HOSTMON-Agent	30015	33,434	0	0.0	10-Nov-2014 16:31:31

Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.



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.

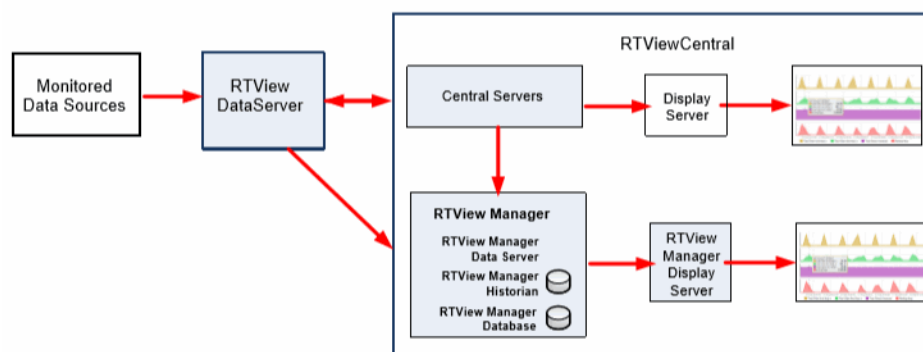
Data Received from Remote Agents Table

AgentName	Name of the agent.
AgentClass	Class of the agent.
Client ID	Unique client identifier.
Total Rows Rcvd	Total number of rows of data received.
Rows Rcvd/sec	Number of rows of data received per second.
Last Receive Time	Last time data was received from the agent.

CHAPTER 9 RTView Manager

Use RTView Manager to monitor RTView Enterprise Monitor. That is, to monitor the performance of RTView processes and applications on RTViewCentral as well as any RTView DataServers to which RTViewCentral is connected.

RTView Manager runs as a separate process on RTViewCentral and has its own data server, database, Display Server and Historian.



RTView Manager requires minimal setup as required connections are auto-discovered from RTViewCentral. To access this monitor, start RTViewCentral and go to **<http://localhost:3070/rtview-manager-classic>**. You can use any one of the following username/password combinations to log in: admin/admin, demo/demo, rtvuser/rtvuser, rtvadmin/rtvadmin or rtvalertmgr/rtvalertmgr.

RTView Manager displays are organized by the following Views:

- **"RTView Manager"**: This series of displays is for monitoring the health of the RTView servers monitoring your system. RTView Manager metrics include connected state, number of clients and other status information for Data Server, Historian and Display Server processes.
- **"JVM "**: For monitoring the health of Java Virtual Machine (JVM) processes. JVM metrics track garbage collection information and trends, including memory usage before and after garbage collection, duration and duty cycles. This, combined with tracking of JVM memory pool trends, enables you to be notified of memory leaks, unusual garbage collection activities and CPU and memory resource issues automatically with minimal user analysis, speeding the discovery of the root cause of any issue. It also monitors a Java Virtual Machine's memory heap, non-heap memory, monitor threads and other key metrics to ensure the JVM has good performance. Detailed metrics including JVM CPU usage, Max Heap, Current Heap, Used Heap and Live Threads can all be tracked over time.
- **"Tomcat "**: For monitoring the health of Tomcat servers, applications and all installed web modules. Performance data provided includes current and historic metrics, number of sessions, request rates, cache hit rates and data transmission metrics.

RTView Manager

The following RTView Servers displays can be found under **Components** tab > **Processes** > **RTView Processes** > **RTView Servers** after installation.

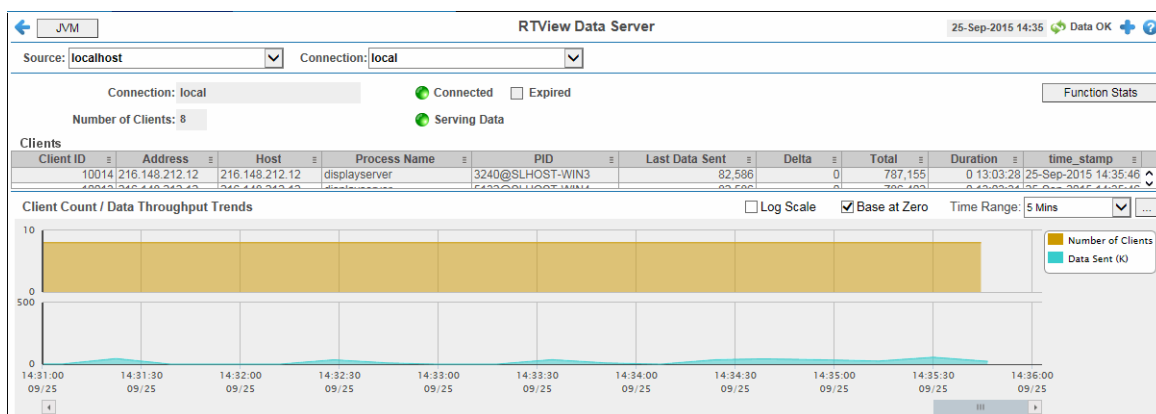
These displays present performance data for all RTView Enterprise Monitor servers. Use these displays to monitor the health of the servers monitoring your system. Displays in this View are:

- **"Data Servers"**: Shows metrics for RTView Data Servers.
- **"Display Servers"**: Shows metrics for RTView Display Servers.
- **"Historian Servers"**: Shows metrics for RTView Historian Servers.
- **"Version Info"**: Shows the version information of each jar used in each connected RTView application.

Data Servers

Track data transfer metrics for RTView Data Servers, client count and throughput trends.

Use the available drop-down menus or right-click to filter data shown in the display.







Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

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.

- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Connection** The connection selected from the **Connection** drop-down menu.
- Number of Clients** The number of clients currently server on this Data Server.

Connected	The Data Server connection state:  Disconnected.  Connected.
Serving Data	 The Data Server is not currently serving data.  The Data Server is currently serving data.
Expired	This server has been marked as expired after no activity.
Function Stats	Opens the RTView Function Stats display which shows detailed performance statistics for RTView functions in the selected Data Server. This button is only enabled if the RTView has a JMX connection defined for the selected Data Server.

Clients

This table describes all clients on the selected server.

Address	The client IP address.
Client ID	The unique client identifier.
Duration	The amount of time for this client session. Format: dd HH:MM:SS <days> <hours>:<minutes>:<seconds> For example: 10d 08:41:38
Host	The client host name.
Last Data Sent	The amount of data, in bytes, last sent to the client.
Delta	The amount of data, in bytes, sent since the last update.
Total	The total amount of data, in bytes, sent to the client.
TIME_STAMP	The date and time this row of data was last updated.

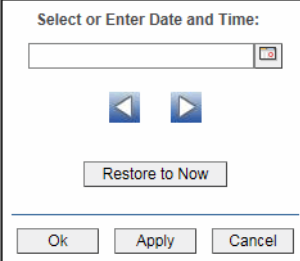
Client Count / Data Throughput Trends

Shows throughput metrics for all clients on the selected server.


Log Scale	Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Use zero as the Y axis minimum for all graph traces.



Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right). At the bottom are three buttons: "Restore to Now", "Ok", "Apply", and "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   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Number of Clients

Traces the number of clients being served by the Data Server.

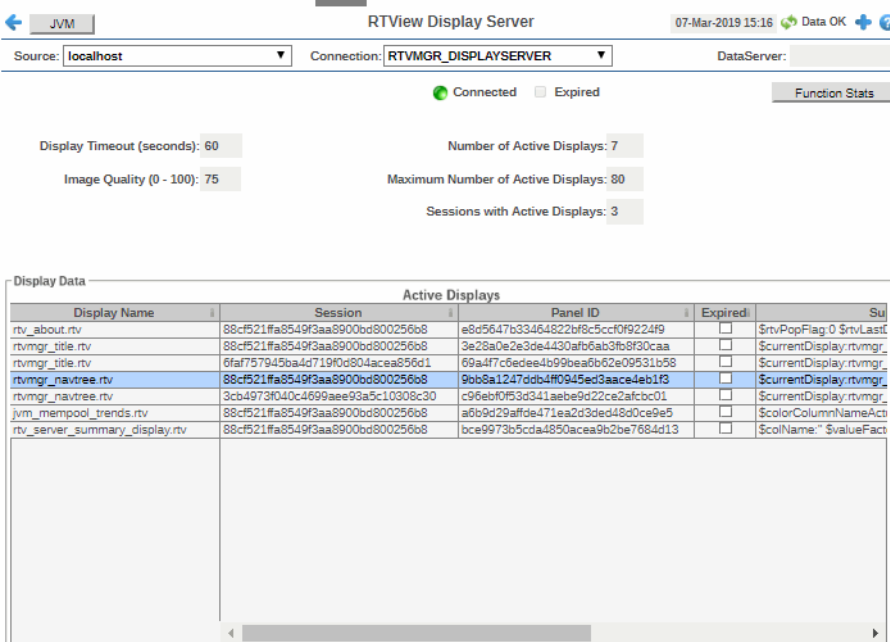
Data Sent

Traces the total amount of data, in Kilobytes, sent to all clients.

Display Servers

Track display utilization metrics for RTView Display Servers.


Use the available drop-down menus or right-click to filter data shown in the display.



The interface shows the RTView Display Server configuration and active displays. The top bar includes a back arrow, a "JVM" tab, the title "RTView Display Server", and a timestamp "07-Mar-2019 15:16" with status icons. Below the title bar, there are dropdown menus for "Source" (set to "localhost"), "Connection" (set to "RTVMGR_DISPLAYSERVER"), and "DataServer". A "Function Stats" button is on the right. Below these are status indicators: "Connected" (green dot) and "Expired" (grey dot). Configuration settings include "Display Timeout (seconds): 60", "Image Quality (0 - 100): 75", "Number of Active Displays: 7", "Maximum Number of Active Displays: 80", and "Sessions with Active Displays: 3".

The "Display Data" section shows a table of active displays:

Display Name	Session	Panel ID	Expired	Source
rtv_about.rtv	88cf521fa8549f3aa8900bd800256b8	e8d5647b33464822bf8c5ccf0f9224f9	<input type="checkbox"/>	\$rtvPopFlag 0 \$rtvLastF
rtvmgr_title.rtv	88cf521fa8549f3aa8900bd800256b8	3e28a0e2e3de4430afb6ab3fb8f90caa	<input type="checkbox"/>	\$currentDisplay.rtvmgr_
rtvmgr_title.rtv	6faf757945ba4d719f0d804acea856d1	69a4f7c6edee4b99beaf6b82e09531b58	<input type="checkbox"/>	\$currentDisplay.rtvmgr_
rtvmgr_navtree.rtv	88cf521fa8549f3aa8900bd800256b8	9bb8a1247ddb4f0945ed3aace4eb1f3	<input type="checkbox"/>	\$currentDisplay.rtvmgr_
rtvmgr_navtree.rtv	3cb4973f040c4699aee93a5c10308c30	c96ebf0f53d341aeb9d22ce2afcb01	<input type="checkbox"/>	\$currentDisplay.rtvmgr_
jvm_mempool_trends.rtv	88cf521fa8549f3aa8900bd800256b8	a6b9d29afde471ea2d3ded48d0ce9e5	<input type="checkbox"/>	\$colorColumnNameActi
rtv_server_summary_display.rtv	88cf521fa8549f3aa8900bd800256b8	bce9973b5cda4850acea9b2be7684d13	<input type="checkbox"/>	\$colName: \$valueFact



Fields and Data

This display includes:

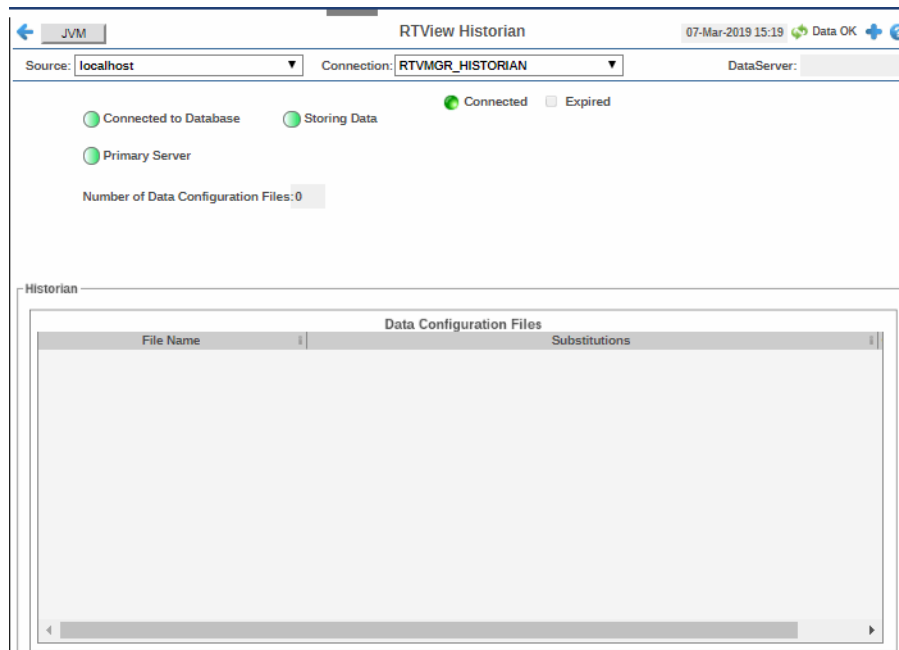
Source	Select the type of connection to the RTView Server.
Connection	Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
Connected	The Display Server connection state: <div> ● Disconnected. ● Connected. </div>
Expired	This server has been marked as expired after no activity.
Function Stats	Opens the RTView Function Stats display which shows detailed performance statistics for RTView functions in the selected Display Server. This button is only enabled if the RTVMGR has a JMX connection defined for the selected Display Server.
Display Timeout (seconds)	The amount of time, in seconds, that a display can be kept in memory after the Display Servlet has stopped requesting it. The default is 60 seconds (to allow faster load time when switching between displays).
Image Quality (0-100)	A value between 0 and 100 , which controls the quality of the generated images. If the value is 100 , the Display Server outputs the highest quality image with the lowest compression. If the value is 0 , the Display Server outputs the lowest quality image using the highest compression. The default is 75 .
Number of Active Displays	The total number of displays currently being viewed by a user.
Maximum Number of Active Displays	The maximum number of displays kept in memory. The default is 20 (to optimize memory used by the Display Server).
Sessions with Active Displays	Number of clients accessing the Display Server.
Display Data / Active Displays	


Display Name	The name of the currently open display.
Session	A unique string identifier assigned to each session.
Panel ID	A unique string identifier assigned to each panel. The Display Server loads each display requested by each client into a panel. This ID can be useful in troubleshooting.
Substitutions	Lists the substitutions used for the display.
Last Ref	The amount of time that has elapsed since the display was last requested by a client.
ID	The client ID.
Preloaded	When checked, indicates that the display (.rtv) file is configured in the DISPLAYSERVER.ini file to be preloaded. The history_config option is used to configure display preloading. Preloading a display makes data immediately available. Preloaded displays are not unloaded unless the Display Server is restarted or the display cache is cleared via JMX. This option can be used multiple times to specify multiple displays to preload.

Historian Servers

Track the status of RTView Historian Servers and data configuration file usage. View the caches that are archived by the Historian application, substitution variables associated with the history cache configuration file, as well as the history cache status. You can also stop and start the Historian, and purge data.







Use the available drop-down menus or right-click to filter data shown in the display.





Fields and Data

This display includes:

Source	Select the type of connection to the RTView Server.
Connection	Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
Connected	The Historian Server connection state:  Disconnected.  Connected.
Expired	This server has been marked as expired after no activity.
Connected to Database	The Historian Server database connection state:  Disconnected.  Connected.
Primary Server	When green, indicates that this Historian, when used within a group of Historians, is the primary group member. If the primary member fails or shuts down, the standby member with the highest priority becomes the primary group member. When red, indicates that the Historian is a secondary server. The Historian Server member state:  The Historian Server is a secondary group member.  This Historian is the primary group member.
Number of Data Configuration Files	The number of configuration files that are used by the history cache.
Historian / Data Configuration Files	

File Name	The name of the history cache configuration file.
Substitutions	Lists the substitutions specified in the history cache configuration file.

Version Info

This display provides detailed version information for all of the connected RTView applications. You can view specific applications by filtering data using the **Source**, **Connection**, **Filter Field**, and **Filter Value** fields at the top of the display. This display provides valuable information about the version of each jar that is used in each connected RTView application that can be used to help Technical Support when issues arise. All RTView applications use multiple jars and this display lists the version information for each jar in the application. The **ApplicationConfiguration** column shows the version of the jar that contains the main class

for the application which is also the version that is printed to the console at startup. The **JarConfiguration** shows the version of the jar specified in the **JarName** field. When **ApplicationConfiguration** and **JarConfiguration** do not match, it indicates that the application is using jars from multiple releases of RTView or that the application is using a patched jar. Rows in the table where the **JarConfiguration** does not match the **ApplicationConfiguration** are highlighted in teal.

Note: RTView applications running versions previous to this enhancement will only have one row in the table and will display "version info not supported in this version" in the **ApplicationConfiguration** column.

RTView Application Versions 25-Sep-2015 14:41 Data OK

Source: All Sources Filter Field: Clear
 Connection: All Connections Filter Value: ☐ RegEx ☐ Not Equal

Detailed Version for All Connected RTView Applications
 Rows where the JarConfiguration does not match ApplicationConfiguration are highlighted in teal

Source	Connection	ApplicationName	JarName	ApplicationConfiguration	JarConfiguration	JarVersionNumber
WIN3	SLMON-DISP-5	RTView Display Server	gmsjagentds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjalertds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjcacheds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjcmdbds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjext.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjflash.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjrmxds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjlog4jds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjmodels.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjlapds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjripeds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjrods.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjrtvhistorian.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjrtvviewer.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0

Page 1 of 8 1 - 200 of 1581 items

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.

Fields and Data

This display includes:

- Source** Select a filter value for the **Source** column.
- Connection** Select a filter value for the **Connection** column.
- Filter Field** Select a table column from the drop-down menu to perform a search in: **ApplicationName**, **JarName**, **ApplicationConfiguration**, **JarConfiguration**, **JarVersionNumber**, **JarVersionDate**, **JarReleaseDate**, and **JarMicroVersion**.
 Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table). Double-clicking on a specific field in the table will populate this field with the selected field's content. For example, double-clicking on the **DataServerName** field in one of the rows displays the entire field's content into this field.
- Clear** Clears entries in the **Filter Field** display list, **Filter Value** field, and **Not Equal** check box.

Filter Value	Enter the (case-sensitive) string to search for in the selected Filter Field .
RegEx	Select to use the Filter Value as a regular expression when filtering. When selected, the Not Equal check box displays.
Not Equal	Works in conjunction with the RegEx field. Selecting this check box searches for values in the specified Filter Field that are NOT equal to the value defined in the Filter Value field. For example, if the Filter Field specified is JarMicroVersion , the Filter Value is specified as 317 , and this check box is selected, then only those rows containing JarMicroVersion fields NOT EQUAL to 317 will display. This field is only enabled when the RegEx check box is checked.
Source	The name of the source of the RTView .
Connection	Lists the name of the JMX connection to the RTView application.
Application Name	Lists the name of the application.
JarName	Lists the name of the jar used in the connected application.
Application Configuration	Lists the configuration string of the application. This string contains the main application version that corresponds to the version information printed to the console at startup.
JarConfiguration	Lists the configuration string for the jar.
JarVersionNumber	Lists the version number for the jar.
JarVersionDate	Lists the version date for the jar.
JarReleaseType	Lists the release type for the jar.
JarMicroVersion	Lists the micro version for the jar.
Expired	When checked, this connection is expired due to inactivity.
time_stamp	The time at which the information in the current row was last received.
DataServerName	The name of the RTView Data Server connection.

JVM

The RTView Manager JVM displays present performance data for monitored Java Virtual Machine (JVM) processes. Use these displays to examine the current and historical performance metrics and resource usage of JVMs. Any JVM that is enabled for monitoring can be included in these displays. The displays include summary overviews and detail pages with historical trends.

You can set alert thresholds on performance and resource metrics for your JVMs, including **CPU Percent**, **Memory Used** and **Gc Duty cycle**. Alerts are shown in the [“All JVMs Heatmap”](#) display. Use the detailed JVM displays to investigate further; for example a **Memory Used** alarm might take you to the [“JVM Summary”](#) display to get historical memory use, or a **Gc Duty Cycle** alarm might take you to the [“JVM GC Trends”](#) display. Displays in this View are:

The following JVM Views can be found under **Components** tab > **Processes /JVM Processes > JVM** once RTView Manager is installed:

- [“All JVMs”](#)
- [“Single JVM”](#)

All JVMs

Check the health status for all your JVMs, then drilldown and investigate issues in the “[Single JVM](#)” displays. Displays in this View are:


- “[All JVMs Heatmap](#)”: Heatmap of alert states for all JVM connections
- “[All JVMs Table](#)”: Table of connection details for all JVM connections.

All JVMs Heatmap

View the most critical alert state for all monitored JVM connections for one or all sources, as well as CPU and memory utilization. The heatmap organizes JVM connections by source and host, and uses color to show the most critical Metric value for each JVM connection associated with the selected source. Each rectangle in the heatmap represents a JVM connection. The rectangle size represents the amount of memory reserved for that process; a larger size is a larger value. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Choose one or **All Sources** from the **Source** drop-down menu. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see detailed JVM connection information (including **PID**). Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected connection in the **JVM Summary** display.





Fields and Data

This display includes:

Source Choose one or **All Sources** to display.








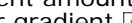
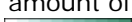
JVM Count The number of JVM connections shown in the display.

Show Inactive Select to show inactive connections.

Connection Select to show JVM connections names.


Metric

Select the Metric to display in the heatmap. Each Metric has a color gradient bar that maps relative values to colors.

Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of 2.  Yellow indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of 1.  Green indicates that no alerts have reached their alert thresholds. Alerts that have not exceeded their specified thresholds have an Alert Severity value of 0.
Alert Count	<p>The number of alerts for the rectangle. The color gradient  bar values range from 0 to the maximum number of alerts in the heatmap.</p>
CPU %	<p>The total percent (%) CPU utilization for the rectangle. The color gradient  bar values range from 0 to the maximum percent (%) CPU utilization in the heatmap.</p>
Memory %	<p>The total percent (%) memory utilization for the rectangle. The color gradient  bar values range from 0 to the maximum percent (%) memory utilization in the heatmap.</p>
Current Heap	<p>The current amount of heap committed for the connection, in kilobytes. The color gradient  bar values range from 0 to the maximum amount in the heatmap.</p>
Used Heap	<p>The total amount of heap used by the connection, in kilobytes. The color gradient  bar values range from 0 to the maximum amount used in the heatmap.</p>

All JVMs Table

View JVM connection details for one or all sources, the most critical alert state for each JVM connection, as well as CPU and memory utilization in a tabular format. Each row in the table is a different connection.

Choose one or **All Sources** from the **Source** drop-down menu. Check the **Show Inactive** box to include inactive connections. The row color for inactive connections is dark red. Click Sort  to order column data. Drill-down and investigate by clicking a row in the table to view details for the selected connection in the **JVM Summary** display.

Heatmap All JVMs - Table View 19-Jan-2017 14:01 Data OK

Source: All Sources

JVM Count: 56 ☒ Show Inactive

All JMX Connections

Connection	Source	Expired	Connected	Alert Severity	Alert Count	Host	Port
ALERT_SERVER	localhost	<input type="checkbox"/>		0	0	localhost	10023
ALERT_SERVER	TBSender	<input type="checkbox"/>		0	0	localhost	10023
ALERTHISTORIAN	localhost	<input type="checkbox"/>		0	0	localhost	10025
ALERTHISTORIAN	TBSender	<input type="checkbox"/>		0	0	localhost	10025
AMXMON-alpha-TB34	localhost	<input type="checkbox"/>		0	0	192.168.200.34	6368
AMXMON-alpha-TB34	TBSender	<input type="checkbox"/>		0	0	192.168.200.34	6368
AMXMON-alpha-TB34-HIST	localhost	<input type="checkbox"/>		0	0	192.168.200.34	6367
AMXMON-beta-TB3-HIST	localhost	<input type="checkbox"/>		0	0	192.168.200.133	6367
BWMON-alpha-TB34	localhost	<input type="checkbox"/>		0	0	192.168.200.34	3368
BWMON-alpha-TB34	TBSender	<input type="checkbox"/>		0	0	192.168.200.34	3368
BWMON-alpha-TB34-HIST	localhost	<input type="checkbox"/>		0	0	192.168.200.34	3367
BWMONITOR-release-WIN-8	localhost	<input type="checkbox"/>		0	0	192.168.200.146	3368
BWMONITOR-TB8	localhost	<input type="checkbox"/>		0	0	192.168.200.138	3368
CONFIG_SERVER	localhost	<input type="checkbox"/>		0	0	localhost	10013
CONFIG_SERVER	TBSender	<input type="checkbox"/>		0	0	localhost	10013
DISPLAYSERVER	localhost	<input type="checkbox"/>		0	0	localhost	10024
DISPLAYSERVER	TBSender	<input type="checkbox"/>		0	0	localhost	10024
DISPLAYSERVER_DARKSTY	localhost	<input type="checkbox"/>		0	0	localhost	10124
DISPLAYSERVER_DARKSTY	TBSender	<input type="checkbox"/>		0	0	localhost	10124
EMSMON_SENDER-alpha-TB	TBSender	<input type="checkbox"/>		0	0	192.168.200.34	3166
EMSMON_SENDER-alpha-TB	localhost	<input type="checkbox"/>		0	0	192.168.200.34	3166

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.

Row Color Code:

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

Fields and Data






This display includes:

Source Choose one or **All Sources** to display.

JVM Count: The number of JVM connections in the table.

Show Inactive Select to include inactive connections.

All JMX Connections Table

Connection	The name of the JVM connection.
Source	The name of the source.
Expired	When checked, this connection is expired due to inactivity.
Connected	The data connection state:  Disconnected.  Connected.
Alert Severity	The maximum level of alerts associated with the connection. Values range from 0 to 2 , where 2 is the greatest Alert Severity.  One or more alerts associated with the connection exceeded their ALARM LEVEL threshold.  One or more alerts associated with the connection exceeded their WARNING LEVEL threshold.  No alerts associated with the connection have exceeded their thresholds.
Alert Count	The current number of alerts for the connection.
Host	The name of the host for this connection.
Port	The port number for the connection.
PID	The connection PID.
CPU %	The amount of CPU, in percent (%) used by this connection.
Max Heap	The maximum amount of heap used by this connection, in kilobytes.
Current Heap	The current amount of committed heap for this connection, in kilobytes.
Used Heap	The current amount of heap used by this connection, in kilobytes.
Mem % Used	The amount of JVM memory used by this connection, in percent (%).
RtvAppType	The type of RTView application, where: 1 is for the Historian, 3 is for the Data Server; 5 is for the Display Server, and 0 is a non-RTView application.
Source	The Data Server that sent this value.
time_stamp	The date and time this row of data was last updated.

Single JVM

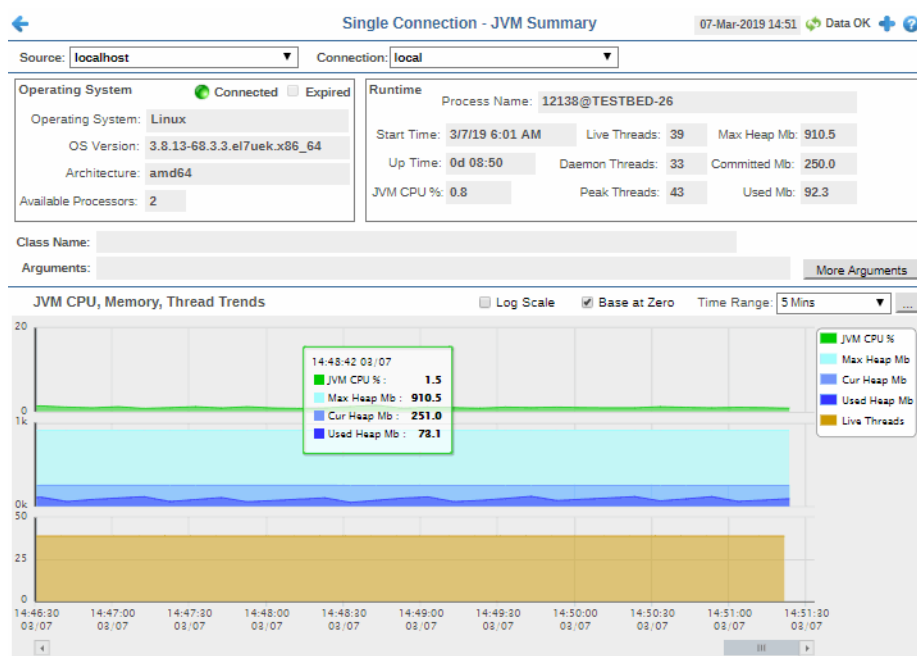
Use these detailed JVM displays to investigate performance issues on a JVM. Displays in this View are:

- **"JVM Summary"**: Table of connection details for a single JVM as well as performance trend graphs.
- **"JVM System Properties"**: Table of system details for a single JVM.
- **"JVM Memory Pool Trends"**: Trend graphs of memory pool utilization.
- **"JVM GC Trends"**: Trend graphs of garbage collection memory utilization.

JVM Summary

Track JVM memory and CPU usage, get JVM system information, application performance metrics, and input arguments for a single connection. Verify whether the memory usage has reached a plateau. Or, if usage is getting close to the limit, determine whether to allocate more memory.

Use the available drop-down menus or right-click to filter data shown in the display.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

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.

Fields and Data

This display includes:

Source Select the type of connection to the RTView Server.


Connection Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.

Operating System

Displays data pertaining to the operating system running on the host on which the JVM resides.

Connected The data connection state:

 Disconnected.

 Connected.

Expired When checked, this server is expired due to inactivity.

Operating System The name of the operating system running on the host on which the JVM resides.


OS Version The operating system version.

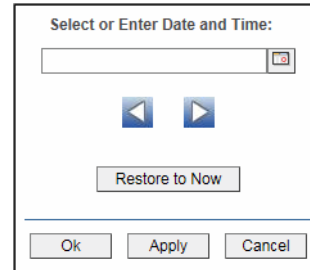
Architecture The ISA used by the processor.


Available Processors The total number of processors available to the JVM.



Runtime

Process Name	Name of the process.
Start Time	The date and time that the application started running.
Up Time	<p>The amount of time the application has been running, in the following format:</p> <p>0d 00:00</p> <p><days>d <hours>:<minutes>:<seconds></p> <p>For example:</p> <p>10d 08:41:38</p>
JVM CPU %	The amount of CPU usage by the JVM, in percent.
Live Threads	The total number of live threads.
Daemon Threads	The total number of live daemon threads.
Peak Threads	The total number of peak live threads since the JVM started or the peak was reset.
Max Heap Mb	<p>The maximum amount of memory used for memory management by the application in the time range specified. This value may change or be undefined.</p> <p>NOTE: A memory allocation can fail if the JVM attempts to set the Used memory allocation to a value greater than the Committed memory allocation, even if the amount for Used memory is less than or equal to the <i>Maximum</i> memory allocation (for example, when the system is low on virtual memory).</p>
Committed Mb	<p>The amount of memory, in megabytes, guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for Committed memory could be less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for Used memory.</p>
Used Mb	The amount of memory currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.
Class Name	Class name used for JVM.
Arguments	The arguments used to start the application.
More Arguments	Additional arguments used to start the application.
JVM CPU, Memory, Thread Trends Shows JVM metrics for the selected server.	

- Log Scale** Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

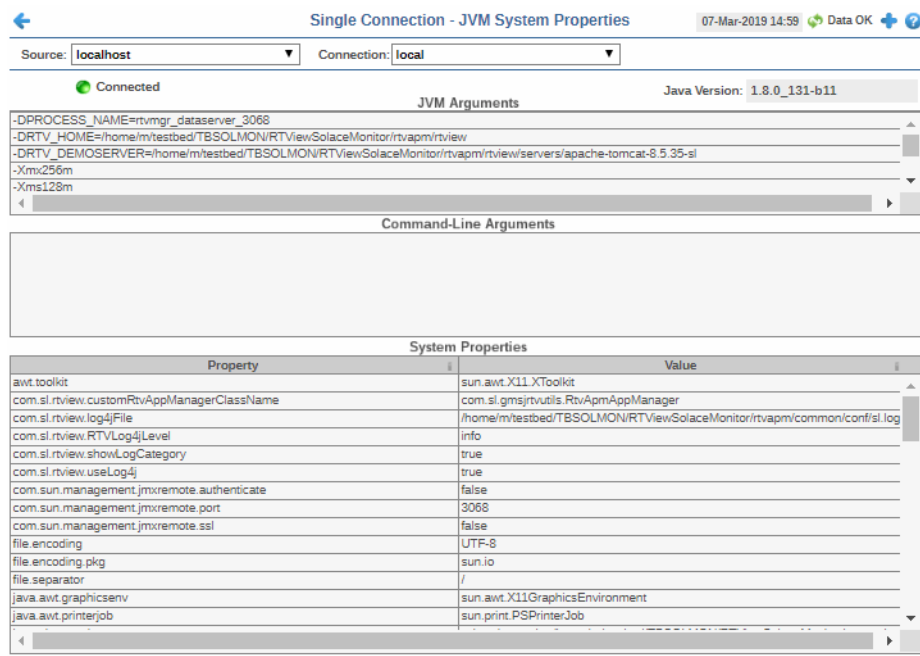
Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- JVM CPU %** Traces the amount of memory, in percent, used by the JVM in the time range specified.
- Max Heap Mb** Traces the maximum amount of memory used for memory management by the application in the time range specified. This value may change or be undefined.
NOTE: A memory allocation can fail if the JVM attempts to set the **Used** memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the **Maximum** memory allocation (for example, when the system is low on virtual memory).
- Cur Heap Mb** Traces the current amount of memory, in megabytes, used for memory management by the application in the time range specified.
- Used Heap Mb** Traces the memory currently used by the application.
- Live Threads** Traces the total number of currently active threads in the time range specified.

JVM System Properties

Track JVM input arguments and system properties for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Fields and Data

This display includes:

Source	Select the type of connection to the RTView Server.
Connection	Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
Connected	The data connection state: <ul style="list-style-type: none"> Disconnected. Connected.
Java Version	The Java version running on the selected server.
JVM Arguments	The JVM arguments in the RuntimeMXBean InputArguments attribute.
Command Line Arguments	Arguments used to start the application.

System Properties


This table lists and describes system property settings.

Property Name of the property.

Value Current value of the property.

JVM Memory Pool Trends

Track JVM heap and non-heap memory usage for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▼, Table open commonly accessed displays.

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.

Fields and Data

This display includes:

Source Select the type of connection to the RTView Server.

Connection Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.

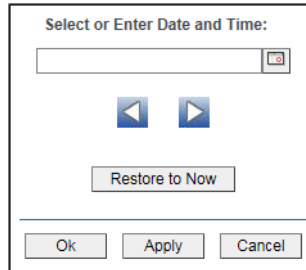
Connected The data connection state:


● Disconnected.



● Connected.

Base at Zero Use zero as the Y axis minimum for all graph traces.

Time Range Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Heap Memory

Maximum	<p>The maximum amount of memory used, in megabytes, for memory management by the application in the time range specified. This value may change or be undefined.</p> <p>NOTE: A memory allocation can fail if the JVM attempts to set the Used memory allocation to a value greater than the Committed memory allocation, even if the amount for Used memory is less than or equal to the Maximum memory allocation (for example, when the system is low on virtual memory).</p>
Committed	<p>The amount of memory, in megabytes, guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for Committed memory could be less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for Used memory.</p>
Used	<p>The amount of memory, in megabytes, currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.</p>
Peak Tenured Used	<p>The amount of memory, in megabytes, used by tenured JVM objects in the time range specified. Tenured refers to JVM objects contained in a pool that holds objects that have avoided garbage collection and reside in the survivor space. Peak tenured refers to the maximum value of the tenured memory over a specified period of time.</p>
Eden Space	<p>Traces the amount of memory used by the JVM eden pool in the time range specified. Eden refers to the JVM eden pool, which is used to initially allocate memory for most objects.</p>
Survivor Space	<p>Traces the amount of memory used by the JVM survivor pool in the time range specified. The JVM survivor pool holds objects that survive the eden space garbage collection.</p>
Tenured Gen	<p>Traces the amount of memory used by tenured JVM objects in the time range specified. Tenured refers to JVM objects contained in a pool that holds objects that have avoided garbage collection and reside in the survivor space. Peak tenured refers to the maximum value of the tenured memory over a specified period of time.</p>

Non-Heap Memory


Maximum	The maximum amount of memory, in megabytes, used for JVM non-heap memory management by the application in the time range specified.
Committed	The amount of memory, in megabytes, guaranteed to be available for use by JVM non-heap memory management. The amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for Committed memory could be less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for Used memory.
Used	The amount of memory, in megabytes, currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.
Objects Pending Finalization	The value of the MemoryMXBean ObjectPendingFinalizationCount attribute.
Verbose	The value of the MemoryMXBean Verbose attribute.
Code Cache	Traces the amount of non-heap memory used in the JVM for compilation and storage of native code.
Perm Gen	Traces the amount of memory used by the pool containing reflective data of the virtual machine, such as class and method objects. With JVMs that use class data sharing, this generation is divided into read-only and read-write areas.

Operations

Run Garbage Collector	Performs garbage collection on the selected server.
Reset Peak Usage	Clears peak usage on the selected server.

JVM GC Trends

Track JVM garbage collection memory usage for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.



Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

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.

Fields and Data


This display includes:


- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Garbage Collector** Select a garbage collection method: **Copy** or **MarkSweepCompact**.
- Max** Shows the maximum amount of memory used for JVM garbage collection in the time range specified.
- Committed** Shows the amount of memory guaranteed to be available for use by JVM non-heap memory management. The amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. **Committed** memory will always be greater than or equal to the amount allocated for **Used** memory.


Base at Zero



Use zero as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Memory Usage (in MB) Before and After Garbage Collection**Maximum**

Traces the maximum amount of memory used by garbage collection in the time range specified. This value may change or be undefined.

NOTE: A memory allocation can fail if the JVM attempts to set the **Used** memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the **Maximum** memory allocation (for example, when the system is low on virtual memory).

Committed

Traces the amount of memory guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. **Committed** memory will always be greater than or equal to the amount allocated for **Used** memory.

Used - Before

Traces the amount of memory used before the last garbage collection.

Used - After

Traces the amount of memory used after the last garbage collection.

Duration

The duration, in seconds, of garbage collection.

Duty Cycle

The percentage of time that the application spends in garbage collection.

Tomcat

The Tomcat displays provide extensive visibility into the health and performance of Tomcat application servers and installed web modules. The following Tomcat Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers** > **Tomcat**. The Tomcat displays come with RTView Enterprise Monitor.

Tomcat has the following Views:

- [“Tomcat Servers”](#)
- [“Tomcat Applications”](#)

Tomcat Servers

These displays present performance data for monitored Tomcat Servers. Use these displays to examine the state and performance of your Tomcat servers as well as all installed web modules. The server displays include summary overviews and detail pages with historical trends. Displays in this View are:

- [“All Tomcat Servers”](#): Table of connection details and performance metrics for all Tomcat connections.
- [“Tomcat Server Summary”](#): Performance metrics for one Tomcat Server, including current and historic performance metrics.

All Tomcat Servers

View Tomcat Server details per connection such as the total number of sessions, bytes sent/received, and processing time. Each row in the table is a different Tomcat Server. The row color for inactive connections is dark red.

Use this display to get Tomcat server session counts, access and request rates, cache hit rates and data transmission metrics.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the **Service Summary** display.

Connection	Source	Sessions Active	Sessions Total	Sessions Expired	Accesses per sec	Accesses Total	Bytes Rcvd per sec	Bytes Rcvd Total
TOMCAT	localhost	4	17	13	1.4	30,302	603.1	433,851.8

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.
- 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.

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.

Fields and Data

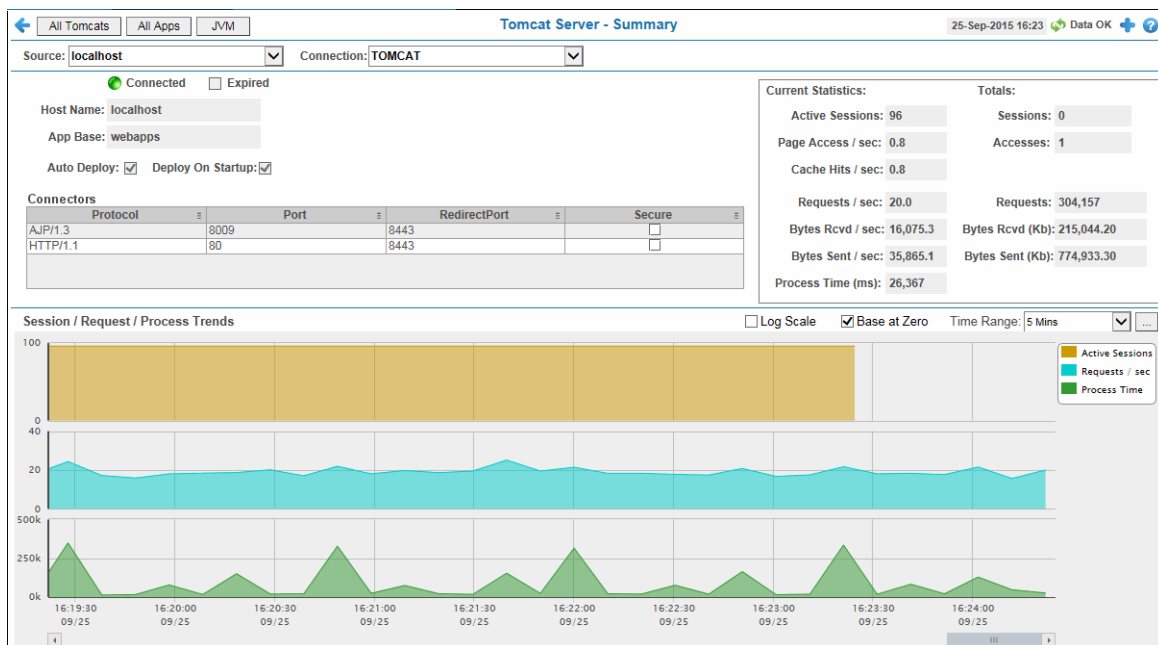
This display includes:


Tomcat Count	The number of Tomcat connections in the table.
Connection	The name of the Tomcat connection.
Source	The host where the Tomcat Server is running.
Sessions Active	The number of currently active client sessions.
Sessions Total	The total number of client sessions since the server was started.
Sessions Expired	The total number of client sessions that expired since the server was started.
Accesses per sec	The number of times pages are accessed, per second.
Accesses Total	The total number of times pages have been accessed since the server was started.
Bytes Rcvd per sec	The number of bytes received per second.
Bytes Rcvd Total	The total number of bytes received since the server was started.

Bytes Sent per sec	The number of bytes sent per second.
Bytes Sent Total	The total number of bytes sent since the server was started.
Cache Hit Rate	The number of times the cache is accessed, per second.
Requests per sec	The number of requests received, per second.
Requests Total	The total number of requests received since the server was started.
Process Time	The average amount of time, in milliseconds, to process requests.
Error Count	The number of errors that have occurred since the server was started.
appBase	The directory in which Tomcat is installed.
Display Name	The name of the currently open display.
Expired	When checked, this connection is expired due to inactivity.
time_stamp	The date and time this row of data was last updated. Format: MM/DD/YY HH:MM:SS <month>/ <day>/<year> <hours>:<minutes>:<seconds>

Tomcat Server Summary

Track the performance of one Tomcat Server and get Tomcat hosting and connection details. You can drill down to this display from the Servers table for detailed information and historical trends for a specific server. The trends include Active Sessions, Requests per Sec, and Process Time.





Fields and Data

This display includes:

- Source** Select the host where the Tomcat Server is running.
- Connection** Select a Tomcat Server from the drop-down menu.
- Connected** The Tomcat Server connection state:
 - Disconnected.
 - Connected.
- Expired** When checked, this server is expired due to inactivity.
- Host Name** The name of the host where the application resides.
- App Base** The directory in which Tomcat modules are installed.
- Auto Deploy** When checked, indicates that the Tomcat option, automatic application deployment, is enabled.

Note: This Tomcat option is set using the **autoDeploy** property in the **server.xml** file, located in the Tomcat **conf** directory. **autoDeploy=true** enables the option.
- Deploy On Startup** When checked, indicates that the option to deploy the application on Tomcat startup is enabled.

Note: This Tomcat option is set using the **deployOnStartup** property in the **server.xml** file, located in the Tomcat **conf** directory. When enabled (**deployOnStartup=true**), applications from the host are automatically deployed.

Connectors

This table shows Tomcat application connection information.


Protocol	The protocol used by the Tomcat application on the host.
Port	The port number used by the Tomcat application on the host.
RedirectPort	The redirect port number used by the Tomcat application on the host.
Secure	When checked, specifies that the Tomcat application uses a secure connection on the host.

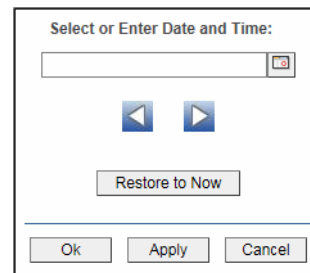
Current Statistics / Totals

Active Sessions	The number of clients currently in session with the servlet.
Sessions	The total number of client sessions since the server was started.
Page Access / sec	The number of times pages are accessed, per second.
Accesses	The total number of page accesses since the server was started.
Cache Hits / sec	The number of times the cache is accessed, per second.
Requests / sec	The number of requests received, per second.
Requests	The total number of requests since the server was started.
Bytes Rcvd / sec	The number of bytes received, per second.
Bytes Rcvd (Kb)	The number of kilobytes received since the server was started.
Bytes Sent / sec	The number of bytes sent, per second.
Bytes Sent (Kb)	The total number of kilobytes sent since the server was started.
Process Time	The amount of time, in milliseconds, for the servlet to process client requests.


Session / Request / Process Trends

Shows metrics for the selected server.

- Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "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   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- Active Sessions** Traces the number of currently active client sessions.
- Requests /sec** Traces the number of requests received, per second.
- Process Time** Traces the average amount of time, in milliseconds, to process requests.

Tomcat Applications

These displays present performance data for monitored Tomcat Applications. Use these displays to examine the state and performance of your Tomcat applications and all installed web modules. The server displays include summary overviews and detail pages with historical trends.

Displays in this View are:

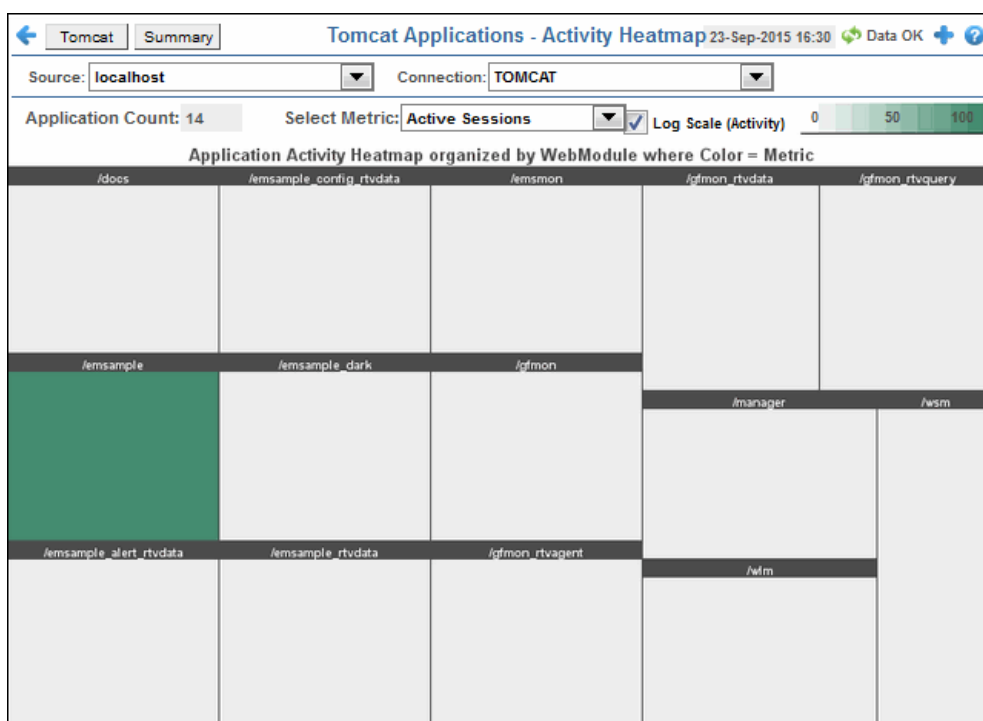
- **"Applications Heatmap"**: Heatmap of performance metrics for all Web modules for one Tomcat Server.
- **"Applications Summary"**: Table and trend graphs of performance metrics for Web modules.

Applications Heatmap

View performance metrics for all monitored Tomcat Web modules for one Tomcat Server. The heatmap organizes Tomcat Web modules by server, and uses color to show the most critical Metric value for each Tomcat connection associated with the selected source. Each rectangle in the heatmap represents a Web module. In this heatmap, the rectangle size is the same for all Web modules. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Use this display to see at-a-glance the health of all your web applications. You can select the heatmap color metric from a list including active sessions, access rate, and total access count.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes ☒ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Web module in the **Application Summary** display.



Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- 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.

Fields and Data

This display includes:

Source Select the host where the Tomcat Server is running.

Connection Select a Tomcat Server from the drop-down menu.

Application Count The number of Tomcat applications in the heatmap.

Log Scale (Activity) Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

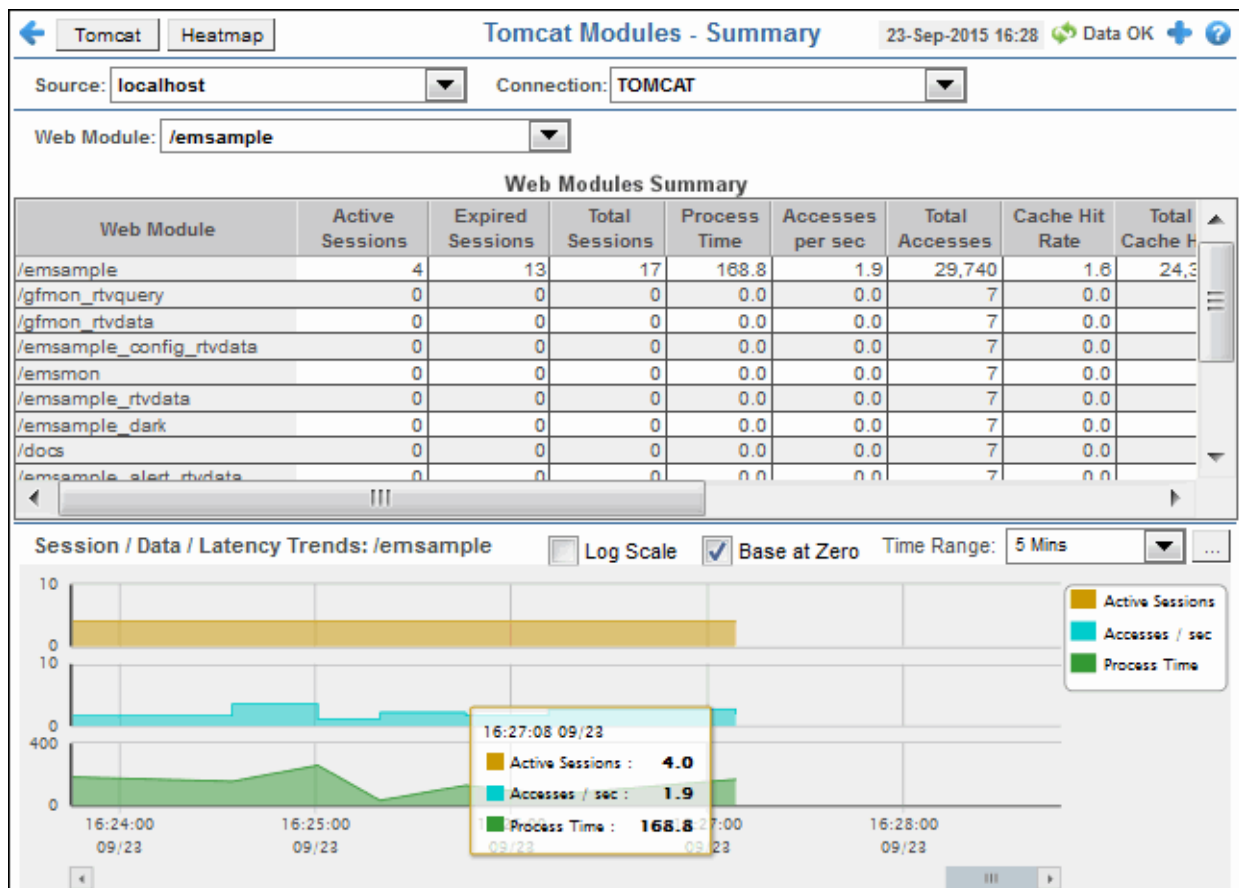
Select Metric Select the metric to display in the heatmap. Each Metric has a color gradient bar that maps relative values to colors.


Applications Summary

Track the performance of all web application modules in a server and view utilization details. The table summarizes the sessions, accesses, cache hit and so forth, for all installed web modules. Each row in the table is a different web application module. The row color for inactive modules is dark red. Select a web application module to view metrics in the trend graph.

Use this data to verify response times of your Web application modules.

Use the available drop-down menus or right-click to filter data shown in the display.





Fields and Data

This display includes:

- Source** Select the host where the Tomcat Server is running.
- Connection** Select a Tomcat Server from the drop-down menu. This menu is populated by the selected Source.
- Web Module** Select a Web module from the drop-down menu. This menu is populated by the selected Connection. The Web Module you select populates the trend graphs.

Web Module Summary

Web Module	The name of the Web module.
Sessions Active	The number of currently active client sessions.
Sessions Total	The total number of client sessions since the application was started.
Sessions Expired	The total number of client sessions that expired since the application was started.
Accesses per sec	The number of times pages are accessed, per second.
Accesses Total	The total number of times pages have been accessed since the application was started.
Bytes Rcvd per sec	The number of bytes received per second.
Bytes Rcvd Total	The total number of bytes received since the application was started.
Bytes Sent per sec	The number of bytes sent per second.
Bytes Sent Total	The total number of bytes sent since the application was started.
Cache Hit Rate	The number of times the cache is accessed, per second.
Requests per sec	The number of requests received, per second.
Requests Total	The total number of requests received since the application was started.
Process Time	The average amount of time, in milliseconds, to process requests.
Error Count	The number of errors occurred since the application was started.
appBase	The directory in which Tomcat is installed.

Expired When checked, this connection is expired due to inactivity.

time_stamp The date and time this row of data was last updated.

Format:

MM/DD/YY HH:MM:SS

<month>/ <day>/<year>

<hours>:<minutes>:<seconds>

Session/Data/Latency Trends

Shows metrics for the selected Web module. The Web module can be selected from the **Web Module** drop-down menu or the **Web Modules Summary** table.


Log Scale

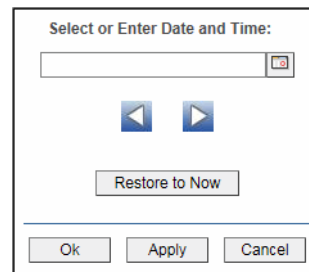
Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.


Base at Zero



Use zero as the Y axis minimum for all graph traces.

Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

Active Sessions

Traces the number of currently active client sessions.

Accesses /sec

Traces the number of times pages are accessed, per second.

Process Time

Traces the average amount of time, in milliseconds, to process requests.

APPENDIX A **Alert Definitions**

This section describes alerts that are available with RTView Enterprise Monitor per solution package. This section includes:

- ["Amazon Web Services"](#)
- ["Apache Kafka"](#)
- ["Docker"](#)
- ["Microsoft SQL Server"](#)
- ["MongoDB"](#)
- ["MySQL Database"](#)
- ["Node.js"](#)
- ["Oracle Coherence"](#)
- ["Oracle Database"](#)
- ["Oracle WebLogic"](#)
- ["RTView Manager and RTView Rules"](#)
- ["RTView Host Agent"](#)
- ["Solace"](#)
- ["TIBCO ActiveMatrix BusinessWorks"](#)
- ["TIBCO ActiveSpaces"](#)
- ["TIBCO Adapters"](#)
- ["TIBCO BusinessEvents"](#)
- ["TIBCO Enterprise Message Service"](#)
- ["TIBCO FTL"](#)
- ["UX"](#)
- ["VMware vCenter"](#)

Amazon Web Services

The following alerts are available for Amazon Web Services. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
AcwInstanceCpuHigh Executes a single warning and single alarm if the CPU used exceeds the specified threshold. Index Type: PerInstance Metric: CPUUtilization	70	80	30	FALSE
AcwInstanceDiskReadBytesHigh Executes a single warning and single alarm if the number of bytes read from the disk exceeds the specified threshold. Index Type: PerInstance Metric: DiskReadBytes	10000	20000	30	FALSE
AcwInstanceDiskReadOpsHigh Executes a single warning and single alarm if the number of disk reads exceeds the specified threshold. Index Type: PerInstance Metric: DiskReadOps	100	200	30	FALSE
AcwInstanceDiskWriteBytesHigh Executes a single warning and single alarm if the number of bytes written to the disk exceeds the specified threshold. Index Type: PerInstance Metric: DiskWriteBytes	1000000	2000000	30	FALSE
AcwInstanceDiskWriteOpsHigh Executes a single warning and single alarm if the number of disk writes exceeds the specified threshold. Index Type: PerInstance Metric: DiskWriteOps	100	200	30	FALSE
AcwInstanceNetworkReadBytesHigh Executes a single warning and single alarm if the number of bytes read from the network exceeds the specified threshold. Index Type: PerInstance Metric: NetworkIn	1000000	20000	30	FALSE
AcwInstanceNetworkWriteBytesHigh Executes a single warning and single alarm if the number of bytes written across the network exceeds the specified threshold. Index Type: PerInstance Metric: NetworkOut	10000	20000	30	FALSE

Apache Kafka

The following alerts are available for Apache Kafka. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
KafkaBrokerBytesInPerSecHigh The number of incoming bytes per second exceeds the defined threshold for the broker. Index Type(s): PerKafkaServer	1600	2000	30	FALSE
KafkaBrokerBytesOutPerSecHigh The number of outgoing bytes per second exceeds the defined threshold for the broker. Index Type(s): PerKafkaServer	1600	2000	30	FALSE
KafkaBrokerExpired The Kafka Broker is not responding. Index Type(s): PerKafkaServer	NaN	NaN	30	FALSE
KafkaBrokerLogFlushLatency95PHigh The current log flush latency exceeds the 95th percentile. Index Type(s): PerKafkaServer	1600	2000	30	FALSE
KafkaBrokerMsgsInPerSecHigh The number of incoming messages per second exceeds the defined threshold for the broker. Index Type(s): PerKafkaServer	1600	2000	30	TRUE
KafkaBrokerNetProcAvgIdlePctHigh The average percent idle for the network processor exceeds the threshold. Index Type(s): PerKafkaServer	.05	.3	30	FALSE
KafkaBrokerNetProcAvgIdlePctLow The average percent idle for the network processor is below the threshold. Index Type(s): PerKafkaServer	.05	.3	30	FALSE
KafkaBrokerOfflinePartitionCountHigh The number of partitions without an active leader is not zero. Index Type(s): PerKafkaServer	NaN	1	30	TRUE
KafkaBrokerReqHdlrAvgIdleHigh The average fraction of time the broker request handler is idle exceeds the threshold. Index Type(s): PerKafkaServer	.5	.3	30	TRUE
KafkaBrokerUncInLderElecsPerSecHigh The available replicas were not in sync during leader election. Data loss has probably occurred. Index Type(s): PerKafkaServer	1600	2000	30	FALSE

KafkaBrokerUnderReplicatedPartnsHigh The number of under-replicated partitions is not zero. Index Type(s): PerKafkaServer	NaN	1	30	FALSE
KafkaClusterLeadersUnbalancedHigh The partition leaders for the cluster are not evenly distributed across the available brokers. Index Type(s): PerKafkaCluster	10	10	30	FALSE
KafkaClusterNoActiveController There is more than one active controller per cluster, which could indicate a split-brain error. Index Type(s): PerKafkaCluster	NaN	NaN	30	FALSE
KafkaClusterPartitionsUnbalancedHigh Partitions supported by the cluster are not evenly distributed across the available brokers. Index Type(s): PerKafkaCluster	10	1	30	FALSE
KafkaClusterSplitBrain One (or more) zookeeper/broker is not acting as part of the main cluster. Index Type(s): PerKafkaCluster	NaN	NaN	30	FALSE
KafkaConsumerBytesPerSecHigh The consumer message load (bytes per second) exceeds the threshold. Index Type(s): PerKafkaConsumer	1600	2000	30	FALSE
KafkaConsumerExpired The consumer is not responding. Index Type(s): PerKafkaConsumer	NaN	NaN	30	FALSE
KafkaConsumerFetchLatencyHigh The consumer fetch latency exceeds the threshold. Index Type(s): PerKafkaConsumer	1600	2000	30	TRUE
KafkaConsumerFetchRateHigh The consumer is pulling records from Kafka at a slower than expected rate. Index Type(s): PerKafkaConsumer	1600	2000	30	FALSE
KafkaConsumerLagIncreasing The consumer lag rate of change is greater than zero for the specified duration, which could mean that lag is steadily increasing. Index Type(s): PerKafkaConsumer	NaN	NaN	300	FALSE
KafkaConsumerMaxLagHigh The consumer is falling too far behind the producer. Index Type(s): PerKafkaConsumer	1600	2000	30	TRUE
KafkaConsumerMsgsPerSecHigh The consumer message load (messages per second) exceeds the threshold. Index Type(s): PerKafkaConsumer	1600	2000	30	TRUE

KafkaConsumerPartitionStalled The consumer lag delta is not negative and the current offset delta is positive for the defined duration for a topic on a partition, which could mean that new messages are being added to the partition but the consumer is not reading them. Index Type(s): PerKafkaConsumer	NaN	NaN	300	FALSE
KafkaConsumerSlow This alert is triggered for a topic when consumer lag delta is not negative and the current offset delta is positive for the specified duration, which could mean that the consumer is slow in reading messages. Index Type(s): PerKafkaConsumer	NaN	NaN	300	FALSE
KafkaProducerExpired The producer is not responding. Index Type(s): PerKafkaProducer	NaN	NaN	30	FALSE
KafkaProducerIncomingByteRateHigh The producer's incoming byte rate exceeds the threshold. Index Type(s): PerKafkaProducer	1600	2000	30	TRUE
KafkaProducerIoWaitTimeMSHigh The producer is waiting for IO longer than expected (on average). Index Type(s): PerKafkaProducer	1600	2000	30	FALSE
KafkaProducerOutgoingByteRateHigh The producer output byte rate exceeds the threshold. Index Type(s): PerKafkaProducer	1600	2000	30	TRUE
KafkaProducerRequestLatencyHigh The producer request latency exceeds the threshold. Index Type(s): PerKafkaProducer	1600	2000	30	TRUE
KafkaProducerRequestRateHigh The producers request rate exceeds the threshold. Index Type(s): PerKafkaProducer	1600	2000	30	TRUE
KafkaProducerResponseRateHigh The producer response rate exceeds the threshold. Index Type(s): PerKafkaProducer	1600	2000	30	TRUE
KafkaZookeeperAvgLatencyHigh The average time for the zookeeper to respond to a request exceeds the threshold. Index Type(s): PerKafkaZookeeper	1600	2000	30	TRUE
KafkaZookeeperExpired The zookeeper is not responding. Index Type(s): PerKafkaZookeeper	NaN	NaN	30	FALSE

KafkaZookeeperNumAliveConnsHigh The total number of connections to a given zookeeper exceeds the threshold. Index Type(s): PerKafkaZookeeper	1600	2000	30	TRUE
KafkaZookeeperOutstandingReqsHigh Clients are making requests faster than the zookeeper can process them. Index Type(s): PerKafkaZookeeper	1600	2000	30	FALSE
KafkaZookeeperRatePktsRcvdHigh The rate that the zookeeper is receiving packets exceeds the threshold. Index Type(s): PerKafkaZookeeper	1600	2000	30	TRUE
KafkaZookeeperRatePktsSentHigh The rate that the zookeeper is sending packets exceeds the threshold. Index Type(s): PerKafkaZookeeper	1600	2000	30	TRUE

Docker

The following alerts are available for Docker. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
DocContainerCpuUsageHigh A Docker Container's CPU usage is above the defined threshold. Index Type(s): PerContainer Metric: cpu.usage	24	50	30	FALSE
DocContainerExpired A Docker Container has expired. Index Type(s): PerContainer Metric: Expired	NaN	NaN	30	FALSE
DocContainerNetBytesInHigh A Docker Container's incoming network data rate is above the defined thresholds. Index Type(s): PerContainer Metric: net.rxbytes.avg	750000	1000000	30	FALSE
DocContainerNetBytesOutHigh A Docker Container's outgoing network data rate is above the defined thresholds. Index Type(s): PerContainer Metric: net.txbytes.avg	750000	1000000	30	FALSE

DocEngineCpuUsageHigh A Docker Engine's CPU usage is above the defined thresholds. Index Type(s): PerEngine Metric: cpu.usage	50	75	30	TRUE
DocEngineExpired A Docker Engine has expired. Index Type(s): PerEngine Metric: Expired	NaN	NaN	30	FALSE
DocEngineNetBytesInHigh A Docker Engine's incoming network data rate is above the defined thresholds. Index Type(s): PerEngine Metric: net.rxbytes.avg	750000	1000000	30	TRUE
DocEngineNetBytesOutHigh A Docker Engine's outgoing network data rate is above the defined thresholds. Index Type(s): PerEngine Metric: net.txbytes.avg	750000	1000000	30	TRUE

Microsoft SQL Server

The following alerts are available for Microsoft SQL Server. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
MssqlInstanceDeadlocksDetected The number of current deadlocks has exceeded its threshold. Index Type(s): PerServer Metric: DeltaNumber of Deadlocks	1	2	0	TRUE
MssqlInstanceLatchWaitsHigh The number of current latch waits has exceeded its threshold. Index Type(s): PerServer Metric: DeltaLatch Waits	15	30	0	TRUE
MssqlInstanceLockWaitsHigh The amount of seconds on lock waits has exceeded its threshold. Index Type(s): PerServer Metric: DeltaLock Waits	15	30	0	TRUE

MssqlInstancePacketErrorsDetected The amount of current packet errors has exceeded its threshold. Index Type(s): PerServer Metric: DeltaPacket Errors	1	2	0	TRUE
MssqlInstanceSqlCpuUsedHigh The percentage of CPU utilization on SQL processing has exceeded its threshold. Index Type(s): PerServer Metric: CPU Util	10	15	0	TRUE
MssqlInstanceUsedMemoryHigh The percentage of memory used by the SQL Server has exceeded its threshold. Index Type(s): PerServer Metric: Memory In Use (%)	65	85	0	TRUE

MongoDB

The following alerts are available for MongoDB. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
MongoCollectionExpired A collection was not able to be contacted for longer than the normal expiration window. Index Type(s): PerCollection Metric: Expired	NaN	NaN	30	FALSE
MongoCollectionNumObjectsHigh The number of objects for the collection exceeds a given threshold. Index Type(s): PerCollection Metric: numberOfObjects	1600	2000	30	FALSE
MongoDatabaseDataSizeHigh The database size for the database exceeds a given threshold. Index Type(s): PerDatabase Metric: dataSize	80000	100000	30	FALSE
MongDatabaseExpired The database was not able to be contacted for longer than the normal expiration window. Index Type(s): PerDatabase Metric: Expired	NaN	NaN	30	FALSE

MongoInstanceExpired The instance was not able to be contacted for longer than the normal expiration window. Index Type(s): PerInstance Metric: Expired	60	80	30	FALSE
MongoInstanceNotConnected The instance was not able to be contacted for longer than the normal expiration window. Index Type(s): PerInstance Metric: connectionStatus	NaN	NaN	30	FALSE
MongoInstanceOpenCursorsHigh The number of Open Cursors for the Instance exceeds a given threshold. Index Type(s): PerInstance Metric: openCursors	160	200	30	TRUE

MySQL Database

The following alerts are available for MySQL Database. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
MysqlBytesReceivedHigh Executes a single warning and a single alarm if the amount of kilobytes received exceeds the specified threshold. Index Type: PerServer Metric: Received	5	10	na	FALSE
MysqlBytesSentHigh Executes a single warning and a single alarm if the amount of kilobytes sent exceeds the specified threshold. Index Type: PerServer Metric: Sent	5	10	na	FALSE
MysqlDelayedWritesHigh Executes a single warning and a single alarm if the number of delayed writes exceeds the specified threshold. This alert only applies to previous versions to MySQL 5.7 as delayed inserts are not supported in later versions. Index Type: PerServer Metric: Delayed Writes	1	2	na	FALSE

MysqlLocksWaitedHigh Executes a single warning and a single alarm if the number of times that requests for a table lock requires a wait before being granted exceeds the specified threshold. Index Type: PerServer Metric: Table_locks_waited	1	2	na	FALSE
MysqlQcacheLowMemPrunesHigh Executes a single warning and a single alarm if the number of queries deleted from the query cache because of low memory exceeds the specified threshold. Index Type: PerServer Metric: Qcache_lowmem_prunes	1	2	na	FALSE
MysqlSlowQueriesHigh Executes a single warning and a single alarm if the number of queries that exceed the number of seconds specified for long_query_time exceeds the specified threshold. Index Type: PerServer Metric: Slow Queries	1	2	na	FALSE
MysqlSlowThreadsHigh Executes a single warning and a single alarm if the number of threads that exceed the number of seconds specified for slow_launch_time to create exceeds the specified threshold. Index Type: PerServer Metric: Slow_launch_threads	1	2	na	FALSE

Node.js

The following alerts are available for Node.js. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
NodeMasterCpuUsageHigh A master node's CPU usage is above the defined thresholds. Index Type(s): PerConnection Metric: Node Master - CPU %	30	50	30	FALSE
NodeMasterExpired A master node has expired. Index Type(s): PerConnection Metric: Node Master - Expired	NaN	NaN	30	FALSE

NodeMasterRequestRateHigh The request rate of a master node is above the defined thresholds. Index Type(s): PerConnection Metric: Node Requests - Requests Per Second	1600	2000	30	FALSE
NodeMasterResponseTimeHigh The response time of a URL is above the defined thresholds. Index Type(s): PerConnection Metric: Node Requests - Avg Response Time	5	10	30	FALSE
NodeProcessCpuUsageHigh A worker node's CPU usage is above the defined thresholds. Index Type(s): PerConnection Metric: Node Processes - CPU Used %	5	50	30	TRUE
NodeProcessExpired A worker node has expired. Index Type(s): PerConnection Metric: Node Processes - Expired	NaN	NaN	30	FALSE
NodeProcessMemUsageHigh A master node's memory usage has exceeded the defined limits. Index Type(s): PerConnection Metric: Node Processes - Memory Used %	90	95	30	TRUE

Oracle Coherence

The following alerts are available with both the solution package and standalone versions for Oracle Coherence.

OcAvailableMemoryLowCluster

A single alert is executed if the average percent memory used over max memory of all nodes in the cluster exceeds the specified thresholds.

OcAvailableMemoryLowNode

For each node in the cluster, an alert is executed if the percent memory used over max memory available for that node exceeds the specified thresholds.

OcAvailableMemoryLowNodeSpike

For each node in the cluster, an alert is executed if the percent memory used exceeds the specified threshold for the percent above average memory used in the previous 24 hours. For example, if the threshold is set to 50% of total memory used, and the average memory consumption on a particular node for the previous 24 hours is 40%, an alert will be executed if current memory usage exceeds 60% of the total.

NOTE: The 24 hour time span (86400 seconds) is controlled by the \$AVERAGE_MEMORY_TIME_WINDOW substitution.

The warning default setting is **115** (percent) of the previous 24 hours and the alarm default setting is **125** (percent) of the previous 24 hours.

By default the alert is disabled.

OcBadCommunicationCluster

A single alert is executed if the average communication failure rate of all nodes in the cluster exceeds the specified thresholds.

OcBadCommunicationNode

For each node in the cluster, an alert is executed if the communication failure rate for that node exceeds the specified thresholds.

OcBadCommunicationNodesInTimeRange

Executes a single warning and a single alert if the percentage of nodes in a cluster exceeds the specified threshold for the BadCommunicationNode alert within a time range specified.

To specify the time range, modify the \$BAD_COMMUNICATION_NODES_TIME_RANGE substitution.

The default time range setting is 5 minutes (300 seconds), the warning default setting is **40** (percent) and the alarm default setting is **50** (percent).

By default the alert is enabled.

OcCacheHitPercentageLow

This alert is executed when the current **Hit%** (total current hits/total current gets) is below the specified threshold for a sampling period and the specified cache(s).

OcCacheQueueSizeHigh

A single alert is executed when the CacheQueueSize for all nodes in the cluster exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is **100** (total objects), Alarm is **200** (total objects) and Duration is **60** (total objects).

OcCacheRateCacheMissesHigh

Executes when the Misses per second exceed the specified threshold and duration. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **2000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

OcCacheRateStoreReadsHigh

Executes when the cache StoreReads rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

OcCacheRateStoreWritesHigh

Executes when the cache StoreWrites rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

OcCacheRateTotalGetsHigh

Executes when the cache total gets rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

OcCacheRateTotalPutsHigh

Executes when the cache DeltaTotalPuts rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

OCCacheSizeHigh

Executes when the number of objects in a cache exceeds the specified threshold. By default the alert is disabled with the following default settings: Warning is **1000** (count), Alarm is **5000** (count) and Duration is **60** (seconds).

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

NOTE: If you want to know when the size of a specific cache exceeds specific thresholds, it might be preferable to use the **Per Cache** or **Per Storage Class** override settings, allowing you set specific thresholds for specific caches.

OCCacheSizeLow

Executes when the number of objects in a cache goes below the specified threshold. By default the alert is disabled with the following default settings: Warning is **1000** (count), Alarm is **5000** (count) and Duration is **60** (seconds).

NOTE: If you want to know when the size of a specific cache goes below specific thresholds, it might be preferable to use the **Per Cache** or **Per Storage Class** override settings, allowing you set specific thresholds for specific caches.

OcCapacityLimitAllCaches

An alert is executed if the percent cache used over cache capacity for any cache in the cluster exceeds the specified thresholds. There is one highWarning and one highAlert threshold. For example, if there are 3 caches in a cluster, where:

cache1 val = 95

cache2 val = 100

cache3 val = 70

and the CapacityLimitAllCaches highWarning is **80** and highAlert is **90**, one high alert is executed.

OcCapacityLimitCache

Executes when the average CPU usage for the cluster / storage class exceeds the specified thresholds and durations. This alert has a per cluster and a per (cluster) storage class override. This alert appears in the Other Category when executed.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **95** (percent), Alarm is **95** (percent) and Duration is **60** (seconds).

OcClusterNodesRcvdFailureRateHigh

Executes when the average network/packet received failure rate for the cluster/storage class exceeds the specified thresholds and durations. The metrics are averaged across all nodes of a storage class in a cluster.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

This alert has a per cluster and a per (cluster) storage class override. Note that this alert appears in the Network Category when executed.

By default the alert is disabled with the following default settings: Warning is **95** (percent), Alarm is **95** (percent) and Duration is **60** (seconds).

OcClusterNodesSentFailureRateHigh

Executes when the average network/packet sent failure rate for the cluster / storage class exceeds the specified thresholds and durations. The metrics are averaged across all nodes of a storage class in a cluster.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

This alert has a per cluster and a per (cluster) storage class override. Note that this alert appears in the Memory Category when executed.

By default the alert is disabled with the following default settings: Warning is **95** (percent), Alarm is **95** (percent) and Duration is **60** (seconds).

OcDepartedNode

For each node in the cluster, an alert is executed if the time a node is absent from the cluster exceeds the specified thresholds. When the departed node rejoins the cluster, the alert is cleared.

OcDepartedNodesPercentage

This scalar alert executes a single warning and a single alert if the percentage of nodes departed from the cluster exceeds the specified thresholds within the specified time periods. The percentage is measured against the total number of nodes in the cluster, including both running and departed nodes.

The time period is set in the **rtview.properties** file using the `$NODES_DEPARTED_TIME_WINDOW` substitution. The time period can also be overridden using the command line interface. For example, the following sets a time window of 300 seconds:

```
-sub: $NODES_DEPARTED_TIME_WINDOW: 300
```

The time period default setting is **600** (10 minutes), the warning default setting is **90** (percent) and the alarm default setting is **95** (percent).

By default the alert is disabled.

OcEndangeredAllCaches

This alert is executed if the StatusHA for the cache service is `NODE_SAFE` (high warning) or `ENDANGERED` (high alert).

OcEndangeredCache

For each node in the cluster, an alert is executed if the StatusHA value is `ENDANGERED`. By default the alert is disabled.

OcExtendConnectionByteBacklogHigh

This limits alert executes a single warning and a single alert if the `OutgoingByteBacklog` for a Proxy Extend Connection exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is **1000** (bytes), Alert is **5000** (bytes).

OcHATargetFailed

This alert executes when the distributed service target status (`HATarget`) is not met. The `HATarget` value is determined using the `PartitionAssignment` MBean in Coherence Versions 12 and above. In prior Coherence versions, the default value of `MACHINE_SAFE` is used. The default value can be overridden by setting the substitution variable `$ocmDefaultHATarget` to the desired value.

OcHighGCDutyCycleNode

This scalar alert executes a single warning and a single alert if a node exceeds the specified duty cycle threshold (the percent of time spent in Garbage Collection).

By default the alert is enabled with the following default settings: Warning is **10** (percent), Alarm is **20** (percent) and Duration is **10** seconds.

OcHighPendingRequestNode

A single alert is executed if the `RequestPendingCount` amount exceeds the specified threshold. This alert allows for setting the warning level, alarm level and duration.

By default the alert is disabled.

OcHighTaskBacklogNode

A single warning and a single alert are executed if the number of backlogged tasks exceeds the specified user threshold. This alert allows for setting the warning level, alarm level and duration.

The default setting executes a warning if the number of backlogged tasks exceeds **10**, and executes an alert if the number of backlogged tasks exceeds **20**.

By default the alert is disabled.

OcHighThreadAbandonedNode

A single alert is executed if the Coherence Thread Abandoned Count amount exceeds the specified threshold. This alert allows for setting the warning level, alarm level and duration.

The default setting executes a warning and an alert if the Thread Abandoned Count amount exceeds **0**. The default duration setting is **60**.

By default the alert is enabled.

OcJmxProcessingTime

This alert is executed if the sum of time for JMX queries and all data processing functions exceeds the specified threshold for the **jmxsampleperiod** property. By default the alert is disabled with the following default settings: Warning is **80** (percent), Alarm is **90** (percent) and Duration is **0** (seconds).

NOTE: The OcJmxProcessingTime alert does not support overrides. For that alert the Override Count is displayed as **-1**.

OcLongGCDurationNode

A single warning and a single alert are executed if any of the last garbage collection times exceed the specified duration.

The default setting executes a warning if the duration exceeds 1 second, and executes an alert if the duration exceeds 2 seconds.

It is possible for GC times to exceed the specified duration and NOT execute an alert. This is possible if it occurs between the alert duration time and an alert condition time.

For example, if your alert duration is 60 seconds, and there is also an alert condition set at 27 seconds into that 60 seconds, the following scenarios could occur (where XX:XX:XX is Hours:Minutes:Seconds):

Scenario 1:

12:00:00 GC amount is below the specified threshold. No alert executed.

12:00:27 GC amount exceeds the specified threshold. Alert ignored for now.

12:01:00 GC amount is below the specified threshold. No alert executed.

Scenario 2:

12:00:00 GC amount is below the specified threshold. No alert executed.

12:00:27 GC amount exceeds the specified threshold. Alert ignored for now.

12:01:00 GC amount remains above the specified threshold. Alert executed.

By default the alert is enabled.

OcLowClientNodeCount

This alert executes if the total number of nodes being monitored, including storage enabled nodes, client nodes, and management (JMX) nodes, exceeds the specified threshold. When the count returns to above the threshold (departed nodes rejoin the cluster), the alert is cleared.

By default the alert is disabled.

OcLowStorageNodeCount

This alert executes if the total number of storage nodes in the cluster exceeds the specified threshold. When the count returns to above 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 the threshold (departed nodes rejoin the cluster), the alert is cleared.

By default the alert is disabled.

OcMemoryUsedPercentageAfterGC

This alert is executed if the percent of memory used on a node after garbage collection exceeds the specified threshold. By default the alert is disabled with the following default settings: Warning is **70** (percent), Alarm is **80** (percent) and Duration is **30** (seconds).

OcNodeSafeCache

For each node in the cluster, an alert is executed if the StatusHA value is **NODE-SAFE**. By default the alert is disabled.

OcNoJmxConnection

This alert is executed if a JMX connection remains disconnected after a specified duration of time. The default duration of time is **60** seconds. By default, this alert is enabled.

OcObjectCountDeltaUpCache

This tabular alert executes a single warning and a single alert for each cache in the cluster if the cache object count delta increases and exceeds the specified threshold. In addition to setting the warning and alarm levels, this alert also allows for setting the duration for each cache.

When this alert is selected in the Active Alert Table, the Per Cache Alert Setting box is displayed (rather than the scalar alert box).

By default the alert is disabled.

OcObjectCountDeltaDownCache

This tabular alert executes a single warning and a single alert for each cache in the cluster where the cache object count delta decreases and exceeds the specified threshold. In addition to setting the warning and alarm levels, this alert also allows for setting the duration for each cache.

When this alert is selected in the Active Alert Table, the Per Cache Alert Setting box is displayed (rather than the scalar alert box).

By default the alert is disabled.

OcProxyNodeByteBacklogHigh

This limits alert executes a single warning and a single alert if the OutgoingByteBacklog for a Proxy Node exceeds the specified threshold. This is often indicates overloaded capacity on an individual proxy node. By default the alert is disabled with the following default settings: Warning is **100** (bytes), Alarm is **50** (bytes).

OcSendQueueSize

For each node in the cluster, an alert is executed if the Send Queue for that node exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is **100** (seconds), Alarm is **200** (seconds) and Duration is **60** (seconds).

OcStoreFailure

This alert is executed if the number of StoreFailures exceeds the specified threshold. By default the alert is disabled with the following default settings: Warning is **1** (second), Alarm is **10** (seconds) and Duration is **30** (seconds).

OcStoreReadMillisHigh

This alert is executed if the current average read per millisecond (total current StoreReadMillis/total current StoreReads) exceeds the specified threshold for a sampling period and the specified cache(s).

Oracle Database

The following alerts are available for Oracle Database. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
OraDatabaseConnectionLoss Executes a single alert if the SQL database connection state is false. Index Type: Per Database Metric: Connected	NaN	NaN	30	FALSE
OraDatabaseQueryError Executes a single alert if the if the last query state is false (an error). Index Type: Per Database Metric: Last Query Status	NaN	NaN	30	FALSE
OraDatabaseResponseTimeHigh Executes a single warning and a single alarm if the time (in milliseconds) to execute a SQL query exceeds the specified threshold. Index Type: Per Database Metric: ResponseTimeMilliSec	200	220	30	FALSE
OraDatabaseSpaceUsedHigh Executes a single warning and a single alarm if the percent utilization of the space allocated to the database exceeds the specified threshold. Index Type: Per Database Metric: PercentUsedSpace	80	90	30	FALSE
OraDatabaseTablespaceUsedHigh Executes a single warning and a single alarm if the percent utilization of the database used by the tablespace exceeds the specified threshold. Index Type: Per Table Space Metric: USED_PERCENT	80	90	30	FALSE
OraInstanceAvgQueryTimeHigh Executes a single warning and a single alarm if the average time (in milliseconds) to perform a query exceeds the specified threshold. Index Type: Per Instance Metric: AVGQUERYTIME	300	400	30	FALSE
OraInstanceCommitRateHigh Executes a single warning and a single alarm if the number of commits per second exceeds the specified threshold. Index Type: Per Instance Metric: RateCOMMITTS	250	300	30	FALSE

OraInstanceNumCurrentLoginsHigh Executes a single warning and a single alarm if the number of database clients exceeds the specified threshold. Index Type: Per Instance Metric: CURRENT_LOGINS	12	15	30	FALSE
OraInstanceDataDictHitRatioLow Executes a single warning and a single alarm if the data dictionary hit ratio goes below the specified threshold. Index Type: Per Instance Metric: DD_HIT_RATIO	95	90	30	FALSE
OraInstanceDiskReadRateHigh Executes a single warning and a single alarm if the number of physical disk reads per second exceeds the specified threshold. Index Type: Per Instance Metric: RatePHYSICAL_READS	250	300	30	FALSE
OraInstanceDiskWriteRateHigh Executes a single warning and a single alarm if the number of physical disk writes per second exceeds the specified threshold. Index Type: Per Instance Metric: RatePHYSICAL_WRITES	250	300	30	FALSE
OraInstanceLatchHitRatioLow Executes a single warning and a single alarm if the latch hit ratio goes below the specified threshold. Index Type: Per Instance Metric: LatchHitPerCent	95	90	30	FALSE
OraInstanceMaxQueryTimeHigh Executes a single warning and a single alarm if the query time (in milliseconds) exceeds the specified threshold. Index Type: Per Instance Metric: MAXQUERYTIME	10000	15000	30	FALSE
OraInstanceNumActiveSessionsHigh Executes a single warning and a single alarm if the number of active sessions for the instance exceeds the specified threshold. Index Type: Per Instance Metric: ACTIVE_SESSIONS	12	15	30	FALSE
OraInstanceNumCurrentLoginsHigh Executes a single warning and a single alarm if the number of current logins for the instance exceeds the specified threshold. Index Type: Per Instance Metric: CURRENT_LOGINS	12	15	30	FALSE

OraInstanceRollbackRateHigh Executes a single warning and a single alarm if the number of rollbacks per second exceeds the specified threshold. Index Type: Per Instance Metric: RateROLLBACKS	5	10	30	FALSE
OraInstanceSqlHitRatioLow Executes a single warning and a single alarm if the SQL hit ratio goes below the specified threshold. Index Type: Per Instance Metric: SQL_HIT_RATIO	95	90	30	FALSE
OraInstanceState Executes a single warning and a single alarm if the database is not in an ACTIVE or OPEN state for queries. Index Type: Per Instance Metric: AlertStatus	NaN	NaN	30	FALSE

Oracle WebLogic

The following alerts are available for Oracle WebLogic. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
WlsAppNewSessionsRateLow The rate per second of newly opened sessions is below the specified threshold. Index Type(s): PerApplication	10	1	30	FALSE
WlsAppOpenSessionsHigh The maximum total number of open sessions for that application has been reached. Index Type(s): PerApplication	7	10	30	FALSE
WlsClusterServersPercentNotRunningHigh The percentage of cluster not running is high. Index Type(s): PerCluster	33	50	30	FALSE
WlsHoggingThreadsHigh The maximum number of hogging threads for that server has been reached. Index Type(s): PerServer	15	20	30	FALSE
WlsJDBCConnectionsWaitingHigh Triggered when the number of threads waiting for a JDBC connection exceeds the threshold. Index Type(s): PerConnection, PerLocation, PerModule, PerName	1	10	0	FALSE

WlsJmsBytesCurrentHigh The current number of bytes stored on this JMS server has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsJmsBytesPendingHigh The current number of bytes pending (unacknowledged or uncommitted) stored on this JMS Server has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsJmsConnectionsCurrentHigh The current number of connections to this JMS WebLogic Server has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsJmsDestinationBytesCurrentHigh The current number of bytes stored in the destination, not including the pending bytes, has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsJmsDestinationBytesPendingHigh The number of pending bytes stored in the destination has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsJmsDestinationConsumersCurrentLow The number of pending bytes stored in the destination has reached its minimum. Index Type(s): PerServer	15	5	30	FALSE
WlsJmsDestinationMessagesCurrentHigh The current number of messages in the destination has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsJmsDestinationMessagesPendingHigh The number of pending messages in the destination has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE

WlsJmsDestinationsCurrentLow The current number of destinations on this JMS Server has reached its minimum. Index Type(s): PerServer Note: To enable this alert, you must uncomment the following options under the Collect all other metrics section in the sample.properties file: <pre># Collect all other metrics; all or none collector.sl.rtvview.cache.config=wls_workmg r_cache.rtv collector.sl.rtvview.cache.config=wls_auxila ry_cache.rtv collector.sl.rtvview.cache.config=wls_jmsser ver_cache.rtv collector.sl.rtvview.cache.config=wls_jmsbri dge_cache.rtv collector.sl.rtvview.cache.config=wls_jmspst ore_cache.rtv</pre>	85	95	30	FALSE
WlsJmsMessagesPendingHigh The current number of messages pending (unacknowledged or uncommitted) stored on this JMS Server has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsJmsServerHealthNotOK The health state of this JMS Server is not OK. Index Type(s): PerServer	NaN	NaN	30	FALSE
WlsLockedUserCurrentHigh The maximum number of current locked users for that server has been reached. Index Type(s): PerServer Note: To enable this alert, you must uncomment the following options under the Collect all other metrics section in the sample.properties file: <pre># Collect all other metrics; all or none collector.sl.rtvview.cache.config=wls_workmg r_cache.rtv collector.sl.rtvview.cache.config=wls_auxila ry_cache.rtv collector.sl.rtvview.cache.config=wls_jmsser ver_cache.rtv collector.sl.rtvview.cache.config=wls_jmsbri dge_cache.rtv collector.sl.rtvview.cache.config=wls_jmspst ore_cache.rtv</pre>	85	95	30	FALSE
WlsOpenSocketsHigh The maximum number of open sockets for that server has been reached. Index Type(s): PerServer	85	95	30	FALSE
WlsPendingRequestCurrentHigh The maximum number of current requests for that server has been reached. Index Type(s): PerServer	85	95	30	FALSE

WlsQueueLengthHigh The number of pending requests in the priority queue has reached its maximum. This is the total of internal system requests and user requests. Index Type(s): PerServer	85	95	30	FALSE
WlsServerCpuHigh The server CPU has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsServerHealthNotOK The server health is not OK. Index Type(s): PerServer	NaN	NaN	30	FALSE
WlsServerHostCpuHigh The CPU percentage of the host server has reached its maximum. Index Type(s): PerServer	85	95	30	FALSE
WlsServerMemoryUsageHigh The maximum used memory established for the server has been reached. Index Type(s): PerServer	85	95	30	FALSE
WlsServerNewSessionsLow The number of new sessions created is below the threshold. Index Type(s): PerServer	15	5	30	FALSE
WlsServerOpenSessionsHigh The maximum number of open sessions for that server has been reached. Index Type(s): PerServer	85	95	30	FALSE
WlsServerPendingUserRequestsHigh The maximum number of pending user requests has been reached. Index Type(s): PerServer	85	95	30	FALSE
WlsServerReloadsHigh The maximum number of reloads for that server has been reached. Index Type(s): PerServer	85	95	30	FALSE
WlsServerStaleData The server has stale data. Index Type(s): PerServer	NaN	NaN	30	FALSE
WlsServerStateNotRunning The state of the server is different from "Running." Index Type(s): PerServer	NaN	NaN	30	FALSE

WlsThreadsTotalHigh The total number of threads for that server has been reached. Index Type(s): PerServer	50	95	30	FALSE
WlsTransactionRolledBackTotalHigh The total number of transactions rolled back has been reached. Index Type(s): PerServer	85	95	30	FALSE

RTView Host Agent

The following alerts are available for RTView Host Agent. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
HostCpuLoadAvg1High Executes a single warning alert and a single alarm alert if the average CPU load per minute exceeds the specified threshold. Index Type(s): PerHost Metric: loadAvg1	50	75	30	FALSE
HostCpuLoadAvg5High Executes a single warning alert and a single alarm alert if the average CPU load per 5 minutes exceeds the specified threshold. Index Type(s): PerHost Metric: loadAvg5	50	75	30	FALSE
HostCpuLoadAvg15High Executes a single warning alert and a single alarm alert if the average CPU load per 15 minutes exceeds the specified threshold. Index Type(s): PerHost Metric: loadAvg15	50	75	30	FALSE
HostCpuPercentHigh Executes a single warning alert and a single alarm alert if the percent CPU load exceeds the specified threshold. Index Type(s): PerHost Metric: hostCpuPercent	50	75	30	FALSE

HostMemoryUsedHigh Executes a single warning alert and a single alarm alert if the percent of physical memory used exceeds the specified threshold. Index Type(s): PerHost Metric: MemUsedPerCent	75	90	5	FALSE
HostNetworkRxRateHigh Executes a single warning alert and a single alarm alert if the inbound network data rate, in kilobytes per second, exceeds the specified threshold. Index Type(s): PerHost Metric: RateRxKBytes	50	75	30	FALSE
HostNetworkTxRateHigh Executes a single warning alert and a single alarm alert if the outbound network transmission rate, in kilobytes per second, exceeds the specified threshold. Index Type(s): PerHost Metric: RateTxKBytes	50	75	30	FALSE
HostProcessCountLow Executes a single warning alert and a single alarm alert if the process count exceeds the specified threshold. Index Type(s): PerHost Metric: Count	80	90	30	FALSE
HostStaleData Executes a single alarm alert and sets the Expired flag to true if data is not received from the given host within the specified expiration time interval. Index Type(s): PerHost Metric: Expired	NaN	NaN	30	FALSE
HostStorageUsedHigh Executes a single warning alert and a single alarm alert if the percent of space used on the storage medium exceeds the specified threshold. Index Type(s): PerStorage Metric: percentused	80	90	5	FALSE
HostSwapUsedHigh Executes a single warning alert and a single alarm alert if the percent of used swap space exceeds the specified threshold. Index Type(s): PerHost Metric: swapUsedPerCent	75	90	30	FALSE
HostVirtualMemoryUsedHigh Executes a single warning alert and a single alarm alert if the percent of used virtual memory exceeds the specified threshold. Index Type(s): PerHost Metric: VMemUsedPerCent	75	90	30	FALSE

RTView Manager and RTView Rules

If RTView Manager and RTView Rules are installed on your system you might see the following alert types for RTView Servers (Data Servers, Display Servers and Historian Servers):

RTView Server Manager Alert Types

JvmCpuPercentHigh	Executes a single warning alert and a single alarm alert if the percent of JVM CPU used exceeds the specified threshold. Index Type: Per JVM Metric:: CpuPercent
JvmGcDutyCycleHigh	Executes a single warning alert and a single alarm alert if the garbage collector duty cycle exceeds the specified threshold. Index Type: Per GC Source Metric: DutyCycle
JvmMemoryUsedAfterGCHigh	Executes a single warning alert and a single alarm alert if the percent of memory used after garbage collection exceeds the specified threshold. Index Type: Per GC Source Metric: PctMemoryUsedAfterGC
JvmMemoryUsedHigh	Executes a single warning alert and a single alarm alert if the percent of memory used exceeds the specified threshold. Index Type(s): Per JVM Metric: MemoryUsedPercent
JvmNotConnected	Executes a single alert if the JVM is disconnected, indicating that it might have crashed. Index Type(s): Per JVM Metric: Connected
JvmStaleData	Executes a single alert if the data update wait time exceeds the specified duration threshold. Index Type(s): Per JVM Metric: Expired
JvmThreadCountHigh	Executes a single warning alert and a single alarm alert if the number of threads exceeds the specified threshold. Index Type(s): Per JVM Metric: ThreadCount
TomcatAccessRateHigh	Executes a single warning alert and a single alarm alert if the number of accesses per second exceeds the specified threshold. Index Type(s): Per Server Metric: RateaccessCount
TomcatActiveSessionsHigh	Executes a single warning alert and a single alarm alert if the number of active sessions exceeds the specified threshold. Index Type(s): Per Server Metric: activeSessions

TomcatAppAccessRateHigh

Executes a single warning alert and a single alarm alert if the number of accesses per second exceeds the specified threshold.

Index Type(s): Per Application

Metric: RateaccessCount

TomcatAppActiveSessionsHigh

Executes a single warning alert and a single alarm alert if the number of active sessions exceeds the specified threshold.

Index Type(s): Per Application

Metric: activeSessions

RTView Rules Alert Types**RtvEmServiceAlert**

This discrete alert is generated when a Service has one or more alerts on any associated CIs.

RtvEmServiceAlertImpactHigh

This limits alert is generated when a Service has an Alert Impact value that exceeds the specified threshold on any associated CI.

Solace

The following alerts are available with both the solution package and standalone versions for Solace. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
SolBridgeInboundByteRateHigh The number of inbound bytes per second across the bridge has reached its maximum. Index Type: PerBridge	8000000	10000000	30	FALSE
SolBridgeInboundMsgRateHigh The number of inbound messages per second across the bridge as a whole has reached its maximum. Index Type: PerBridge	40000	50000	30	FALSE
SolBridgeOutboundByteRateHigh The number of outbound bytes per second across the bridge has reached its maximum. Index Type: PerBridge	8000000	10000000	30	FALSE
SolBridgeOutboundMsgRateHigh The number of outbound messages per second across the bridge has reached its maximum. Index Type: PerBridge	40000	50000	30	FALSE
SolClientInboundByteRateHigh The number of outbound bytes per second for the client has reached its maximum. Index Type: PerClient	8000000	10000000	30	FALSE

SolClientInboundMsgRateHigh The number of outbound messages per second for the client as a whole has reached its maximum. Index Type: PerClient	40000	50000	30	FALSE
SolClientOutboundByteRateHigh The number of outbound bytes per second for the client has reached its maximum. Index Type: PerClient	8000000	10000000	30	FALSE
SolClientOutboundMsgRateHigh The number of outbound messages per second for the client as a whole has reached its maximum. Index Type: PerClient	40000	50000	30	FALSE
SolClientSlowSubscriber One or more clients are consuming messages too slowly; endpoints may drop messages! Index Type: PerClient	1	NaN	30	FALSE
SolCspfNeighborDown State is not "OK" for one or more CSPF neighbors. Index Type: PerNeighbor	1	NaN	30	FALSE
SolEndpointPendingMsgsHigh The number of pending messages on a queue has reached its maximum. Index Type: PerEndpoint	8000	10000	30	FALSE
SolEndpointSpoolUsageHigh The endpoint is consuming too much message router memory for storing spooled messages. (Threshold units are megabytes.) Index Type: PerEndpoint	40	50	30	FALSE
SolEventModuleBrokerAlert This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the SYSTEM scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of SYSTEM scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.	FALSE			
SolEventModuleClientAlert This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the CLIENT scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of CLIENT scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.	FALSE			
SolEventModuleVpnAlert This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the VPN scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of VPN scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.	FALSE			

SolGuaranteedMsgingHbaLinkDown For Guaranteed Messaging only, the Operational State for each HBA Fibre-Channel should be Online (e.g., not Linkdown). Index Type: PerHbaLink	0	NaN	30	FALSE
SolGuaranteedMsgingMatePortDown For Guaranteed Messaging only, the Mate Link Ports for ADB should have status OK. Index Type: PerADB	0	NaN	30	FALSE
SolGuaranteedMsgingNoMsgSpoolAdActive For Guaranteed Messaging only with Redundancy, at least one message router in an HA pair should show "AD-Active." Index Type: PerPair	0	NaN	30	FALSE
SolMsgRouterActiveDiskUtilHigh The utilization of the active disk partition for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterByteEgressUtilHigh The egress rate (bytes/sec) utilization (current egress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterByteIngressUtilHigh The ingress rate (bytes/sec) utilization (current ingress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterConnectionUtilHigh The connection utilization for the message router (current number of connections divided by max allowed) is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterCpuTemperatureHigh CPU temperature margin is above threshold. Index Type: PerApplianceSensor	-30	-15	30	FALSE
SolMsgRouterCspfNeighborDown Link-detect = no for CSPF neighbor. Index Type: PerAppliance	1	NaN	30	FALSE
SolMsgRouterDelvrdUnAckMsgUtilHigh The delivered unacked messages as a percentage of all messages delivered for the application is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterFailoverDetected The backup message router in a HA pair has assumed control. Index Type: PerAppliance	1	NaN	30	FALSE

SolMsgRouterFanSensorCheckFailed The speed measured for one or more fans is below threshold. Index Type: PerApplianceSensor	5000	2657	30	FALSE
SolMsgRouterInboundByteRateHigh The number of inbound bytes per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterInboundMsgRateHigh The number of inbound messages per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterIngressFlowUtilHigh The ingress flow utilization (current flows divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterInterfaceDown Link-detect = no for one or more enabled network interfaces. Index Type: PerSolInterface	NaN	NaN	30	FALSE
SolMsgRouterMsgCountUtilHigh The message count utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterMsgEgressUtilHigh The message egress rate utilization (current message egress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterMsgIngressUtilHigh The message ingress rate utilization (current message ingress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterNABUsageHigh Network Acceleration Blade memory usage is excessive. Index Type: PerNAB	60	80	30	FALSE
SolMsgRouterNotConnected The message router is not ready for collecting performance monitoring data. Index Type: PerAppliance	NaN	NaN	30	FALSE
SolMsgRouterOutboundByteRateHigh The number of outbound bytes per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterOutboundMsgRateHigh The number of outbound messages per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE

SolMsgRouterPendingMsgsHigh The total number of pending messages for this message router has reached its maximum. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterPowerSupplyFailed A power supply has failed. Index Type: PerAppliance	0	NaN	30	FALSE
SolMsgRouterSpoolUtilization The amount of spool space used for messages is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterStandbyDiskUtilHigh The utilization of the standby disk partition for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSubscriptionUtilHigh The subscription utilization (current number of subscriptions divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSwapUsedHigh The amount of swap space used by the message router operating system is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSyslogAlert This alert executes when a Solace Syslog Warning or Critical message is received. To get Syslog event alerts (in RTView Enterprise Monitor or the standalone Monitor), go to the Alert Administration display and enable the SolMsgRouterSyslog alert.	-	-	-	-
SolMsgRouterTemperatureSensorCheckFailed A chassis temperature measurement is above threshold. Index Type: PerAppliance	40	45	30	FALSE
SolMsgRouterTranSessionCntUtilHigh The transacted session count utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterTranSessionResUtilHigh The transacted session resource utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterVoltageSensorCheckFailed A power supply voltage is high or low. Index Type: PerApplianceSesor	NaN	NaN	30	FALSE

SolSparseMessageSpoolFile This is a Limits Alert that issues a Warning alert and is enabled by default. Important: Do not modify thresholds for this alert as they were set up by Solace development team) This alert is defined to determine when there is a Sparse Message Spool File Condition. When disk space usage is several multiples of persistent store usage, then there is likely a large number of message spool files residing on the disk where each file contains few messages. This is referred to as a sparse message spool file condition, and requires urgent attention to mitigate and avoid the disk reaching capacity. For further information, refer to Solace documentation.				TRUE
SolVpnConnectionCountHigh The number of connections to the server has reached its maximum. Index Type: PerVPN	60	80	30	FALSE
SolVpnInboundByteRateHigh The number of inbound bytes per second for the vpn has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
SolVpnInboundDiscardRateHigh The number of discarded inbound messages per second for the server is excessive. Index Type: PerVPN	1	5	30	FALSE
SolVpnInboundMsgRateHigh The number of inbound messages per second for the vpn as a whole has reached its maximum. Index Type: PerVPN	40000	50000	30	FALSE
SolVpnOutboundByteRateHigh The number of outbound bytes per second for the VPN has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
SolVpnOutboundDiscardRateHigh The number of discarded outbound messages per second for the server is excessive. Index Type: PerVPN	1	5	30	FALSE
SolVpnOutboundMsgRateHigh The number of outbound messages per second for the server as a whole has reached its maximum. Index Type: PerVPN	40000	50000	30	FALSE
SolVpnPendingMsgsHigh The total number of pending messages for this destination has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
SolVpnSubscriptionCountHigh The number of endpoints in this VPN has reached its maximum. Index Type: PerVPN	8000	10000	30	FALSE

TIBCO ActiveMatrix BusinessWorks

The following alerts are available with both the solution package and standalone versions for TIBCO® ActiveMatrix BusinessWorks™. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
BW6AppErrorState BW6 application status is not Running or Stopped (status is Impaired, AppError or StartFailed) Index Type: PerApp Metric: State	NaN	NaN	30	FALSE
Bw6AppExpired BW6 application expired due to application inactivity. Index Type: PerApp Metric: Stopped	NaN	NaN	30	FALSE
Bw6AppNodeCpuUsedHigh BW6 AppNode CPU usage exceeded limit. CPU Usage is the CPU time in use by all processes expressed as a percentage of the total CPU time available. Index Type: PerAppNode Metric: CPU Usage%	50	80	30	FALSE
Bw6AppNodeMemUsedHigh BW6 AppNode memory usage exceeded limit. Memory usage is the percentage of total JVM memory currently consumed by this appnode. Index Type: PerAppNode Metric: Memory Usage%	50	80	30	FALSE
Bw6AppNodeStopped BW6 AppNode stopped purposefully (for example, an administrator stopped the AppNode process). Index Type: PerAppNode Metric: State	NaN	NaN	10	FALSE
Bw6AppNodeUnreachable BW6 AppNode stopped abnormally (for example, the AppNode process crashed). Index Type: PerAppNode Metric: State	NaN	NaN	10	FALSE
Bw6AppProcessCreatedRateHigh BW6 Process created rate for application exceeded limit. Index Type: PerApp Metric: App Created Rate	50	80	30	FALSE
Bw6AppProcessElapsedTimeHigh BW6 Process delta elapsed time rate of increase for application exceeded limit. Index Type: PerApp Metric: App Elapsed Rate	200	400	30	FALSE

Bw6AppProcessExecutionTimeHigh BW6 Process delta execution time rate of increase for application exceeded limit. Index Type: PerApp Metric: App Execution Rate	200	400	30	FALSE
Bw6AppProcessFailedRateHigh BW6 Process failed rate for application exceeded limit. Index Type: PerApp Metric: App Failed Rate	50	80	30	FALSE
Bw6AppStopped BW6 application stopped. Index Type: PerApp Metric: Stopped	NaN	NaN	30	FALSE
Bw6ProcessActivityErrorRateHigh BW6 Process error rate exceeded limit. Index Type: PerProcess Metric: Process Failed Rate	50	80	30	FALSE
Bw6ProcessCreatedRateHigh BW6 Process error rate exceeded limit. Index Type: PerProcess Metric: Process Failed Rate	50	80	30	FALSE
Bw6ProcessElapsedTimeHigh BW6 Process delta elapsed time rate of increase exceeded limit. Index Type: PerProcess Metric: Delta Exec Rate	200	400	30	FALSE
Bw6ProcessExecutionTimeHigh BW6 Process delta execution time rate of increase exceeded limit. Index Type: PerProcess Metric: Delta Time Rate	200	400	30	FALSE
Bw6ProcessFailedRateHigh BW6 Process suspended rate exceeded limit. Index Type: PerProcess Metric: Suspended Rate	50	80	30	FALSE
Bw6ProcessHung The delta elapsed time is greater than zero but the delta execution time is zero. Index Type: PerProcess Metric: Hung/Not Hung	NaN	NaN	10	FALSE
Bw6ProcessSuspendRateHigh BW6 Process failed rate exceeded limit. Index Type: PerProcess Metric: Failed Rate	50	80	30	FALSE

BwActivityErrorRateHigh BW5 Activity error rate exceeded limit. The rate is calculated by taking the delta of total error returns in this update period and dividing by the length of the period. Index Type: PerActivity Metric: RateErrorCount	50	80	30	FALSE
BwActivityExecutionTimeHigh BW5 Activity execution time rate of increase exceeded limit. The rate is calculated by taking the delta of total execution time in this update period and dividing by the length of the period. Index Type: PerActivity Metric: RateExecutionTime	200	400	30	FALSE
BwEngineCpuUsedHigh BW Engine CPU usage (% of total) exceeded limit. CPU Usage is the CPU time used by the BW engine expressed as a percentage of the total CPU time available. Index Type: PerEngine Metric: CPU Usage%	50	80	30	FALSE
BwEngineMemUsedHigh BW Engine memory usage (% of total) exceeded limit. Memory usage is the percentage of total JVM memory currently consumed by this engine. Index Type: PerEngine Metric: PercentUsed	50	80	30	FALSE
BwEngineStopped BW Engine has stopped running. Index Type: PerEngine Metric: Stopped	NaN	NaN	30	FALSE
BwEngineUnreachable BW engine stopped abnormally. Index Type: PerEngine Metric: State	NaN	NaN	30	FALSE
BwProcessAbortRateHigh BW Process aborted rate exceeded limit. The rate is calculated by taking the delta of total aborts in this update period and dividing by the length of the period. Index Type: PerProcess Metric: RateAborted	50	80	30	FALSE
BwProcessAvgElapsedTimeHigh BW Process Average Elapsed Time exceeded limit. Value is calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. Index Type: PerProcess Metric: Process Avg Elapsed Time	100	200	30	FALSE
BwProcessAvgExecutionTimeHigh BW Process average execution time exceeded limit. Index Type: PerProcess Metric: AverageExecution	0	0	0	FALSE

BwProcessCreatedRateHigh BW Process creation rate exceeded limit. The rate is calculated by taking the number of process instances created in the interval and dividing by the length of the interval in seconds. Index Type: PerProcess Metric: Processes Created/sec	100	200	30	FALSE
BwProcessCreatedRateLow BW Process creation rate per second went below limit. Index Type: PerProcess Metric: App Created Rate	0	0	0	FALSE
BwProcessElapsedTimeHigh BW Process elapsed time rate of increase exceeded limit. The rate is calculated by taking the delta of total elapsed time in this update period and dividing by the length of the period. Index Type: PerProcess Metric: RateTotalElapsed	50	80	30	FALSE
BwProcessExecutionTimeHigh BW Process execution time rate of increase exceeded limit. The rate is calculated by taking the delta of total execution time in this update period and dividing by the length of the period. Index Type: PerProcess Metric: RateTotalExecution	50	80	30	FALSE
BwProcessHung The delta elapsed time is greater than zero but the delta execution time is zero. Index Type: PerProcess Metric: Hung/Not Hung	NaN	NaN	10	FALSE
BwProcessSuspendRateHigh BW Process suspended rate exceeded limit. The rate is calculated by taking the delta of total suspends in this update period and dividing by the length of the period. Index Type: PerProcess Metric: RateSuspended	50	80	30	FALSE
BwProcessTotalCpuPercentHigh BW Process CPU percent utilization exceeded limit. This is the percent CPU used by all process instances executing over the interval. Index Type: PerProcess Metric: Process Total CPU Percent	50	75	30	FALSE
BwServerCpuUsedHigh BW Server CPU usage (% of total) exceeded limit. CPU Usage is the CPU time in use by all processes expressed as a percentage of the total CPU time available. Index Type: PerServer Metric: CPU Usage%	60	85	30	FALSE

BwServerFreeMemLow BW Server free memory available (in megabytes) is below limit. Free memory means available physical (RAM) memory. Index Type: PerServer Metric: Memory Free Mbytes	15	5	30	FALSE
BwServerInactive BW Server has become inactive. The period of time specified by the substitution variable \$bwserverExpirationTime has passed since data was last received from the server. Index Type: PerServer Metric: Expired	NaN	NaN	30	FALSE
BwServerMemUsedHigh BW Server memory usage (% of total) exceeded limit. Memory usage is the virtual memory in use expressed as a percentage of the available virtual memory. The meaning of available virtual memory is system-dependent: on Windows it refers to pagefile space; on Unix systems it refers to swap space. Index Type: PerServer Metric: Virtual Memory Used%	50	80	30	FALSE
HawkAlert Display Hawk alerts throughout the Monitor. To enable Hawk Alerts to be included in alert counts and displayed throughout the Monitor, scroll down to HawkAlert in the Active Alert Table and select the Alert Enabled checkbox. It is possible to filter unwanted alerts from the cache data so that those alerts are not included throughout the Monitor. To filter unwanted alerts out of the Hawk cache data, enter the following into the sample.properties file (located in the project directory you created). NOTE: Unwanted alerts are filtered out according to the AlertText. sl.rtvew.sub=\$hawkAlertTextFilterOut:AlertText For example, to filter out all Hawk Alerts in which the AlertText contains Source you would enter the following: sl.rtvew.sub=\$hawkAlertTextFilterOut:Source The default time to remove cleared Hawk Alerts from the table is 3600 seconds. To adjust this setting, edit the following in sample.properties : sl.rtvew.sub=\$hawkAlertTextFilterOut:3600 Index Type: PerServer Metric: Hawk	NaN	NaN	-1	TRUE
JvmCpuPercentHigh The percentage of CPU that has been reached by the JVM is above the limit. Index Type: PerJVM Metric: CpuPercent	50	75	30	FALSE
JvmGcDutyCycleHigh The duty cycle is out the upper limit. Index Type: PerGC Metric: DutyCycle	50	75	30	FALSE

JvmMemoryUsedHigh The memory used out the upper limit Index Type: PerJVM Metric: MemoryUsedPercent	50	75	30	FALSE
JvmNotConnected The JVM in not connected. Index Type: PerJVM Metric: Connected	NaN	NaN	30	FALSE
JvmStaleData Cut in reception from that JVM. Index Type: PerJVM Metric: Expired	NaN	NaN	30	FALSE

TIBCO ActiveSpaces

The following alerts are available for TIBCO ActiveSpaces. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
TasMemberCpuHigh The CPU usage is above the defined thresholds. Index Type(s): PerMember	80	95	30	FALSE
TasMemberEntriesHigh The number of objects inserted into the space is above the defined thresholds. Index Type(s): PerMember	8000	10000	30	FALSE
TasMemberEvictsRateHigh The rate at which 'evicts' are occurring is above the defined thresholds. Index Type(s): PerMember	80	100	30	FALSE
TasMemberExpireRateHigh The rate at which 'expires' are occurring is above the defined thresholds. Index Type(s): PerMember	80	100	30	FALSE
TasMemberGetRateHigh The rate at which 'gets' are occurring is above the defined thresholds. Index Type(s): PerMember	80	100	30	FALSE
TasMemberJvmMemoryUsedHigh The percent JVM memory used is above the defined thresholds. Index Type(s): PerMember	80	95	30	FALSE

TasMemberMemoryUsedHigh The percent memory used is above the defined thresholds. Index Type(s): PerMember	80	95	30	FALSE
TasMemberPutRateHigh The rate at which 'puts' are occurring is above the defined thresholds. Index Type(s): PerMember	80	100	30	FALSE
TasMemberSeederCapacity The percentage utilization (number of entries/capacity)*100 of the seeder is high for the given space. "Capacity per seeder" must be set in the space definition for this alarm to be effective. Index Type(s): PerMemberandSpace	80	90	30	FALSE
TasMemberTakeRateHigh The rate at which 'takes' are occurring is above the defined thresholds. Index Type(s): PerMember	80	100	30	FALSE
TasMetaspaceEntriesHigh The number of objects inserted into the metaspace is above the defined thresholds. Index Type(s): PerMetaspace	8000	100000	30	FALSE
TasMetaspaceEvictsRateHigh The rate at which 'evicts' are occurring is above the defined thresholds. Index Type(s): PerMetaspace	80	100	30	FALSE
TasMetaspaceExpireRateHigh The rate at which 'expires' are occurring is above the defined thresholds. Index Type(s): PerMetaspace	80	100	30	FALSE
TasMetaspaceGetRateHigh The rate at which 'gets' are occurring is above the defined thresholds. Index Type(s): PerMetaspace	80	100	30	FALSE
TasMetaspacePutRateHigh The rate at which 'puts' are occurring is above the defined thresholds. Index Type(s): PerMetaspace	80	100	30	FALSE
TasMetaspaceTakeRateHigh The rate at which 'takes' are occurring is above the defined thresholds. Index Type(s): PerMetaspace	80	100	30	FALSE
TasQueryDurationHigh The query duration (in seconds) is above the defined threshold (in seconds). Index Type(s): PerSpace	4	5	30	FALSE

TasSpaceEntriesHigh The number of objects inserted into the space is above the defined thresholds. Index Type(s): PerSpace	8000	100000	30	FALSE
TasSpaceEvictsRateHigh The rate at which 'evicts' are occurring is above the defined thresholds. Index Type(s): PerSpace	80	100	30	FALSE
TasSpaceExpireRateHigh The rate at which 'expires' are occurring is above the defined thresholds. Index Type(s): PerSpace	80	100	30	FALSE
TasSpaceGetRateHigh The rate at which 'gets' are occurring is above the defined thresholds. Index Type(s): PerSpace	80	100	30	FALSE
TasSpacePutRateHigh The rate at which 'puts' are occurring is above the defined thresholds. Index Type(s): PerSpace	80	100	30	FALSE
TasSpaceSeederCountLow Not enough seeders are available. Index Type(s): PerSpace	NaN	NaN	30	FALSE
TasSpaceState The state of the space is "not ready". Index Type(s): PerSpace	NaN	NaN	30	FALSE
TasSpaceTakeRateHigh The rate at which 'takes' are occurring is above the defined thresholds. Index Type(s): PerSpace	80	100	30	FALSE

TIBCO Adapters

The following alerts are available for TIBCO Adapters. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
TadAdapterDeltaErrorsHigh The number of errors incurred by the adapter in last measurement interval is above the defined threshold. Index Type(s): PerAdapter	1600	2000	30	FALSE

TadAdapterExpired The data from this adapter has not been updated since the last measurement interval. The data shown from this adapter is currently stale. Index Type(s): PerAdapter	NaN	NaN	0	FALSE
TadAdapterMsgsRcvdRateHigh The number of messages received by this adapter since the last measurement interval is above the defined threshold. Index Type(s): PerAdapter	1600	2000	60	FALSE
TadAdapterMsgsSentRateHigh The number of messages sent by the adapter in the last measurement interval is above the defined threshold. Index Type(s): PerAdapter	1	2	60	FALSE

TIBCO BusinessEvents

The following alerts are available with both the solution package and standalone versions for TIBCO® BusinessEvents®.

TbeBackingStoreEraseRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which entries are erased from the backing store exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeBackingStoreLoadRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which entries are loaded from the backing store exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeBackingStoreStoreRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which entries are written to the backing store exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeClusterMalformed	<p>This alert executes for any cluster where the member count is not equal to the expected cluster size. The expected cluster size is a count of the number of nodes that have the same cluster name, as discovered by reading the cluster MBean for each node in the connection property file. The MemberCount attribute is also read from the same cluster MBean, and is the number of nodes in the (sub)cluster which the current node has joined.</p> <p>The condition where these counts differ can occur if there are missing connections in the property file (for example, some nodes are unmonitored). It can also occur if, due to network or other anomalies, some nodes do not join the "main" cluster, but instead form a "sub-cluster" of one or more nodes. This condition is commonly referred to as "split-brain".</p>
TbeDestinationStatusRecvdEventsRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which events are received from the channel exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .

TbeNodeConceptsGetRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which concepts are received from the cache exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeNodeConceptsPutRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which concepts are written to the cache exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeNodeConceptsRemoveRateHigh	This alert executes a single warning alert and a single alarm alert if the rate which concepts are removed from the cache exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeNodeConnectionLoss	This discrete alert executes when the JMX Connection to the TIBCO BusinessEvents agent is lost (the TCP connection flag for an engine is false).
TbeNodeEventsGetRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which events are received from the cache exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeNodeEventsPutRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which events are written to the cache exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeNodeEventsRemoveRateHigh	This alert executes a single warning alert and a single alarm alert if the rate which events are removed from the cache exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .
TbeObjectTableExtIdSize	This alert executes a single warning alert and a single alarm alert if the number of external object IDs exceeds the specified threshold. The warning default threshold is 9000 and the alarm default threshold is 10000 .
TbeObjectTableSize	This alert executes a single warning alert and a single alarm alert if the number of objects maintained by the cache exceeds the specified threshold. The warning default threshold is 9000 and the alarm default threshold is 10000 .
TbeRuleFiringRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which rules are executing exceeds the specified threshold. The warning default threshold is 80 and the alarm default threshold is 95 .

TIBCO Enterprise Message Service

The following alerts are available with both the solution package and standalone versions for TIBCO® Enterprise Message Service™. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
EmsConsumerStalled Indicates consumers are stalled or are no longer consuming messages (have not received a message within a defined threshold). The server must be running for a minimum time (5 minutes by default) before this alert is triggered. Thresholds are in seconds. Note: This alert does not allow overrides. Index Type(s): PerConsumer: ID/ PerServerConsumer: URL; ID Metric: elapsedSinceLasAckInSec	85	95	30	FALSE
EmsConsumerStuck Indicates a consumer is stuck because there are existing messages that can be consumed (currentMsSentCount > 0), but none of the messages have been consumed within the defined warning and alert thresholds (elapsedSinceLasAckInSec > threshold). Alert and warning thresholds are in seconds. Index Type(s): PerConsumer: ID/ PerServerConsumer: URL; ID Metric: currentMsgSentCount, elapsedSinceLasAckInSec	85	95	30	FALSE
EmsQueueConsumerIdleTimeHigh The idle time of the queue consumer has reached its maximum. This alert is triggered when there is no change in the number of incoming messages for a queue within a specified period of time (in seconds). Index Type(s): PerQueue; PerServerQueue Metric: ConsumerIdleTime	60	80	30	FALSE
EmsQueueInboundDeltaHigh The number of new incoming messages for the EMS Queue has reached its maximum. Index Type(s): PerQueue; PerServerQueue Metric: DeltainboundTotalMessages	60	80	30	FALSE
EmsQueueMsgLatencyHigh The time, in seconds, needed to process all pending messages based on the current outbound message rate exceeded its threshold. This alert does not take into account queues with outbound message rate equals to zero. Index Type(s): PerServerQueue: URL; name Metric: messageLatency	60	80	30	FALSE

EmsQueueProviderIdleTimeHigh The queue idle time exceeded the specified threshold. A queue is idle when the number of inbound messages remains unchanged. Index Type(s): PerServerQueue: URL; name Metric: ProviderIdleTime	60	80	30	FALSE
EmsQueuesConsumerCountHigh The number of consumers of a queue exceeded the specified high threshold. Index Type(s): PerServerQueue: URL; name/ PerQueue: name Metric: consumerCount	60	80	30	FALSE
EmsQueuesConsumerCountLow The number of consumers of a queue is below the specified threshold. Index Type(s): PerServerQueue: URL; name/ PerQueue: name Metric: consumerCount	15	5	30	FALSE
EmsQueuesInMsgRateHigh The rate of inbound messages on the queue exceeded the specified threshold. Index Type(s): PerServerQueue: URL; name/ PerQueue: name Metric: inboundMessageRate	60	80	30	FALSE
EmsQueuesOutMsgRateHigh The number of outbound messages on the queue exceeded the specified threshold. Index Type(s): PerServerQueue: URL; name Metric: outboundMessageRate	60	80	30	FALSE
EmsQueuesPendingMsgsHigh The number of pending messages on the queue exceeded the specified threshold. Index Type(s): PerServerQueue: name; PerServerQueue: URL; name Metric: pendingMessageCount	60	80	30	FALSE
EmsQueuesProducerCountHigh The number of producers to a queue exceeded the specified high threshold. Index Type(s): PerQueue: name/ PerServerQueue: URL; name Metric: producerCount	60	80	30	TRUE
EmsQueuesProducerCountLow The number of producers to a queue is below the specified threshold. Index Type(s): PerQueue: name/ PerServerQueue: URL; name Metric: producerCount	15	5	30	TRUE

EmsServerAsyncDBSizeHigh The size of the Async database, in bytes, for the EMS Server reached its maximum. Index Type(s): PerServer: URL Metric: asyncDBSize	50	100	30	FALSE
EmsServerInboundDeltaHigh The number of new incoming messages for the EMS Server has reached its maximum Index Type(s): PerServer Metric: DeltainboundMessageCount	60	80	30	FALSE
EmsServerSyncDBSizeHigh The size of the Sync database, in bytes, for the EMS Server reached its maximum. Index Type(s): PerServer: URL Metric: syncDBSize	50	100	30	FALSE
EmsServerConnectionCountHigh Alert is triggered when the number of connections to the server reaches the specified threshold. Index Type(s): PerServer: URL Metric: connectionCount	60	80	30	FALSE
EmsServerInMsgRateHigh The number of inbound messages on the server exceeded the specified threshold. Index Type(s): PerServer: URL Metric: inboundMessageRate	2	80	30	FALSE
EmsServerMemUsedHigh The percent memory used on the server exceeded the specified threshold. Index Type(s): PerServer: URL Metric: messageMemoryPct	60	80	30	FALSE
EmsServerNotStarted The server state is empty. The server is not started. Index Type(s): PerServer: URL Metric: NotStarted	NaN	NaN	30	FALSE
EmsServerOutMsgRateHigh The number of outbound messages on the server exceeded the specified threshold. Index Type(s): PerServer: URL Metric: outboundMessageRate	60	80	30	FALSE
EmsServerPendingMsgsHigh The number of pending messages in the server queue exceeded the specified threshold. Index Type(s): PerServer: URL Metric: pendingMessageCount	60	80	30	FALSE

EmsServerPendingMsgSizeHigh The size, in KB, of the pending messages stored on this EMS Server reached its maximum. Index Type(s): PerServer: URL Metric: pendingMessageCount	60	80	30	FALSE
EmsServerRouteState One or more routes on the server are not active. Index Type(s): PerServer: URL Metric: Alert State	NaN	NaN	30	FALSE
EmsServerStaleData The server stopped receiving data. Index Type(s): PerServer: URL Metric: Expired	NaN	NaN	30	FALSE
EmsTopicConsumerIdleTimeHigh The idle time of the topic consumer has reached its maximum. This alert is triggered when there is no change in the number of incoming messages for a topic within a specified period of time (in seconds). Index Type(s): PerTopic; PerServerTopic Metric: ConsumerIdleTime	60	80	30	FALSE
EmsTopicInboundDeltaHigh The number of new incoming messages for the EMS Topic has reached its maximum. Index Type(s): PerTopic; PerServerTopic Metric: DeltainboundTotalMessages	60	80	30	FALSE
EmsTopicMsgLatencyHigh The time, in seconds, needed to process all pending messages based on the current outbound message rate exceeded its threshold. This alert does not take into account topics with outbound messages rates equal to zero. Index Type(s): PerServerTopic Metric: messageLatency	60	80	30	FALSE
EmsTopicProviderIdleTimeHigh The topic idle time exceeded the specified threshold. A topic is idle when the number of inbound messages remains unchanged. Index Type(s): PerServerTopic: URL; name Metric: ProviderIdleTime	60	80	30	FALSE
EmsTopicsConsumerCountHigh The number of consumers for the topic exceeded the specified threshold. Index Type(s): PerServerTopic: URL; name Metric: consumerCount	60	80	30	FALSE
EmsTopicsConsumerCountLow The number of consumers for the topic is below the specified threshold. Index Type(s): PerServerTopic Metric: consumerCount	60	80	30	FALSE

EmsTopicsInMsgRateHigh The number of inbound messages for the topic exceeded the specified threshold. Index Type(s): PerServerTopic Metric: inboundMessageRate	60	80	30	FALSE
EmsTopicsOutMsgRateHigh The rate of outbound messages for the topic exceeded the specified threshold. Index Type(s): PerServerTopic Metric: outboundMessageRate	60	80	30	TRUE
EmsTopicsPendingMsgsHigh The number of pending messages on the queue for the topic exceeded the specified threshold. Index Type(s): PerTopic Metric: pendingMessageCount	50	75	30	FALSE
EmsTopicsProducerCountHigh The number of active producers for this topic exceeded the specified high threshold. Index Type(s): PerTopic/PerServerTopic Metric: producerCount	60	80	30	TRUE
EmsTopicsProducerCountLow The number of producers for the topic is below the specified threshold. Index Type(s): PerTopic/PerServerTopic Metric: producerCount	60	80	30	TRUE
EmsTopicsSubscriberCountHigh The number of subscribers for the topic exceeded the specified threshold. Index Type(s): PerServerTopic Metric: subscriberCount	50	75	30	FALSE
JvmCpuPercentHigh The percent JVM CPU usage exceeded the specified threshold. Index Type(s): PerJVM Metric: CpuPercent	30	40	30	FALSE
JvmGcDutyCycleHigh The JVM Garbage Collection contains an item that exceeded the specified duty cycle threshold (the percent of time spent in Garbage Collection). Index Type(s): PerGC Metric: TimeUsedPercent	50	75	30	FALSE
JvmMemoryUsedHigh The percent JVM memory used exceeded the specified threshold. Index Type(s): PerJVM Metric: MemoryUsedPercent	50	75	30	FALSE

JvmNotConnected The JVM is not connected. Index Type(s): PerJVM Metric: Connected	NaN	NaN	30	FALSE
JvmStaleData The JVM stopped receiving data. Index Type(s): PerJVM Metric: Expired	NaN	NaN	30	FALSE

TIBCO FTL

The following alerts are available for TIBCO FTL. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
TftlClientCpuTime Executes a single warning and a single alarm if the CPU response time to the client exceeds the specified threshold. Index Type: PerClient Metric: Delta_CPU_TIME	160000	200000	30	FALSE
TftlClientCpuUsage Executes a single warning and a single alarm if the CPU usage by the client exceeds the specified threshold. Index Type: PerClient Metric:	160000	200000	30	FALSE
TftlClientExpired Executes a single alert if the response time to the client exceeds the specified threshold. Index Type: PerClient Metric:	NaN	NaN	30	FALSE
TftlClientMemory Executes a single warning and a single alarm if the memory usage by the client exceeds the specified threshold. Index Type: PerClient Metric: PROCESS_RSS_KB	160000	200000	30	FALSE
TftlClientMsgsRcvdRate Executes a single warning and a single alarm if the number of messages received by the client per second exceeds the specified threshold. Index Type: PerClient Metric: RateMESSAGES_RECEIVED	160000	200000	30	FALSE

TftlClientMsgsSentRate Executes a single warning and a single alarm if the number of messages sent by the client per second exceeds the specified threshold. Index Type: PerClient Metric: RateMESSAGES_SENT	160000	200000	30	FALSE
TftlClientNotRunning Executes a single if the client status is not "RUNNING". Index Type: PerClient Metric: Delayed Writes	NaN	NaN	30	FALSE
TftlClientVirtualMemory Executes a single warning and a single alarm if the virtual memory usage by the client exceeds the specified threshold. Index Type: PerClient Metric:	160000	200000	30	FALSE
TftlServerClientCount Executes a single warning and a single alarm if the number of clients on the FTL server exceeds the specified threshold. Index Type: PerServer Metric:	160	200	30	FALSE
TftlServerCpuTime Executes a single warning and a single alarm if the FTL server CPU response time exceeds the specified threshold. Index Type: PerServer Metric:	160	200	30	FALSE
TftlServerCpuUsage Executes a single warning and a single alarm if the FTL server CPU usage exceeds the specified threshold. Index Type: PerServer Metric:	60	80	30	FALSE
TftlServerExpired Executes a single warning and a single alarm if the FTL server response time exceeds the specified threshold. Index Type: PerServer Metric:	NaN	NaN	30	FALSE
TftlServerInboxSendFaults Executes a single warning and a single alarm if the number of times the FTL server fails to queue messages to the appropriate inbox exceeds the specified threshold. Index Type: PerServer Metric: SEND_TO_INBOX_FAILURES	160	200	30	FALSE

TftlServerMemory Executes a single warning and a single alarm if the FTL server memory usage exceeds the specified threshold. Index Type: PerServer Metric: PROCESS_RSS_KB	160	200	30	FALSE
TftlServerOnBackup Executes a single alert if the primary FTL server is down and now running on the backup FTL server. Index Type: PerServer Metric:	NaN	NaN	30	FALSE
TftlServerSatelliteCount Executes a single alert if the number of satellite servers is lower than expected. Note: Set threshold to one less than number of deployed satellites. Index Type: Response Time Metric: Table_locks_waited	NaN	5	30	FALSE
TftlServerVirtualMemory Executes a single warning and a single alarm if the FTL server virtual memory usage exceeds the specified threshold. Index Type: Response Time Metric: Table_locks_waited	160	200	30	FALSE

UX

The following are the Monitor alerts you can enable to be aware of any web application that is unresponsive, performing slowly, generating errors or returning invalid information. By default, Monitor alerts are disabled.

Monitor alerts execute when the UX Robot performs its routine runs on URLs. The **uxmon.properties** file defines which URLs the UX Robot checks and reports on. There are two types of Monitor alerts, UX-ROBOT alerts and UX-URL alerts.

- UX-ROBOT alerts apply to multiple URLs.
- UX-URL alerts apply to a single URL.

UXRobotError	During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the number of URL errors exceed the specified threshold. The warning default threshold is 1 and the alarm default threshold is 10 . For example, the URL error message "no such URL" indicates an issue at the Web Server that serves the URL. Using the default settings, a warning alert executes if the UX Robot encounters 1 or more URL errors and an alarm alert executes if the UX Robot encounters 10 or more URL errors.
UXRobotResponseSlow	During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the total response time for all specified URLs exceeds the specified threshold. The warning default threshold is 1000 milliseconds and the alarm default threshold is 2000 milliseconds.

UXRobotSearchSentinel	<p>During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the number of false URL responses (responses without the specified searchString in the URL line) exceeds the specified threshold. The warning default threshold is 1 and the alarm default threshold is 10.</p> <p>For example, using the default settings, a warning alert executes if the UX Robot encounters 1 or more false responses from URLs and an alarm alert executes if the UX Robot encounters 10 or more false responses from URLs.</p>
UXRobotTimeout	<p>During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the number of URL timeouts exceeds the specified maxTimeoutMS threshold. The warning default threshold is 1 and the alarm default threshold is 15.</p> <p>For example, the URL error message "no such URL" indicates an issue at the Web Server that serves the URL. Using the default settings, a warning alert executes if the UX Robot encounters 1 or more URL errors and an alarm alert executes if the UX Robot encounters 15 or more URL errors.</p>
UXURLError	<p>During a UX Robot run, this UX-URL alert executes a single alert if the UX Robot receives an error message from a URL. The default setting is TRUE.</p> <p>For example, the URL error message "no such URL" indicates an issue at the Web Server that serves the URL.</p>
UXURLResponseSlow	<p>During a UX Robot run, this UX-URL alert executes a single warning alert and a single alarm alert if the response time for a URL exceeds the specified threshold. The warning default threshold is 1000 milliseconds and the alarm default threshold is 2000 milliseconds.</p>
UXURLSearchSentinel	<p>During a UX Robot run, this UX-URL alert executes an alert if the UX Robot receives a false URL response (a response without the specified searchString in the URL line). The default setting is FALSE.</p>
UXURLTimeout	<p>During a UX Robot run, this UX-URL alert executes an alert if the URL response time exceeds the specified maxTimeoutMS threshold. UX Robot receives a false URL response (a response without the specified searchString). The default setting is TRUE.</p>

VMware vCenter

The following alerts are available for VMware vCenter. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
VmwareHostCpuUtilizationHigh The Host's CPU utilization is above the defined threshold. Index Type(s): PerVmHost	50	75	2	TRUE
VmwareHostDiskBytesReadHigh The disk read rate (kBytes/second) is above the defined thresholds. Index Type(s): PerVmHost	1024	2048	2	TRUE

VmHostDiskBytesWrittenHigh The disk write rate (kBytes/second) is above the defined thresholds. Index Type(s): PerVmHost	1024	2048	2	TRUE
VmwHostInBytesHigh The inbound byte rate (KB/second) is above the defined thresholds. Index Type(s): PerVmHost	1024	2048	2	TRUE
VmwHostInPktDropLossHigh The percentage of inbound packets dropped is above the defined threshold. Index Type(s): PerVmHost	1	3	2	TRUE
VmwHostInPktErrorLossHigh The percentage of inbound packets discarded for any error is above the defined threshold. Index Type(s): PerVmHost	1	3	2	TRUE
VmwHostMemoryUsageHigh The percentage memory utilization (used/configured) is above the defined threshold. Index Type(s): PerVmHost	70	80	2	TRUE
VmwHostOutBytesHigh The outbound byte rate (KB/second) is above the defined threshold. Index Type(s): PerVmHost	1024	2048	2	TRUE
VmwHostOutPktDropLossHigh The percentage of outbound packets dropped is above the defined thresholds. Index Type(s): PerVmHost	1	3	2	TRUE
VmwHostOutPktErrorLossHigh The percentage of inbound packets discarded for any error is above the defined threshold. Index Type(s): PerVmHost	1	3	2	TRUE
VmwHostStatusBad The overall status is not "green." Index Type(s): PerVmHost	NaN	NaN	2	TRUE
VmwHostSwapUsedHigh The amount of swap space used by a host is above the defined thresholds. Index Type(s): PerVmHost	10240	40960	2	TRUE
VmwVmCpuUtilizationHigh The virtual machine CPU utilization is above the defined thresholds. Index Type(s): PerVm	50	75	2	TRUE
VmwVmDiskBytesReadHigh The disk read rate (KB/second) is above the defined thresholds. Index Type(s): PerVm	1024	2048	2	TRUE

VmwVmDiskBytesWrittenHigh The disk write rate (KB/second) is above the defined thresholds. Index Type(s): PerVm	1024	2048	2	TRUE
VmwVmDiskUsageHigh The amount of disk space used by the virtual machine is above the defined threshold. Index Type(s): PerVm	85	95	30	TRUE
VmwVmInBytesHigh The inbound byte rate (KB/second) is above the defined threshold. Index Type(s): PerVm	1024	2048	2	TRUE
VmwVmInPktDropLossHigh The percentage of inbound packet loss due to dropped packets is above the defined threshold. Index Type(s): PerVm	1	3	30	TRUE
VmwVmMemoryUsageHigh The percentage of memory utilization (active/configured) is above the defined thresholds. Index Type(s): PerVm	70	80	2	TRUE
VmwVmOutBytesHigh The outbound byte rate is above the defined threshold. Index Type(s): PerVm	1024	2048	2	TRUE
VmwVmOutPktDropLossHigh The percentage of outbound packet loss due to dropped packets on the virtual machine is above the defined threshold. Index Type(s): PerVm	1	3	2	TRUE
VmwVmStatusBad The overall status for this virtual machine is not "green." Index Type(s): PerVm	NaN	NaN	2	TRUE
VmwVmSwapUsedHigh The amount of host memory swapped out for the virtual machine by the host's virtual machine kernel is above the defined threshold. This metric is not related to any swapping the may occur in the guest operating system. Index Type(s): PerVm	3072	4096	2	TRUE

APPENDIX B Limitations

This section includes:

- [“iPad Safari Limitations”](#)
- [“TIBCO ActiveMatrix BusinessWorks”](#)

iPad Safari Limitations

- In the iPad settings for Safari, **JavaScript** must be **ON** and **Block Pop-ups** must be **OFF**. As of this writing, the Thin Client has been tested only on iOS 4.3.5 in Safari.
- The iPad does not support Adobe Flash, so the Fx graph objects (obj_fxtrend, obj_fxpie, obj_fxbar) are unavailable. The Thin Client automatically replaces the Fx graph objects with the equivalent non-Fx object (obj_trendgraph02, obj_pie, obj_bargraph). Note that the replacement objects behave the same as the Fx objects in most cases but not in all. In particular, obj_trendgraph02 does not support the sliding cursor object nor the **legendPosition** property. Custom Fx objects are not supported on the iPad.
- The Thin Client implements scrollbars for table objects and graph objects. However, unlike the scrollbars used on desktop browsers, the scrollbars used on the iPad do not have arrow buttons at each end. This can make it difficult to scroll precisely (for example, row by row) on objects with a large scrolling range.
- At full size, users may find it difficult to touch the intended display object without accidentally touching nearby objects and performing an unwanted drill-down, sort, scroll, and so forth. This is particularly true of table objects that support drill-down and also scrolling, and also in panel layouts that contain the tree navigation control. In those cases, the user may want to zoom the iPad screen before interacting with the Thin Client.
- If the iPad sleeps or auto-locks while a Thin Client display is open in Safari, or if the Safari application is minimized by clicking on the iPad's home button, the display is not updated until the iPad is awakened and Safari is reopened. In some cases it may be necessary to refresh the page from Safari's navigation bar.

Because the iPad uses a touch interface there are differences in the Thin Client appearance and behavior in iOS Safari as compared to the conventional desktop browsers that use a cursor (mouse) interface, such as Firefox and Internet Explorer. These are described below.

- **Popup browser windows:** An RTView object's drill-down target can be configured to open a display in a new window. In a desktop browser, when the RTView object is clicked the drill-down display is opened in a popup browser window. But in iOS Safari 4.3.5, only one page is visible at a time, so when the RTView object is touched a new page containing the drill-down display opens and fills the screen. The Safari navigation bar can be used to toggle between the currently open pages or close them.
- **Mouseover text:** When mouseover text and drill-down are both enabled on an RTView object (for example, a bar graph), in iOS Safari the first touch on an element in the object (for example, a bar) displays the mouseover text for that element and the second touch on the same element performs the drill-down.
- **Resize Mode and Layout:** By default, the Display Server runs with **resizeMode** set to **crop**. In **crop** mode, if a display is larger than the panel that contains it only a portion of the display is visible. In a desktop browser, scrollbars become available to allow the user to scroll to view the entire display. In iOS Safari, scrollbars do not appear but the display can be scrolled by dragging two fingers inside the display. (Dragging one finger scrolls the entire page, not the display).

If the Display Server is run with **resizeMode** set to **scale** or **layout**, the display is resized to fit into the panel that contains it. If a desktop browser is resized after a display is opened, the display is resized accordingly. On the iPad, the Safari browser can only be resized by reorienting the iPad itself, between portrait mode and landscape mode.

The panel layout feature is supported in the Thin Client. However, unlike a desktop browser which resizes to match the layout size, the size of Safari is fixed. So if the Display Server is run with **resizeMode** set to **crop** or **scale** mode, there may be unused space at the edges of the display(s) or, in **crop** mode, the panels and displays may be cropped.

This means that **layout** mode should be used for best results on the iPad. For layout mode to be most effective, displays should use the **anchor** and **dock** object properties. Please see RTView documentation for more information.

- **Scrolling:** The Thin Client implements scrollbars for table objects and graph objects. The scrollbars are activated by dragging with one finger.

If an RTView display is viewed in **crop** mode and is too large to be displayed entirely in Safari, scrollbars do not appear (as they would in a desktop browser) but the display can be scrolled by dragging with two fingers inside the display.

Scrollbars do not ever appear in a text area control. If the text area contains more text than is visible, use the two finger drag in the text area to scroll the text.

Regardless of the size of a listbox control, it can only display a single item (typically, the selected item). When the listbox is touched, the list of items appear in a popup list. In other words, on iOS Safari the listbox control and the combobox control behave identically.

- Context menu: The Thin Client context menu is opened by a right mouse button click in a desktop browser. It is opened in iOS Safari by touching any location on a display and holding that touch for 2 seconds. The menu appears in the top left corner of the display, regardless of where the display is touched. The items **Export Table to Excel**, **Drill Down**, and **Execute Command** are not included on the context menu in Safari. All other items are available. The **Export Table to HTML** item is enabled if a table object is touched (unless the table object's drillDownTarget is configured to open another display). After an **Export to PDF/HTML** is performed, the exported content opens on another page in Safari. From there, the content can either be opened by another application (for example, the iBooks application opens PDF) and emailed, or it can be copied and pasted into an email.

TIBCO ActiveMatrix BusinessWorks

Servers

AIX

- Status will be **LIMITED**.
- CPU Usage, Free Memory and Virtual Memory Usage will not be available.

Business Works 5.7.1 Engine Status

The BW Engine microagent has a method **GetExecInfo** that includes a field called **Status**, which may have the following values:

- ACTIVE
- SUSPENDED
- STANDBY
- STOPPING
- STOPPED

In Business Works 5.7.1 (but not earlier or later versions) this method fails to return any data and, in some cases when the starts, it may not know an engine's exact status. For example, if an engine is deployed but not active it could be SUSPENDED or STOPPED, or else it could be ACTIVE or STOPPING. In these cases the sets the status to UNKNOWN. An UNKNOWN status will be resolved once the engine is stopped and restarted; henceforth the status will display as STOPPED or ACTIVE.

BWSE Components

- JVM memory metrics are available for BWSE components running in AMX 3.x environments only.
- The BW Version column in the All Engines Table display is blank for BWSE components.
- The Deployment column in the All Engines Table display is UNKNOWN for BWSE components. This is because the AMX environment controls in which node or nodes a BWSE component is running, therefore the concept of "deployment" in traditional BusinessWorks does not apply.
- BWSE components only appear in the All Engines Table display when they are running in a node.

APPENDIX C Third Party Notice Requirements

This section includes:

- "RTView EM"
- "RTView Core®"

RTView EM

** The UX Monitor Solution Package uses HTMLUnit 2.14 and is licensed under the Apache License Version 2.0

** The UX Monitor Solution Package uses Log4J 1.2 and is licensed under the Apache License Version 2.0

** The RTView Host Agent Solution Package uses Hyperic SIGAR and is licensed under the Apache License Version 2.0

** The Solution Package for Apache Kafka uses Jackson 2.8.4 and is licensed under the Apache License Version 2.0

** The Solution Package for Apache Kafka uses Netty 3.10.5.Final and is licensed under the Apache License Version 2.0

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below)

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

=====

** The UX Monitor Solution Package uses Netbeans App Framework 1.0.3 and is licensed under GNU Lesser General Public License Version 3

=====

** BrowserLauncher2 1.3

This product uses BrowserLauncher 1.3 and is distributed pursuant to the terms of the Lesser General Public License. The source code for BrowserLauncher2 1.3 can be obtained from:
<http://browserlaunch2.sourceforge.net/>

GNU LESSER GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

0. Additional Definitions.

As used herein, "this License" refers to version 3 of the GNU Lesser General Public License, and the "GNU GPL" refers to version 3 of the GNU General Public License.

"The Library" refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An "Application" is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A “Combined Work” is a work produced by combining or linking an Application with the Library. The particular version of the Library with which the Combined Work was made is also called the “Linked Version”.

The “Minimal Corresponding Source” for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The “Corresponding Application Code” for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

1. Exception to Section 3 of the GNU GPL.

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

2. Conveying Modified Versions.

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

- a) under this License, provided that you make a good faith effort to ensure that, in the event an Application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or
- b) under the GNU GPL, with none of the additional permissions of this License applicable to that copy.

3. Object Code Incorporating Material from Library Header Files.

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

- a) Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the object code with a copy of the GNU GPL and this license document.

4. Combined Works.

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

- a) Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the Combined Work with a copy of the GNU GPL and this license document.
- c) For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.
- d) Do one of the following:

- 0) Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.
- 1) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (a) uses at run time a copy of the Library already present on the user's computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.
- e) Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 4d0, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option 4d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

5. Combined Libraries.

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.
- b) Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

6. Revised Versions of the GNU Lesser General Public License.

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy's public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.

RTView Core®

** Apache Tomcat is delivered for convenience only as a separate application and is licensed under the Apache License Version 2.0

** Apache HttpClient is embedded in the RTView Core libraries and is licensed under the Apache License Version 2.0

** JEval 0.9.4 is licensed under the Apache License Version 2.0

** Jetty 9.2.19 is licensed under the Apache License Version 2.0

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean anyform resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below)

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean anywork of authorship, including the original version of the Work and anymodifications or additions to that Work or Derivative Worksthereof, that is intentionally submitted to Licensor for inclusion inthe Work by the copyright owner or by an individual or Legal Entityauthorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives,including but not limited to communication on electronic mailinglists, source code control systems, and issue tracking systems that aremanaged by, or on behalf of, the Licensor for the purpose of discussingand improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyrightowner as "Not a Contribution."

"Contributor" shall meanLicensor and any individual or Legal Entity on behalf of whom a Contribution hasbeen received by Licensor and subsequently incorporated within theWork.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor herebygrants to You a perpetual, worldwide, non-exclusive, no-charge,royalty-free, irrevocable copyright license to reproduce, prepareDerivative Works of, publicly display, publicly perform,sublicense, and distribute the Work and such Derivative Works in Sourceor Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor herebygrants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combinationof their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement,then any patent licenses granted to You under this License forthat Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce anddistribute copies of the Work or Derivative Works thereof in anymedium, with or without modifications, and in Source or Objectform, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions.

Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

=====

** TreeMap Algorithms v1.0 is used without modifications and licensed by MPL Version 1.1. The source for TreeMap Algorithms can be obtained from <http://www.cs.umd.edu/hcil/treemap/>

** iTextAsian 1.0 is licensed by MPL Version 1.1 and the source can be obtained from: <http://itextpdf.com/download.php>

MOZILLA PUBLIC LICENSE

Version 1.1

1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. Source Code License.

2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4. Intellectual Property Matters

(a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

EXHIBIT A -Mozilla Public License.

``The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Portions created by _____ are Copyright (C) _____
_____. All Rights Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[_____] License"), in which case the provisions of [_____] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [_____] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [_____] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [_____] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

=====

****MD Datejs**

Copyright © 2006-2010 Coolite Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

****jQuery**

Copyright © 2009 John Resig

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

**** JCalendar 1.3.2**

This product uses JCalendar 1.3.2. JCalendar is distributed pursuant to the terms of the Lesser General Public License. The source code for the JCalendar may be obtained from <http://www.toedter.com/en/jcalendar/index.html>

=====

**** BrowserLauncher2 1.3**

This product uses BrowserLauncher 1.3 and is distributed pursuant to the terms of the Lesser General Public License. The source code for BrowserLauncher2 1.3 can be obtained from: <http://browserlaunch2.sourceforge.net/>

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the library's name and an idea of what it does.

Copyright (C) year name of author

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public

License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

signature of Ty Coon, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!