CA APM

CA Technologies Prerelease Documentation.
Authorized Use Only.

Workstation User Guide

Version 9.1.0



This documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the “Documentation”) is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION “AS IS” WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with “Restricted Rights.” Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2011 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Internal Use Only. Reference use of documentation. Limited and use only.

CA Technologies Product References

Wily Technology, the Wily Technology Logo, Introscope, and All Systems Green are registered trademarks of CA.

The following are trademarks of CA:

- Blame
- Blame Game
- CA Service Quality Management
- CA SiteMinder®
- CA Unicenter® Network and Systems Management
- CA Unicenter® Service Desk
- CA Unicenter® Service Metric Analysis
- CA Wily Customer Experience Manager
- CA Wily Customer Experience Manager for Siebel Applications
- CA Wily Customer Experience Manager Integration Pack for CA CMDB
- CA Wily Customer Experience Manager Integration Pack for Unicenter® Network and Systems Management
- CA Wily Introscope®
- CA Wily Introscope® ChangeDetector
- CA Wily Introscope® Environment Performance Agent
- CA Wily Introscope® ErrorDetector
- CA Wily Introscope® for Microsoft® .NET
- CA Wily Introscope® for SAP ABAP
- CA Wily Introscope® for SAP® NetWeaver
- CA Wily Introscope® LeakHunter
- CA Wily Introscope® PowerPacks
- CA Wily Introscope® SNMP Adapter
- CA Wily Introscope® SQL Agent
- CA Wily Introscope® Transaction Tracer
- CA Wily Manager for CA SiteMinder®
- CA Wily Portal Manager
- CA Wily SOA Performance Manager

CA Technologies Proprietary Documentation.
Authorized Use Only.

- CA Wily Transaction Generator
- Get Wily
- SmartStor

Java is a trademark of Sun Microsystems in the U.S. and other countries. All other names are the property of their respective holders.

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

If you would like to provide feedback about CA Technologies product documentation, complete our short customer survey, which is available on the CA Support website at http://ca.com/docs.

CA Technologies Prerelease Documentation. Authorized Use Only.

Contents

Chapter 1: Introduction 13

About Application Performance Management 13
- Understanding the APM Architecture 14
- Understanding Directory and File Naming Conventions 14

Introscope and the Workstation 15
- How the Workstation fits in an Introscope installation 16
- The Workstation, Java Web Start, and WebView 16

What's New in the Workstation 17
- Application Triage Map Enhancements 18

Administering the Workstation 19
- Start the Workstation 19
- End Your Workstation Session 25
- Execute Workstation Functions from the Command Line 25
- Configuring HTTP tunneling for the Workstation 27
- Configuring the Workstation to use SSL 27

Introscope Workstation Elements 29
- About the Workstation Console 29
- About the Workstation Investigator 30
- About the Management Module Editor 33
- About the Dashboard Editor 33
- About Data Viewers 33
- About Alerts and Alert Indicators 37

Managing Users 39
- User Permissions 39
- User Preferences 39
- Managing Language Settings 40

About this Guide and Other CA APM Guides 41
- Product Documentation 41

Chapter 2: Using the Workstation Console 45

Navigating Among Dashboards in the Console 45
- Dashboard Drop-down List 45
- Navigate Using Hyperlinks 46
- Creating Dashboard Favorites 46
- Launching Investigator from Console 47
- Launching Console from Investigator 47

CA Technologies Prerelease Documentation
Authorized Use Only.

Find More Information from Dashboards 47
Filtering by agent with the Console Lens 48
Preconfigured CA APM Dashboards 50
Enhanced Dashboards 52
Where the Sample Dashboards are Installed 53
Overall Status Indicators on Dashboards 53
The Sample Intro to Introscope Dashboard 53
The Sample Overview Dashboard 54
The Sample Problem Analysis Dashboard 55
Navigation Details 56
Performance Dashboards 58
Capacity Dashboards 62
Live and Historical Data in the Workstation Console 63
Viewing Live Query Data in the Workstation Console 63
Enable and Disable Live Mode 64
Viewing Historical Data 64

Chapter 3: Using the Workstation Investigator 69

High-level Views in the Investigator 69
General Investigator Features 69
Agent-Centric View 72
Application-Centric and Business-Centric Views 75
The Triage Map Tab 75
Navigation in the By Frontend Node 77
Navigation in the By Business Service Node 79
Other Application Triage Map Display Elements 79
Application Triage Map Controls 84
List of Physical Locations 86
Limits on Map Display 87
Using the Application Triage Map 87
By Frontend Tree and Metrics 88
Frontend View of the Application Triage Map 90
By Business Service Tree View 94
By Business Service Application Triage Map 96
Using alerts 100
Create and Edit Application Triage Map Alerts 104
Create and Edit Resource Metrics and Alerts 109
Historical Mode in the Application Triage Map 110
The Metric Browser Tab 113
Metrics in the Metric Browser Tab 114
Frontends and Backends 114

CA Technologies Prerelease Documentation
Authorized Use Only

View Data from External Sources Using the Location Map 119
Administering agent connections from the Workstation 128
How User Permissions Affect What You Can View 130
Views in the Metric Browser Tab 131
LeakHunter metrics 145
Using tooltips to view metric names and values in a Data Viewer 146
How time range affects data points 147
Manipulating the contents of Data Viewers 147
The APM Status Console 151
APM Status Console Interface 151
Use the APM Cluster Map 154
Use the Important Events Table 155
Use the List of Active Clamps 155
The Denied Agents List 156
Viewing CA APM Customer Experience Metrics in the Workstation 157
Viewing CA APM Customer Experience Metrics in the Investigator 158
Viewing CA APM Customer Experience Metrics in the Console 158

Chapter 4: Monitoring System Performance and Problems 163

Understanding nominal performance 163
Monitoring performance with the GC Heap metrics 163
Monitor Performance with the GC Monitor Metrics 164
Monitoring Status with the Application Triage Map 165
Monitor Performance with the Location Map 168
Monitor Performance with Frontends Metrics 170
Monitoring performance with backends metrics 171
Monitor performance with the APM Status Console 172
Reading and understanding notifications 173
Alert notifications in dashboards 173
Alert messages 174
Alert notifications in What's Interesting events 175
Other Kinds of Notifications 175
Respond to a Notification 176
Confirm the problem 176
Using Hyperlinks to Find More Information 178
Diagnose the Problem with the Metric Browser Tab 179
Using Live and Historical Metrics 180
Using Search 181
Using Transaction Trace 182
Using thread dumps 183
Diagnosing Problems with CA APM Customer Experience Transactions 185

CA Technologies Prerelease Documentation
Authorized Use Only

Incident-related terminology 186
Wily CEM problem resolution triage metrics on Introscope 187
Incident overview 188
Using Web Start to Drill Down from CA APM Customer Experience to the Investigator 191
Using WebView to Find More Information About an Incident 196
Incident troubleshooting to find root cause 196

Chapter 5: Using the Introscope Transaction Tracer 203

About the Transaction Tracer 203
Automatic Transaction Trace Sampling 204
Transaction Trace overhead 204
Transaction Tracer compatibility with agents from previous releases 205
Starting, Stopping, and Restarting a Transaction Trace 205
Starting a Transaction Trace Session 206
Stopping a Transaction Trace session 207
Restarting a Transaction Trace session 207
Transaction Trace session options 208
Turn Off Low-Threshold Execution Time Warnings 208
Reviewing agents targeted for tracing 208
Using the Transaction Trace Viewer 209
Summary view 210
Trace view 211
WYISC--Sequence View for WUG 214
Correlation IDs in cross-process transactions 214
Clamped Transactions 214
Viewing errors with Transaction Tracer 216
About the Tree view in Transaction Tracer 216
Aggregated Data for Multiple Transactions 217
Using Dynamic Instrumentation 217
Temporarily Instrumenting One, More or All Called Methods 219
Viewing and understanding traces on instrumented methods 220
Viewing Metrics Collected on a Temporarily Instrumented Method 221
Convert Temporary Instrumentation to Permanent 222
Removing Temporary or Permanent Instrumentation 224
Exporting Instrumentation 226
Modifying Instrumentation Level 227
Printing a Transaction Trace window 229
Querying Stored Events 229
Query Syntax 229
Querying Historical Events 230
Saving and exporting Transaction Trace information 234

CA Technologies Prerelease Documentation. Authorized Use Only.

Saving Transaction Trace data 235

Chapter 6: Introscope Reporting 239

Creating Report Templates 239
Adding Report Elements to Reports 241
Defining properties in the Report Editor 243
Setting custom group definitions 252
Time Series Bar Charts 256
Working with report templates 259
Copying or deleting report templates 259
Generating reports from report templates 259
Introscope sample report templates 261
Application Capacity Planning report 261
Production Application Health 261
QA/Test Application Performance 261

Chapter 7: Creating and Using Management Modules 263

About Management Modules 263
Permissions, Domain Enforcement and Element Editing 264
Creating and working with Management Modules 264
Elements in the Management Module Editor 265
Using hyperlinks in the Management Module Editor 267
Agent expressions for Management Modules 268
Naming Management Modules and elements 268
Administering Management Modules 269
Defining agent expressions for a Management Module 272
Configure Metric Groupings 273
Metric name structure 273
Creating a new metric grouping 274
Create and Edit Dashboards 278
About dashboard objects 278
Creating dashboards 280
Editing a dashboard 281
Domain enforcement in dashboard editing 284
Create Data Viewers in a Dashboard 284
Creating an empty data viewer and adding data 288
Setting data-viewing properties of a data viewer 290
Creating dashboard text and graphics 295
Adding shapes and lines to a dashboard 296
Drawing connector lines and adding arrowheads 296
Coloring shapes, lines and connectors 296

CA Technologies Internal Use Only. Authorized Release Documentation.

Creating and editing text 296
Inserting an image on a dashboard 297
Embed an Investigator Tab View in a Dashboard 298
Manipulating dashboard objects 299
Creating and managing custom hyperlinks 302
Dashboard links support agent lens 302
Creating a custom link to a dashboard 302
Creating custom link to an external Web page 303
Defining default links 304
Editing custom links 305
Removing links 305
Monitoring performance with alerts 306
About Simple Alerts 306
Creating Simple Alerts 309
Configuring Simple Alert settings 311
Adding actions 315
About Summary Alerts 316
Creating a Summary Alert 318
About Alert Notification options, messages, and exceptions 322
Alerts and the SmartTrigger feature 322
Generating alert state metrics 325
Working with Alert Downtime Schedules 326
Creating actions and notifications 331
Using Calculators 336
Creating Calculators 336
Calculators and weighted averages 338
Changing operation types in Management Module calculators 338
Using JavaScript calculators 338
Writing JavaScript calculators 338
Running JavaScript Calculators on the MOM 341
Turning off the automatic update for Collectors 342
Deploying Management Modules 343
Updating deployed Management Modules 343
Using the Management Module Hot Deploy Service 343

Appendix A: Introscope Metrics 345

How Introscope monitors application performance 346
Common terms 346
Types of metrics 348
Viewing metrics 349
The Five Basic Introscope Metrics 349

CA Technologies Prerelease Documentation
Authorized Use Only

Average Response Time (ms) 351
Concurrent Invocations 353
Errors Per Interval 354
Responses Per Interval 355
Stall Count 357
Other common metrics 358
Memory-Related Metrics 358
Utilization metrics 365
Socket metrics 366
Thread Dump Metrics 368
Thread pool metrics 368
Connection pool metrics 369
Event metrics 371
Resource Metrics 372
Customer Experience Metrics 373
Customer Experience Transaction Metrics 373
Using perflog.txt 376
Other metrics 376
Application triage map metrics 376
EJB 377
Servlets 378
JDBC 378
JSP (Java Server Pages) 379
RMI (Remote method invocations) 381
Database metrics (SQL) 382
XML (Extensible Markup Language) 383
J2EE Connector 384
JTA (Java Transaction API) 384
JNDI (Java Naming and Directory Interface) 385
JMS (Java Messaging Service) 386
Java Mail 387
CORBA 387
Struts 388
Instance Counts 388
Data About Machines 388
Agent node 389
Agent metrics 390
Enterprise Manager node 390
Data Store node 391
Database sub-node 392
Health Sub-node 393
Internal Sub-node 393

CA Technologies Prerelease Documentation
Authorized Use Only

Problems sub-node 395
Tasks sub-node 395
Harvest metrics 395
Incoming Data Capacity (%) 396
Collector metrics 396
Query metrics 398
Converting Spool to Data metric 399
Overall Capacity (%) metric 399
SmartStor Capacity (%) metric 400
Heap Capacity (%) metric 400
Write Duration (ms) metric 400
Number of Agents metric 400
Number of Metrics metrics 400
Historical Metric Count metric 401
Number of Historical Metrics metric 401

Appendix B: Introscope Extensions 403

SNMP Adapter 403
Creating an SNMP collection 403
Publishing a MIB 405
ErrorDetector 406
Reading and understanding error metrics 407

Index 413

CA Technologies Prerelease Documentation
Authorized Use Only.

Chapter 1: Introduction

Welcome to the CA APM Workstation Guide.

CA APM enables you to manage your application's performance. You use the Workstation to view and manipulate data that is stored by the Enterprise Manager.

This guide describes the Workstation components you use on a daily basis to monitor and manage your application, including the Workstation Console, Investigator, Sample Dashboards, Transaction Tracer, and Reporting.

Note: To see what's new in this user guide, see the section What's new in the Workstation (see page 17).

Note: Portions of this guide offer examples of commands, code, XML or other text printed in plain text. If you use the PDF version of this guide as a source from which to copy such text for use as a template or example for your implementation, you may copy extraneous characters that are invisible vestiges of the PDF conversion process. To avoid this issue, use the HTML version of this guide, contained in the Workstation online help system, as a source for plain text.

This section contains the following topics:

About Application Performance Management (see page 13)
Introscope and the Workstation (see page 15)
What's New in the Workstation (see page 17)
Administering the Workstation (see page 19)
Introscope Workstation Elements (see page 29)
Managing Users (see page 39)
About this Guide and Other CA APM Guides (see page 41)

About Application Performance Management

CA APM provides an effective and comprehensive application performance management strategy that enables you to understand the end-user experience and measure service level agreements (SLAs). You can map all transactions to the end-to-end infrastructure, and conduct incident triage and root-cause diagnoses in a complete and integrated solution.

With CA APM, you can:

- Understand the real user experience.
- Set and manage service level agreements on business services.
- Gain 100 percent transaction visibility.

CA Technologies Confidential
Authorized Use Only

- Determine the source of problems quickly.
- Conduct triage, identify stakeholders, and perform root-cause analyses.
- Prioritize incidents based on true business impact.
- Provide proactive and predictive application monitoring.
- Increase reporting and enable continuous improvement.

Understanding the APM Architecture

Understanding Directory and File Naming Conventions

This guide uses the following conventions in file names and directory paths:

Where you see	It indicates
<Agent_Home>	The top-level directory where the Introscope Agent is installed. This directory is typically named *wily*.
<EM_Home>	The top-level directory where the Introscope Enterprise Manager is installed.
<version>	Version-specific identifier included in file names or displayed in the user interface. For example, the following file name: *com.wily.introscope.soa.dependencymap_<version>.jar* represents a version-specific file name, such as: *com.wily.introscope.soa.dependencymap_v9.1.0.0.jar*
Forward slash (/) path separators	The path separator used in directory names on your operating environment. The forward slash (/) is used on UNIX platforms and in examples throughout this guide, but you should use the separator appropriate to the platform you are using.
Dollar sign ($) environment variables	The environment variable notation used on your platform. The dollar sign ($) is used in UNIX environments and in examples throughout this guide, but you should use the character appropriate to your environment.

CA Technologies Proprietary Documentation. Authorized Use Only.

Introscope and the Workstation

Introscope, through the ProbeBuilder, adds Introscope probes to a Java Application. Using AutoProbe automates this process, with the ProbeBuilder dynamically adding probes to the Java Application when the application starts.

The probes measure specific pieces of information about an application without changing the application's business logic. An Introscope agent is installed on the same machine as the instrumented application. After the probes have been installed in the bytecode, the Java application is referred to as an *instrumented application.* When the Java application with probes is running, it is called a managed application.

As a managed application runs, probes relay collected data to the agent. The agent then collects and summarizes the data and sends it to the Enterprise Manager.

Data collected by the Enterprise Manager can be accessed through one or more Workstations. You can use the Workstation to view performance data, and configure the Enterprise Manager to perform such tasks as collecting information for later analysis, and creating alerts.

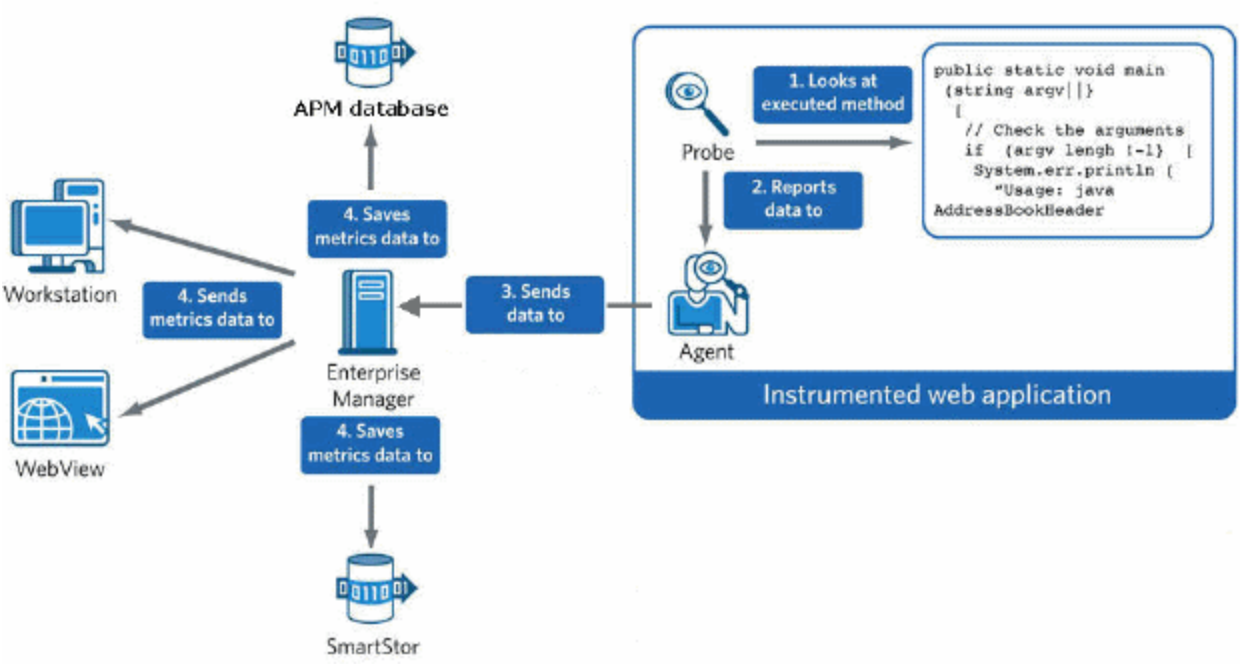
As a managed application runs, Introscope agents collect performance data in real time, and send the information to the Enterprise Manager. The Workstation allows you to configure the Enterprise Manager, organize metrics, define actions based on their values, and display the information you choose in the most convenient format for you.

CA Technologies Prerelease Documentation. Authorized Use Only

How the Workstation fits in an Introscope installation

The Workstation tools help you do the following to better monitor application performance:

- Filter and view performance metrics for various elements of the system your application runs on.
- Drill down to uncover the root cause of system performance issues.
- Create graphical displays of metrics.
- Create reports of system performance data.



The Workstation, Java Web Start, and WebView

Many users find it easy to access the Workstation by using either Java Web Start or WebView. It is important to understand the difference between these two methods.

- Java Web Start uses a command or browser to download and invoke a full Workstation client. This allows you to use full Workstation functionality.

  For more information about Java Web Start, see Launching the Workstation using specific parameters (see page 20).

- WebView allows users to launch a simplified, in-browser, read-only alternative to the full Workstation tools.

  For more information about WebView, see the *CA APM WebView User Guide.*

For more information about launching the Workstation, see Starting the Workstation (see page 19).

What's New in the Workstation

The Introscope Workstation has several new features and some changed functionality. This section provides an overview of these new and changed things:

- The application triage map has many enhancements (see page 18); see the section below for a list of these.
- A new APM Status console interface allows users with apm_status_console_control permission to see the status of stand-alone and clustered Enterprise Managers. See The CA APM Status Console (see page 151).
- Resource metrics (see page 372) are now available on frontend locations. New functionality includes:
  - A new Resources element (see page 93) in the application triage map
  - New Resources tab views (see page 144) in the metric browser and triage map tabs
- A new Location Map allows you to triage problems in the physical and virtual system infrastructure. See View Data from External Sources Using the Location Map (see page 119).
- The Triage Map Configuration Management Module saves definitions of alerts and other objects you create.

  It was formerly named the Admin Definitions Management Module.

Other enhancements and changes

- GC Monitor, a new node in the Browse tree, shows enhanced information about garbage collection. See:
  - Monitoring Performance with the GC Monitor Metrics (see page 164)
  - List and Description of GC Monitor Metrics (see page 359)
- A new Thread Dumps tab allows you to collect a Java thread dump and display current and historical thread dump data. See:
  - Thread Dumps tab (see page 144)
  - Using thread dumps (see page 183)

CA Technologies Prerelease Documentation. Authorized Use Only

- Introscope provides these new metrics:
  - Query memory in transit (bytes) supportability metric in the Metric Browser under Enterprise Manager | Internal
  - <Agent name> | Threads | Deadlock Count
  - Agent connection metrics in the Metric Browser under Enterprise Manager | Connections
    - Max Number of Agent Connection Limit Exceeded Per Interval
    - Max Number of Disallowed Agent Connection Limit Exceeded Per Interval
    - Number of Disallowed Agents
- These Introscope metrics have changed:
  - The Queries Exceeding Max Data Points Read From Disk Limit Per Interval metric now reports on CLW and JDBC queries, not all queries.
  - Queries Exceeding Max Data Points Returned Limit Per Interval metric now reports only on CLW and JDBC queries, not all queries.

Deprecated features

Tooltips in the Workstation tree displays

Tooltips are no longer available from nodes in either the triage map tree or the metric browser tree.

Sampling properties

The following

- introscope.agent.tracer.sampling.maxrate
- introscope.agent.tracer.sampling.initial.period
- introscope.agent.tracer.sampling.reset.period

Application Triage Map Enhancements

The application triage map has a number of new features in release 9.1:

- The application triage map now displays Customer Experience metrics, including a new Customer Experience icon, metrics and alerts on Customer Experience locations. It features better linkage between the map and the Customer Experience browser UI. See Customer Experience metrics (see page 95).
- The application triage map now displays Resource metrics on frontend resources. See:
  - Resources element (see page 93) as it appears in the map
  - Resource metrics (see page 372)
  - Resource alerts

CA Technologies Prerelease Documentation
Authorized Use Only

- You can configure alert thresholds to appear on metric graphs in the Investigator and the Console. See Alert Threshold Line Display (see page 102).
- The Application Triage Map can now be displayed in dashboards. See Adding an application triage map to dashboards (see page 287).

Other changes

Workstation tabs have been renamed:

- The Browse tab is now named the **metrics browser tab**.
- The Map tab is now named the **triage map tab**.

Administering the Workstation

This section has information about starting and stopping the Workstation, and configuring it for tunneling and for SSL.

Start the Workstation

To start Workstation:

1. Launch the Workstation using one of these methods:
   - On Windows, you can:
     - Run *Introscope Workstation.exe.*
     - Click Start > Introscope > Introscope Workstation
   - Using a browser with a URL like:

     http://*<EM_Host>*:8081/workstation

     where *EM_Host* is the host name of the Enterprise Manager.

     You can also use a more complex URL that specifies which page of the Workstation to start at. See Launching the Workstation using specific parameters (see page 20).

     Note: Your first time launching the Workstation, you will be prompted to launch workstation.jnlp up or Save file.
     - Launching workstation.jnlp is recommended.
     - Saving the file and checking "Do this automatically for files like this from now on" is especially *not* recommended, as it will prevent you from properly launching the Workstation via a URL in the future.
   - Using the command line—see Executing Workstation functions from the command line (see page 25).

CA Technologies Prerelease Documentation. Authorized Use Only.

2. In the login dialog, enter:
   - the host name or IP address

     Note: You can use the IP address instead of the host name only if both your client computer and the host computer support the same IP protocol.
   - the Port number
   - your User Name
   - your Password

     Tip: The Workstation will remember the last five login attempts, so if you have entered host and user information previously, the Host, Port and User ID fields are drop-down fields from which you can select credentials.
3. Click Connect, or to make the current host and user information the default for future logins, click Set Defaults.

   If authentication was successful, the Console opens. If authentication was unsuccessful, a message notifies you of the failure and the login window reopens.

Note: If a user tries to log in but does not have permissions defined in *domains.xml* or *server.xml*, Workstation login fails.

Launching the Workstation Using Specific Parameters

You can launch Workstation using specific parameters that specify which Workstation page you want to access. These parameters can be used in either:

- a Java launch command issued on the command line
- a URL that launches the Workstation using Java web start

Example 1

For example, in the command line, the *-page* and *-agent* options would be used like this:

```
java -client -Xms64m -Xmx256m -Dsun.java2d.noddraw=true -jar launcher.jar -
consoleLog -noExit -product com.wily.introscope.workstation.product -
name "Introscope Workstation" -install ".\\product\\workstation" -
configuration ".\\product\\workstation\\configuration" -page
investigator -agent "SuperDomain|localhost|WebLogic|WebLogic Agent"
```

In a URL, the same combination would be:

http://*<localhost>*:8081/workstation?page=investigator&agent=SuperDomain|localhost|WebLogic|WebLogic%20Agent

CA Technologies Prerelease Documentation. Authorized Use Only.

Note the way each of the above examples uses URL encoding to handle the space character in the agent name.

- In the command line example, quotes are used around the entire agent name because the name contains a space.
- In the URL example, a space character is rendered as *%20*.

Use standard URL encoding to escape special characters in agent or metrics names that you want Introscope to treat as a literal part of the name rather than as part of the URL.

Example 2

If the agent name is “MyAgent%1” you use this string in the URL:

```
MyAgent%251
```

in which %25 is the URL encoding for the literal % character.

Example 3

If the agent name is “WhatIsThisAgent??” you use this string in the URL:

```
WhatIsThisAgent%3F%3F
```

in which %3F is the URL encoding for the literal ? character.

The table below lists other parameters you can use.

Options	Description
-loginimmediate	Suppresses the login screen and logs into Workstation immediately using specified hostname and port number, or default values.
-loginhost <hostname>	Specifies login host name; defaults to *localhost* if unspecified.
-loginport <portnumber>	Specifies login port number; defaults to 5001 if unspecified.
-loginresponse <values>	Specifies authentication values for username and password in a comma-separated list.
-page	The name of the Workstation screen to be launched. You must include this parameter with every request to the Workstation Command Line Interface. Supported values: ■ investigator ■ historicalquery ■ console

CA Technologies Prerelease Documentation
Authorized Use Only.

-agent	The fully qualified agent name to display in the Investigator window. Required if the *page* parameter is *investigator*. Use URL encoding to render special characters in agent names.
-metric	The metric path to display in the Investigator window, for a specified agent. You must specify an agent if you use this parameter. Use URL encoding to render special characters in metric names.
-start	The start time, in standard Java format of milliseconds, for a historical time range in the Investigator window, or the start time for a Transaction Tracer Historical Query, depending on the value of the *page* parameter. **Note:** The *start/end* or *guid* parameters are required if the *page* parameter is *historicalquery*.
-end	The end time, in standard Java format of milliseconds, for a historical time range in the Investigator window, or the end time for a Transaction Tracer Historical Query, depending on the value of the page parameter. The *start/end* or *guid* parameters are required if the *page* parameter is *historicalquery*. The following example uses Java timestamp values. You can convert calendar dates to Java timestamp values using widely available converters, including some available on the internet. http://*<localhost>*:8081/workstation?page=historicalquery&start=1135686483474&end=1136686483474
-guid	The unique identifier for a transaction to display in the Transaction Tracer Historical Query window. The start/end or guid parameters are required if the page parameter is *historicalquery*. For example: *http://<localhost>:8081/workstation?page=historicalquery&guid=aRx345*
-agentSpecifier	Filters data to limit the dashboard display to data from the agent you specify. Can be used only when the page parameter = *console*. The argument to the AgentSpecifier parameter must contain the agent name including the Enterprise Manager host name. Special characters, such as the | symbol which separates elements of the agent name, must be escaped with backslashes. Substitute the string *%20* for spaces in agent names. In this example, the dashboard will display only data from WebLogic Agent: *http://<localhost>:8081/workstation?page=console&agentSpecifier=machine1|WebLogic|WebLogic%20Agent&metric=GC%20Heap:Bytes%20In%20Use*

CA Technologies Proprietary & Confidential Documentation
Authorized Use Only

-dashboardName	Specifies a dashboard to display. Can be used only when the page parameter = *console*. Substitute the string *%20* for spaces in dashboard names. In this example, the URL will jump to the dashboard named GC Memory In Use: *http://<localhost>:8081/workstation?page=console&dashboardName=GC%20Memory%20In%20Use&metric=GC%20Heap:Bytes%20In%20Use*

Executing one of the URLs above (or launching a Workstation with an equivalent Java command line) starts a Workstation instance and opens the appropriate window. Subsequent URL requests open a new window in the existing Workstation instance.

JVM requirements for Java Web Start

Using Java Web Start to launch the Workstation requires:

- JVM 1.5 or above to recognize the .jnlp file extension contained in the Java Web Start instructions.
- JVM 1.6.0_05 or above to actually launch the Workstation.

Using Java Web Start downloads a temporary copy of the Workstation client from the Enterprise Manager to your computer. Computers using proxy authentication to connect to an Enterprise Manager might not automatically download the correct JVM if it is missing. If you encounter this problem, install the correct version of JVM manually before attempting to use Java Web Start.

On the client system, Java Web Start will launch the workstation using a Java version defined by these two files in the Enterprise Manager:

- *<iscroot>\product\enterprisemanager\plugins\com.wily.introscope.workstation.webstart_9.0.0\WebContent\jnlp\workstation.jsp*
- *<iscroot>\product\enterprisemanager\plugins\com.wily.introscope.workstation.webstart_9.0.0\WebContent\jnlp\com.wily.introscope.workstation.feature_9.0.0.jsp*

In both of these files, you will find a j2se node with a version attribute to use in determining the Java version to be used to launch the Workstation:

```
<j2se version="1.5*&1.5.0_15+ 1.6*&1.6.0_05+" ... />
```

Note: To parse these instructions, the client system must have an installed JVM of 1.5.0 or later. If it has JVM 1.4.2 or earlier, or no JVM at all, it will not be able to read these instructions. In this case you should download JVM 1.5.0_15 from the java.sun.com website, then proceed.

The attribute lists Java version ranges in preferred order from first to last. Each range is separated by a single space. Java Web Start checks the client system for all version ranges in the listed order and installs the first qualifying version that it finds.

CA Technologies Internal Use Only. Pre-release Documentation.

If you have JVM 1.5 or above, the Enterprise Manager is able to read the Java Web Start instructions, download the required version of JVM 1.6, and use it to download and start the temporary copy of the Workstation client.

Note these exceptions:

- If you have JVM 1.4.2 or earlier, or no JVM at all, Java Web Start may not automatically download the correct version. Download JVM 1.5.0_15 manually from http://www.java.com/en/download/manual.jsp, then proceed.
- If the client system is Windows x64, the Java version must be the 32-bit Java 1.6 update 18 or later. Download the 32-bit JVM 1.6.0_18 manually from http://www.java.com/en/download/manual.jsp, then proceed.

If Java Web Start does not find a pre-installed JVM that matches the listed version ranges, it downloads the newest available version permitted by all of the version ranges taken together, regardless of order.

For example: When the j2se version is set to *"1.5*&1.5.0_15+ 1.6*&1.6.0_05+" ...*

- Web Start will first check the client system for Java 1.5, update 15 or later.

```
1.5*      = version starts with 1.5
&         = and
1.5.0_15+ = version is 1.5 update 15 or later
```

- If a valid update level of Java 1.5 is not found, Web Start will then check the client system for Java 1.6, update 5 or later.

```
1.6*      = version starts with 1.6
&         = and
1.6.0_05+ = version is 1.6 update 5 or later
```

- If a valid update level of Java 1.6 is not found, Web Start downloads the newest available version permitted by any of the version ranges.
- It would never download Java 1.7 since the version range does not permit it.

Java 1.5.0_15 and 1.6.0_05 have been deliberately chosen as minimum update levels since these updates are the earliest to address critical security issues in Java (see http://nvd.nist.gov/nvd.cfm?cvename=CVE-2008-1185). Users can customize the ranges to allow earlier update levels at their own risk.

Note: If you edit the version range list, be sure to edit it in both files: *workstation.jsp* and *com.wily.introscope.workstation.feature_9.0.0.jsp*.

Connecting to alternate Enterprise Managers

You can start multiple Workstation application instances on different Enterprise Manager hosts from a single browser, using the parameters specified in Launching the Workstation using specific parameters (see page 20). To connect to an alternate or different Enterprise Manager, change the *loginHost* parameter as appropriate.

CA Technologies Prerelease Documentation. Authorized Use Only.

End Your Workstation Session

You can log out of the Workstation in addition to quitting the application.

Logging out from the Workstation

Logging out from the Workstation ends the current session, but does not shut it down, so that you can log in again from the Authentication dialog. This is useful if you want to log in with different connection parameters, such as a different host, port, user name, or password.

Workstation saves the number of open Investigator and Console windows when you log out, and the same configuration appears when you next log in.

To log out from the Workstation:

- Select Workstation > Logout.

Exiting the Workstation

Exiting the Workstation logs you out of the Workstation and stops the Workstation process.

When you exit the Workstation, it saves the number of open Investigator and Console windows, so the same configuration appears when you next log in.

To exit the Workstation:

- Select Workstation > Exit Workstation.

Execute Workstation Functions from the Command Line

You can execute Workstation functions from a command line. This is useful if you need to execute these functions from a script for the purpose of batching or scheduling the functions.

For more information about Command Line Workstation, see the *CA APM Configuration and Administration Guide.*

CA Technologies Prerelease Documentation
Authorized Use Only.

To execute Workstation functions from the command line:

1. Change to the Enterprise Manager home or *<EM_Home>* directory.
2. Execute the Workstation start command, using the examples below as models.

   Here is an outline of the command:
   ```
   java [optional arguments] -jar launcher.jar [Eclipse arguments]
   ```

   Here is an example of a full Workstation start command:
   ```
   java -client -Xms64m -Xmx256m -Dsun.java2d.noddraw=true -jar launcher.jar
   -consoleLog -noExit -product com.wily.introscope.workstation.product -name
   "Introscope Workstation" -install ".\\product\\workstation" -configuration
   ".\\product\\workstation\\configuration"
   ```

Follow these guidelines:

- On UNIX, change escaped backslashes to forward slashes.
- If adding your own optional JVM arguments, insert them before the *-jar* argument. The following arguments appear in the example.
  - *-client*—Runs the JVM in client mode
  - *-Xms*—initial Java heap size
  - *-Xmx*—maximum java heap size for the application to use
  - *-Dsun.java2d.noddraw=true*—Optional. Helps resolve potential difficulties between drivers and Java APIs.

Modifying the Eclipse arguments (everything from *-consoleLog* onward) is not recommended except at the request of CA Support.

Additional parameters available for using Command Line Workstation are listed in the table in Launching the Workstation using specific parameters (see page 20).

Configuring the Command Line Workstation Log

You can configure CA APM to log Command Line Workstation (CLW) commands to the Enterprise Manager console and the IntroscopeEnterpriseManager.log file, which is located in the *<EM_Home>*/logs directory.

To configure Command Line Workstation log:

1. Open the IntroscopeEnterpriseManager.properties file located in the *<EM_Home>*\config directory.
2. Configure these properties in the IntroscopeEnterpriseManager.properties file to enable the logging of CLW commands in the log file and on the Enterprise Manager console:
   a. Set log4j.additivity.Manager.CLW=true.

CA Technologies Pre-Release Documentation Authorized Use Only.

Note: The default value for this property is false.

b. Set log4j.logger.Manager.CLW=DEBUG.

The default value for this property is INFO.

Configuring HTTP tunneling for the Workstation

You can configure the Workstation to connect through a proxy server to the Enterprise Manager. This is necessary for a forward-proxy server configuration where the Workstation is running behind a firewall that only allows outbound HTTP traffic routed through the proxy server.

Note: Because tunneling imposes additional CPU and memory overhead on the managed host and Enterprise Manager beyond that expected for a direct socket connection, do not set up Workstation HTTP tunneling if a direct socket connection to the Enterprise Manager is feasible.

IMPORTANT: HTTP/1.1 is required to enable Workstation HTTP tunneling.

To use Workstation tunneling:

- Edit the HTTP Tunneling Proxy Server section of *IntroscopeWorkstation.properties* to specify the tunneling connection:

  a. Uncomment the lines beginning with *transport.http...*

  b. Provide the host, port, username and password of the proxy server.

```
################################
# HTTP Tunneling Proxy Server
#----------------------
# These properties apply if the Workstation is tunneling over HTTP
# and must connect to the Enterprise Manager through a proxy server (forward
proxy).
# If the proxy server cannot be reached at the specified host and port,
# the Workstation tries a direct HTTP tunneled connection to the Enterprise
Manager
# before failing the connection attempt.
#transport.http.proxy.host=
#transport.http.proxy.port=
# These properties apply if the proxy server requires authentication.
#transport.http.proxy.username=
#transport.http.proxy.password=
```

Configuring the Workstation to use SSL

The Workstation ordinarily uses HTTP to connect to the Enterprise Manager. You can configure connections through HTTPS/SSL, optionally using certificates.

CA Technologies Pre-release Documentation
Internal Use Only

To configure the Workstation to connect to the Enterprise Manager using SSL, you edit the *IntroscopeWorkstation.properties* file for the following properties:

Property	Description
truststore location path	Path to the location of a truststore containing trusted Enterprise Manager certificates. Note that on Windows, a backslash must be escaped with another backslash. Example: *transport.tcp.truststore= C:\\Introscope\\config\\internal \\server\\keystore*
transport.tcp.trustpassword	Password for the certificate truststore Example: *transport.tcp.trustpassword=password*
workstation certificate location path	Path to the location of the trusted certificate for the Workstation. Escape backslashes as in the example above.
transport.tcp.keypassword	Keystore password Example: *transport.tcp.keypassword= password*
transport.tcp.ciphersuites	List of cipher suites, separated by commas. If this property is blank, Workstation will use the default list. Example: *transport.tcp.ciphersuites= SSL_DH_anon_WITH_RC4_128_MD5, SSL_RSA_WITH_NULL_MD5*

Things to note:

- Specify a truststore to configure the Workstation to authenticate the server (Enterprise Manager). If no truststore is specified, the server is automatically trusted.
- Specify a keystore only if the Enterprise Manager has been configured to require client authentication.

CA Technologies Prerelease Documentation. Authorized Use Only.

Introscope Workstation Elements

You use the Workstation to view metric data in different forms. Authorized users can perform administrative and configuration functions. The Workstation presents information in these windows:

Console

Shows data in dashboards, which contain Data Viewers.

Investigator

Presents tree views and map views of agents, applications, resources, and metrics.

Management Module Editor

Presents a tree view of Management Modules and elements, allowing you to create and edit Management Modules.

Dashboard Editor

Enables users with *write* permission for a Domain (or SuperDomain) to create and edit Data Viewers and other dashboard objects such as imported images, shapes, lines, and text.

Data Viewers

Visual presentation of data based on the type.

About the Workstation Console

The Console is the default view when you start the Workstation, and contains dashboards that show performance data in graphical views. Dashboards are basic tools for viewing management data in Introscope.

The Sample Management Module provides a set of sample dashboards. Authorized users can create custom dashboards using the Dashboard Editor—see Creating dashboards (see page 280).

You can have more than one Console window open at the same time.

To open a new Console window:

- Select Workstation > New Console.

For more information about how to view information using the Workstation Console, see Chapter 2, Using the Workstation Console (see page 45).

CA Technologies Prerelease Documentation. Authorized Use Only

About the Workstation Investigator

You use the Investigator to view application and system status, to search, and to browse metric data using a tree structure.

You can have more than one Investigator window open at the same time.

To open a new Investigator window:

- Select Workstation > New Investigator.

The Investigator opens, showing data for your Java or .NET application.

You can also open an Investigator window from the Console by double-clicking on some dashboard elements, depending on how the element was created. See Using hyperlinks to navigate (see page 46).

Application-centric and Agent-centric Views

Investigator displays your application infrastructure in two main ways—application-centric and agent-centric. Each has a top-level tab, *triage map* and *metric browser,* respectively.

Triage map tab

The triage map tab shows an application-centric or business process-centric view of your monitored applications. You use it to do the following tasks:

- View deployed applications and business-centric metrics, in both live and historical modes.
- Discover dependencies between application layers and constituent pieces of each layer.
- Monitor high-level health indicators for applications and their constituent frontends, backends, and middleware.
- Monitor aggregated health metrics for applications.

Metric browser tab

The metric browser tab shows an agent-centric view of your monitored applications. You use it to do the following tasks:

- View applications and metrics organized in a tree hierarchy.
- Monitor detailed metrics for each layer of technology.
- Use Transaction Tracing and dynamic instrumentation to triage anomalies in application performance.

CA Technologies Prerelease Documentation. Authorized Use Only.

Note: The Workstation does not display the triage map tab if the Enterprise Manager you logged in to has been configured as a collector on a cluster. To use the triage map tab tools on a clustered application, log in to the MOM Enterprise Manager.

How applications appear in different views

Frontend applications appear slightly differently in the triage map tab and in the metric browser tab. Given an application named test0, the frontend appears as follows:

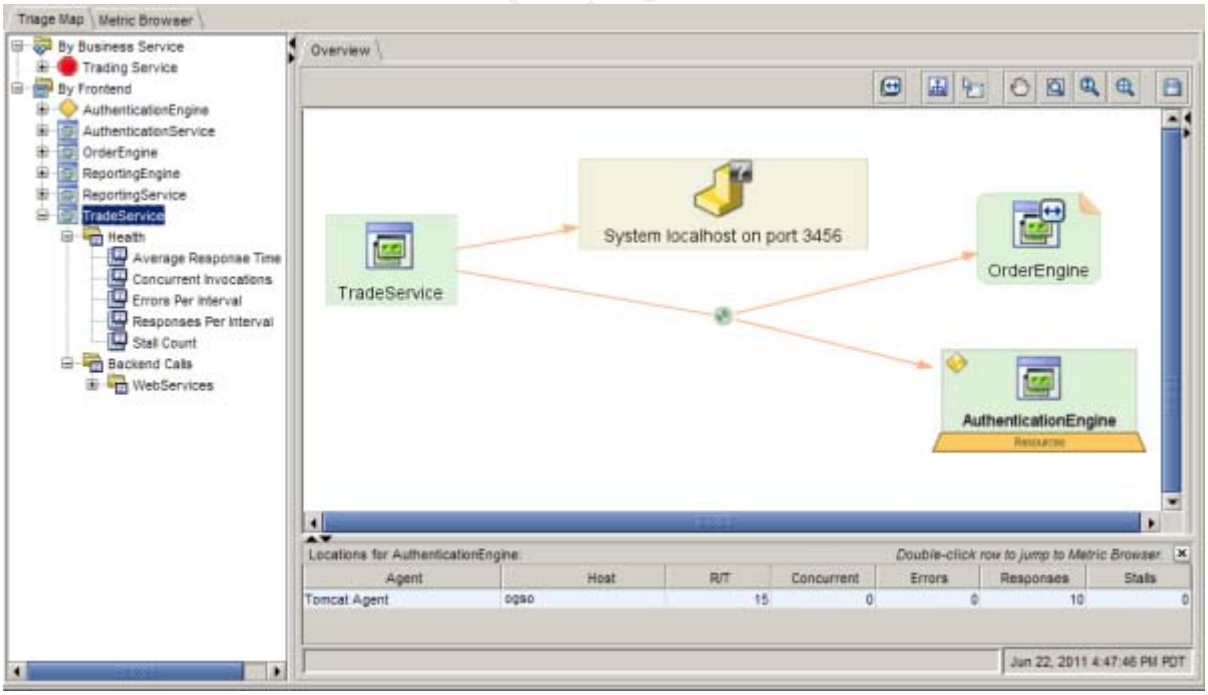
- In the triage map tab, test0 appears as a Frontend Application.
- In the metric browser tab, test0 appears as an App under the Frontends node.

How metrics are aggregated differently in tab views

The application-centric view in the triage map tab displays aggregated health metrics, while the agent-centric view in the metric browser tab displays metrics returned only from the single host where the agent is configured.

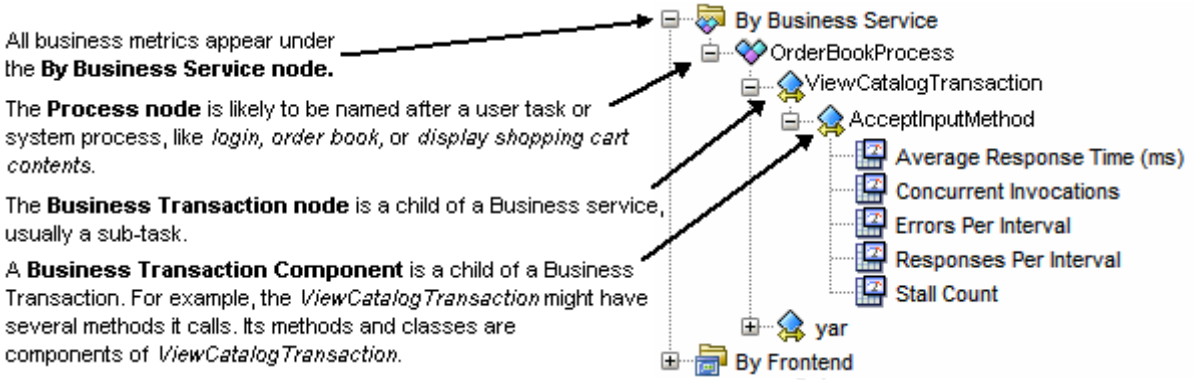
The Application Triage Map

When the triage map tab is active, you can view a visual display of an application. This application-centric visual display or "application triage map" allows you to view application components and their dependencies, view health indicators for components and subcomponents, and drill into underlying metrics.



How business metrics appear

On the triage map tab, the Workstation displays business metrics under the By Business Service folder:



The triage map tab also displays a business-centric dependency map, as seen in By Business Service application triage map.

Business metrics have the form:

```
<Host Name>|<Process Name>|<Agent Name>|By Business Service|<Business Service>|<Business Transaction>|<Business Transaction Component>
```

The metrics which Investigator displays for each business transaction component depend on how each business service, business transaction, and business transaction component have been configured. The process of configuring business metrics is documented in the *CA APM Transaction Definition Guide.*

A note on alerts in historical mode

When the application triage map displays historical data, alert indicators in the triage map tree continue to display *current* status, not historical status.

More information

More information about reading and understanding the application triage map is available. See:

- Navigating in the triage map tab (see page 75)
- Responding to a notification (see page 176)

CA Technologies Proprietary/Confidential Document. Authorized Use Only.

About the Management Module Editor

You use the Management Module Editor to create or edit a Management Module, which contains a set of Introscope monitoring configuration information. Management Modules are listed for each domain, and contain objects, known as elements, that contain and organize data with monitoring logic—alerts, actions, and dashboards.

Note: If you have a full CA APM license, you can create, edit, or delete information in the Management Module Editor. If you do not have a full license, you can only view information here.

The Management Module Editor tree lists the Management Modules deployed to the Enterprise Manager, by domain, and the elements in each Management Module.

The right side of the Management Module Editor presents the current configuration settings for the element selected in the tree.

An authorized user can modify elements in the Management Module Editor.

More information:

Creating and Using Management Modules (see page 263)

About the Dashboard Editor

The Dashboard Editor provides tools for creating and laying out Data Viewers, shapes, lines, text boxes, and connectors. Users with appropriate permissions can create and edit dashboards and dashboard objects such as imported images, shapes, lines, and text—see Creating and editing dashboards (see page 278).

About Data Viewers

Data Viewers in the Metric Browser Tab viewer pane or in a dashboard display data from an Introscope-enabled application in a visual form. Data appears in a Data Viewer based on the type of data—for example, metrics appear as graphs, and alerts (formerly referred to as traffic lights) appear as indicators. Data Viewers can display data from a metric, a resource, or an element, such as an alert.

Note: The time displayed on data viewers is the clock time on the computer hosting the Enterprise Manager. The time zone (for example, PST for Pacific Standard Time) is part of this displayed time.

CA Technologies Prerelease Documentation
Authorized Use Only

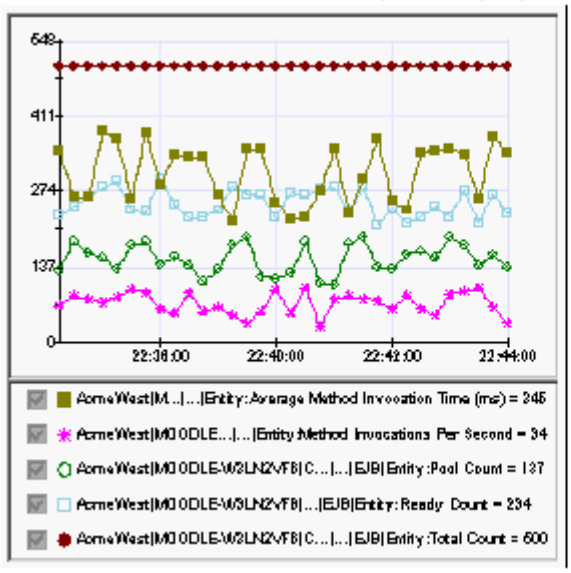
Data Viewer Types

Data types have a default data viewer type and alternative viewers.

Data type	Default Data Viewer type	Can also be viewed as
Metric	Graph	Dial Meter, Bar Chart, Graphic Equalizer, String Viewer, Text Viewer
Metric Grouping	Graph	Bar Chart, String Viewer
Alert	Alert indicator	Graph, Bar Chart, or String Viewer
Calculator	Graph	Dial Meter, Bar Chart, Graphic Equalizer, String Viewer

Depending on the type of metric or element, the Workstation can display the data in a Data Viewer with the view display types shown here.

Graph

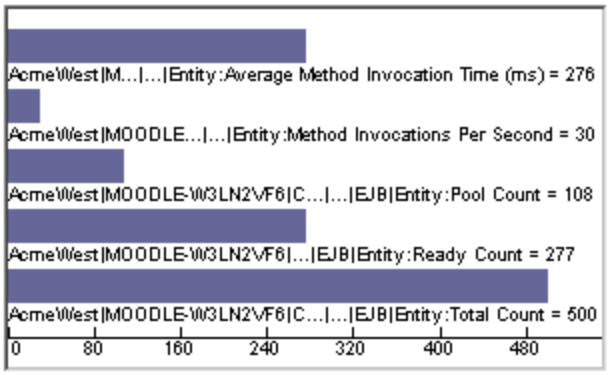


Graphs plot values over time. In real-time views, the Graph dynamically displays the most recent time period that fits in the graph.

If the graph displays an alert, caution and danger thresholds appear as yellow and red lines, respectively.

You can change the scale of graph charts while viewing live data, to see data in a more readable view—see Changing the scale of graph charts (see page 148).

Bar Chart

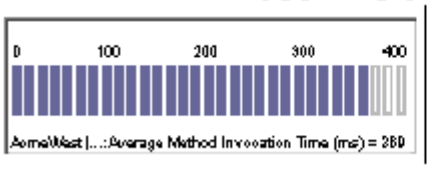


Bar charts display current data values as horizontal bars.The bar chart is the default view for Top N Filtered Views.

If a bar chart is showing an alert, the bars will be either green, yellow or red to correspond to alert status.

The bar chart is available for live data viewing only.

Graphic Equalizer

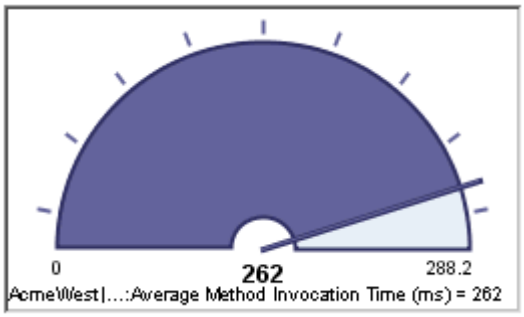


Graphic equalizers show the current value of the data, as well as recent high levels.

A graphic equalizer can only display data for a single metric.

The Graphic Equalizer viewer type is only seen in a WebView Console dashboard.

Dial Meter

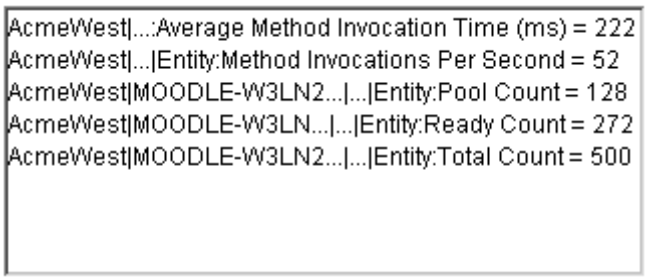


CA Technologies Confidential/Proprietary Information. Authorized Pre-release Documentation. Use Only

Dial meters depict current data values as positions on a half-round dial.

The dial meter viewer type is only seen in a WebView Console dashboard.

String Viewer



String viewers can display a value as a line of text. String viewers allow some values to display in a relatively small space. You can also use a String Viewer for simple values that do not change, such as Launch Time or IP Address.

Note: With live metrics from connected agents, most data is valid only for the most recent 15 second time slice. So when an agent disconnects, string metrics show no value. However, a few constant metrics, such as the Agent's original Launch Time, remain valid whether or not the Agent is or is not presently connected, and so will always appear until the agent is unmounted.

Text Viewer



Text viewers show the text for data when new values are added to old ones—for example, a system or exception log.

About Alerts and Alert Indicators



Alert indicators show whether a metric has crossed a threshold:

- Green disc = status normal
- Yellow diamond = caution threshold was crossed
- Red octagon = danger threshold was crossed
- Gray disc = the alert has no data.

Alert indicators can appear as they do above, as an array of three indicators in which the active indicator tells the status. More often, they appear as a single indicator which changes color and shape when its status changes.

Alert indicators can appear in several modes and locations:

- in the application triage map, as shown in the illustration in By Business Service Application Triage Map (see page 148)
- in dashboards
- in the Overview tab: see Application Overview (see page 133)
- as threshold lines on a graph: see Alert Threshold Line Display (see page 102)
- as colors in table cells, where the functionality is supported; see the illustration in the topic Resources Element (see page 93)
- in place of tree nodes

Understanding the difference between *alerts* and *alert indicators*

It is important to understand exactly what an alert is. Be sure to distinguish between:

- the **alert** itself, the definition of which includes saved attributes like:
  - threshold values
  - the metric grouping to which it is linked
  - the Management Module to which it belongs

CA Technologies Proprietary Documentation. Authorized Use Only.

- the **alert indicator**, which is a graphical display of alert status
- an **action** which might be associated with the alert.

An alert commonly is linked with an action, but actions are separate Management Module objects. They are associated with one another as part of the task of configuring an alert. An *alert notification* is one of the possible actions you can associate with an alert.

For more information about how alerts are configured, see:

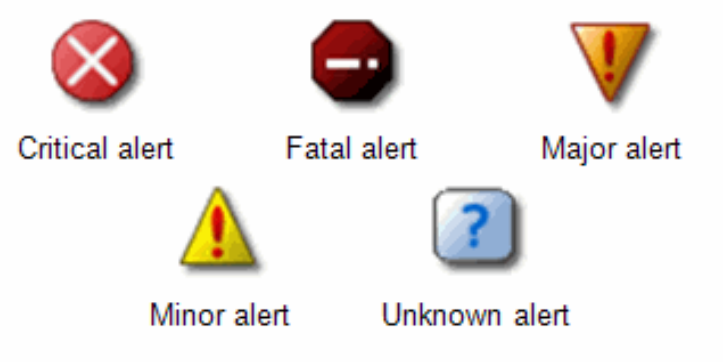
- Monitoring Performance with Alerts (see page 306)
- Using Alerts (see page 100)
- Creating and Editing Application Triage Map Alerts (see page 104)

For more information about alert actions and notifications, see:

- Reading and Understanding Notifications (see page 173)
- Configuring Simple Alert Settings (see page 311)
- Add Actions to Alerts (see page 315) and
- Creating Actions and Notifications (see page 331)

How Catalyst Alert Indicators Appear

Alert indicators imported from Catalyst have a different appearance from Introscope alert indicators.



These alert indicators appear on elements imported from Catalyst. For more information, see Viewing Data using the Location Map (see page 119).

CA Technologies Prerelease Documentation. Limited Use Only.

Managing Users

You manage users through user permissions and user preferences. However, most permissions are set on the Enterprise Manager level. For information about how to set user and group permissions, see the *CA APM Security Guide*.

User Permissions

Workstation users are assigned a user name, password, and certain permissions. Permissions are granted at the Domain and Enterprise level.

Some Workstation functions require specific permissions. For example, to publish a MIB (Management Information Base, a directory of information used by network management protocols), a user must have publish_mib permission for the server. Your Introscope administrator assigns these to you.

If you do not have sufficient permissions for a function, the function is disabled.

For more information about user permissions, see the *CA APM Installation and Upgrade Guide*.

User Preferences

You use Introscope user preferences to specify:

- a home dashboard
- whether to display Management Module names beside dashboard names in the Console
- low-threshold execution-time warnings for Transaction Tracer

Setting a Home Dashboard

Dashboards are pre-configured windows that present graphical views of current or historical performance and availability metrics.

To change your home dashboard:

1. Select Workstation > User Preferences.
2. Select a dashboard by doing one of these:
   - Select a dashboard from the drop-down list.
   - Click Choose, enter a search string to narrow the selection, and select from the remaining list.
3. Click Apply.

CA Technologies Prerelease Documentation. Authorized Use Only

Displaying the Management Module and Domain Names

You can use the same name for dashboards that are in different Management Modules, and use the same name for Management Modules that are in different Domains.

You can set User Preferences to display the name of the Management Module and Domain that contain the dashboard.

To display the Management Module name next to the dashboard name:

1. Select Workstation > User Preferences.
2. Select Show Module and Domain name with dashboard name.
3. Click Apply.

   The Management Module and domain that contain the dashboard appear after the dashboard name.

Note: Domain information does not appear if you have access to only one Domain.

Turning Off Low-Threshold Execution Time Warnings

If you are running the Transaction Tracer and set the threshold execution time to less than one second—to perform a deep analysis, for example—you might see continual warnings. The warnings indicate increased overhead because of increased traces, so you might want to turn them off in a production environment.

To turn off the warnings about low-threshold execution time:

1. Select Workstation > User Preferences.
2. Click the Transaction Tracer tab.
3. Check the Don't warn when threshold is less than 1 second checkbox.

For more information about Transaction Tracing, see Using the Introscope Transaction Tracer (see page 203).

Managing Language Settings

When using the Workstation tools:

- User dialogs reflect the regional language set in the Control Panel on your computer.
- You can set properties in Introscope reports to use a specific language setting separate from the regional language set for your computer.

CA Technologies Proprietary Documentation - Authorized Use Only.

More information:

Defining properties in the Report Editor (see page 243)

About this Guide and Other CA APM Guides

Product Documentation

The following sections list the titles in the CA APM documentation set.

- CA APM documentation (see page 41)
- Introscope documentation (see page 42)
- CA APM Customer Experience documentation (see page 43)
- CA APM extensions documentation (see page 43)

You can view and search all of the titles in the CA APM documentation set from the CA APM bookshelf on the CA Support website.

CA Wily APM documentation

This is the documentation unique for CA APM. This documentation is pertinent to both Introscope and CA APM Customer Experience.

- *CA APM Release Notes* — Release summary and readme information for CA APM.

  Note: In previous releases, this manual was titled the *CA APM ReadMe.*
- *CA APM Readme* — Important last-minute release information for CA APM; available on the product download site or the Wily Community Site.

  Note: In previous releases, this manual was titled the *CA APM Known Issues.* This document is not in the product build.
- *CA APM Glossary* — Terms and concepts for CA APM.
- *CA APM Installation and Upgrade Guide* — Installation requirements; installing CA APM, including installing Enterprise Manager, APM database, Workstation, WebView, CA APM Customer Experience, TIM; upgrading from previous releases.
- *CA APM Overview Guide* — A broad overview of CA APM components and architecture. Explains terms and concepts used in a CA APM deployment.
- *CA APM Configuration and Administration Guide* — Combines configuration and administration information for Introscope and for CA APM Customer Experience. Introscope properties and TESS default properties references are documented in appendixes.

CA Technologies Prerelease Documentation. Authorized Use Only.

- *CA APM Security Guide* — Choosing and configuring CA APM, Introscope, and CA APM Customer Experience security solutions.
- *CA APM Sizing and Performance Guide* — Sizing, tuning, and capacity planning for your CA APM deployment and components.
- *CA APM Transaction Definition Guide* — Transaction definition processes and procedures for CA APM; describes the necessary steps to record, define, and verify customer transactions.

Introscope Documentation

This is the documentation unique for Introscope. CA APM documentation is also pertinent for Introscope.

- *CA APM ChangeDetector User Guide* — Using CA APM ChangeDetector to monitor and report changes in application files and configuration.
- *CA APM Environment Performance Agent Implementation Guide* — Implementing Environment Performance Agent (EPAgent) with Introscope to monitor system information, including process availability, disk statistics, web application server and web server logs, Solaris KStat and HTTP service availability. The guide provides instructions for installing, configuring, and using EPA.
- *Introscope Extension for CA SYSVIEW Guide* — Installing, configuring, and using the SYSVIEW Agent extension, which allows you to manage application performance of distributed applications accessing mainframe back ends and trace transactions from distributed applications to mainframe CICS transactions. It also lets you monitor the health metrics of critical mainframe components.
- *CA APMJava Agent Implementation Guide* — Installation, configuration, and use of the CA APM Java Agent.
- *CA APM.NET Implementation Agent Guide* — Installation, configuration, and use of the CA APM .NET Agent.
- *CA APM for SOA Implementation Guide* — Installation and configuration information for using CA APM for SOA and SOA platform extensions with Introscope.
- *Introscope WebView Guide* — Using WebView to view Introscope data in dashboards and Investigator.
- *Introscope WebView Guide for SAP* — This guide on how to view Introscope data in dashboards, Investigator, and Transaction Tracer, is for WebView for SAP users.
- *Introscope Workstation User Guide* — Using Introscope dashboards, Investigator tree and application triage map, Transaction Tracer, and reporting. Introscope metrics overview and descriptions.

CA Technologies Product Documentation. Authorized Use Only.

CA APM Customer Experience Documentation

This is the documentation unique for CA APM Customer Experience. The CA APM documentation is also pertinent for CA APM Customer Experience.

- *CA APM Integration Guide* — Integration of CA APM and CA Technologies products (such as CA SiteMinder, CA NSM, CA Service Desk, CA APM TG, and CA CMDB) with CA APM Customer Experience.
- *CA APM TG Implementation Guide* — Using the CA APM TG (WTG) to monitor the availability, health, and performance of web sites and web services from the perspective of a user attempting to access web sites. How to use the WTG Agent to generate synthetic transactions that you can monitor using CA APM.

Introscope PowerPack documentation

This is the documentation for PowerPacks.

- *PowerPack for CA SiteMinder Web Access Manager Guide* — Installing, configuring, and using the PowerPack for CA SiteMinder Web Access Manager.
- *PowerPack for IBM WebSphere Application Server Distributed Environments Guide* — Installing, configuring and using the CA Wily PowerPack for IBM WebSphere Application Server Distributed.
- *PowerPack for Oracle Database Guide* — Installing, configuring, and using the CA Wily PowerPack for Oracle Database.
- *PowerPack for SharePoint Portal Guide* — Installing and configuring the PowerPack for SharePoint Portal to monitor your SharePoint Portal components during development, QA, staging, and production.
- *PowerPack for SiteMinder SNMP Collector Guide* — Installing and configuring the PowerPack for CA Wily SiteMinder SNMP Collector. Understanding the associated metrics.
- *PowerPack for Web Servers Guide* — Installing, configuring, and using the CA Wily PowerPack for Web Servers.
- *PowerPack for WebLogic Portal* — Installing, configuring, and using the CA Wily PowerPack for WebLogic Portal.
- *PowerPack for WebLogic Server Guide* — Installing, configuring, and using the Introscope PowerPack for WebLogic Server.
- *PowerPack for IBM WebSphere MQ Connectors and Messaging System Guide* — Using Introscope to view metrics from IBM WebSphere MQ.
- *PowerPack for WebSphere Portal Guide* — Installing, configuring, and using the CA Wily PowerPack for WebSphere Portal.

CA Technologies Internal Use Only Authorized Release Documentation

Workstation Online Help

Workstation online help consists of documentation available in both HTML and PDF formats.

Online Documentation

To open Workstation Help:

1. From the Help menu, select Help and Documentation.

   The top-level Help window appears.

2. Select from the two available Help links.

Note: In the UNIX environment, the Help system is hard-coded to use the Mozilla browser. You must have Mozilla in your PATH environment variable for the links displayed in the top-level Help window to be functional. For example:

```
PATH=/usr/lib/mozilla.sh:$PATH
export PATH
```

Help Troubleshooting

Difficulty bringing up Searchable Help is commonly due to a network problem of some kind. Possible issues include DNS, domain suffixes, or VMWare configurations. A workaround for all of these causes is to use the IP address instead of the hostname of the target Enterprise Manager host computer when you log in to Workstation. For more information, see Knowledge Base article 1673.

PDF-Format Documentation

The same books viewable in online help are available in PDF format for download.

To download PDF-format documentation:

1. From the Help menu, select Help and Documentation.

   The top-level Help window appears.

2. Click Individual Introscope Guides (PDF format).

3. View a book:

   - Click the name of an individual book to view it in a PDF-enabled browser.

   or

   - Right-click the name of an individual book and save it to your computer.

CA Technologies Pre-release Documentation. Authorized Use Only.

Chapter 2: Using the Workstation Console

This chapter describes how to use the Workstation Console.

The Workstation Console displays metric information in dashboards. Dashboards are pre-configured windows that present graphical views of current or historical performance and availability metrics.

When you open the Console, it shows live performance and availability data. You can view historical data by selecting a time range.

This section contains the following topics:

Navigating Among Dashboards in the Console (see page 45)
Preconfigured CA APM Dashboards (see page 50)
Live and Historical Data in the Workstation Console (see page 63)

Navigating Among Dashboards in the Console

You can select Console dashboards in several different ways:

- Dashboard drop-down list
- Forward and backward buttons
- History list
- Home button
- Hyperlinks

Dashboard Drop-down List

You can select dashboards from the drop-down list at the top of the Console page. You can type all or part of the dashboard name, to narrow the selections in the list.

After you have viewed several dashboards, you can navigate among them:

- using forward and back arrows
- using the drop-down list next to the forward arrow and back arrow.
- If you have defined a home dashboard in your user preferences, you can open it by clicking the Home button.

CA Technologies Prerelease Documentation. Authorized Use Only.

Navigate Using Hyperlinks

You can use hyperlinks to navigate between Introscope dashboards and the Investigator:

- Automatic hyperlinks—Introscope automatically links a Data Viewer to the metric grouping it is based upon. The Links menu for the viewer contains a link to the underlying metric grouping definition in the Management Module Editor. Similarly, dashboards that contain Data Viewers based on the same metric grouping are automatically linked, and you can navigate between them using the Links menu.
- Custom hyperlinks—You can define custom links for dashboard items, to link to other dashboards or to web pages. You can define custom links if you have dashboard editing permission.

Note: Some out-of-the-box Console dashboards—for example, EM Capacity—do not automatically contain links to underlying data. Edit these default dashboards or create new dashboards with links. For information about creating and editing custom links, see Creating and managing custom hyperlinks (see page 302).

To see a list of available dashboard links:

1. Right-click a dashboard object.
2. Select Properties > Links.

If no links are available for an object, the Links menu is disabled

To follow dashboard links:

1. Hover your cursor over a dashboard object that has a hyperlink.

   The pointer changes to a hand.
2. Double-click the object to follow the link to its default target.

Creating Dashboard Favorites

To simplify access to dashboards that you use often, you can add them to the Console Favorites menu.

To add a dashboard to your favorites:

1. navigate to the dashboard.
2. Select Favorites > Add to Favorites.

Note: Favorite links are not retained when you rename or delete a favorite dashboard. Update the link, or delete the old link and create a new one.

CA Technologies Prereleased Documentation
Authorized Use Only

To delete a dashboard from favorites:

1. In the Console, select Favorites > Organize Favorites.
2. Select a dashboard.
3. Click Delete.

To edit the list of favorites:

1. In the Console, select Favorites > Organize Favorites.
2. Select a dashboard.
3. Click Edit.

Launching Investigator from Console

If you are viewing live data in the Console and launch the Workstation Investigator from this Console, you can view live data in the Investigator also. However, in the Investigator, the default value for time range is 8 minutes and resolution is 15 seconds. You do not have the option of entering a custom time range and resolution for the live mode in the Investigator.

If you are viewing historical data in the Console and launch the Workstation Investigator from this Console, you can view historical data in the Investigator also for the same time range and resolution that you selected for the historical data in the Console.

Launching Console from Investigator

If you are viewing live data in the Investigator and launch the Workstation Console from this Investigator, you can view live data in the Console also. However, in the Console, the default value for time range is 8 minutes and resolution is 15 seconds. You can enter a custom time range and resolution for the live mode in the Console.

If you are viewing historical data in the Investigator and launch the Workstation Console from this Investigator, you can view historical data in the Console also for the same time range and resolution that you selected for the historical data in the Investigator.

Find More Information from Dashboards

When you want more information about the data presented on dashboards, you can use shortcuts to get more information.

To view more information using dashboards:

- Right-click a graph or an alert, click Links, and navigate to the corresponding alert in the Management Module or another dashboard associated with the graph or alert.

CA Technologies Pre-release Documentation. Authorized Use Only

- Position the cursor over interactive elements in an application triage map dashboard element. Interactive elements on the map include map nodes, connector lines, and alert indicators. See more information about Using the Application Triage Map (see page 87).
- Double-click a metric from the chart displaying the Top N (for example, top 10 or 25 slowest) metric data to view its details in the Investigator.

Filtering by agent with the Console Lens

You use the Console Lens to filter metric data for the agents that are reporting data. In a dashboard that shows data for more than one agent, you can use the Console Lens to view data only for selected agents.

When you apply the Console Lens, that filtering remains in effect until you close the Console window, log out from the Workstation, or use the Clear Lens command.

Applying the Console Lens

To apply the Console Lens:

1. Click the Lens button, or select Dashboard > Lens.

   If the Console is in Live mode, the dialog lists the currently connected agents. If you are viewing a time range of historical data, the dialog box lists agents connected for the selected historical range.

2. In the Select Agent dialog, select a single agent, or select multiple agents (click and drag, or CTRL/click) on which to filter.

   Note: You can begin typing an agent name, hostname, or process name in the Search field. As you type, the agent list filters to match what you type.

3. Click Apply or press Enter.

   The dashboard refreshes to show only data for the selected agent(s). The arrow on the lens changes from light blue to black when a lens is applied.

Unsupported Widgets

Some dashboard widgets do not support the lensing feature:

- Graphs powered by calculators
- Graphs based on a Virtual Agent powered by a simple alert. This includes the Top 10 Connected Agents graph on the Overview dashboard.

CA Technologies Pre-Release Documentation. Authorized Use Only.

Clearing the Console Lens

To clear the Console Lens:

1. Click Lens.
2. Clear the Lens by clicking the Clear button on the Apply Console Lens dialog.

More information about adding Investigator Views to a dashboard

Create and Edit Dashboards (see page 278)

Console Lens and tab views in dashboards

The effect a Console Lens has on an Investigator View in a dashboard depends on the type of tree item with which the view is associated.

If the Investigator item associated with the view is:	and...	then
a domain	a single agent is selected in the lens...	...the item association changes to a single agent selection. If the view doesn't support agent selection, an error message appears.
an agent	a single agent is selected in the lens...	...the item association changes to a single agent selection.
a metric	a single agent is selected in the lens...	...the same metric on the selected agent becomes the current selection. If that metric does not exist an error message appears.
a metric path	a single agent is selected in the lens...	...the same metric path on the selected agent becomes the current selection. If that path doesn't exist, an error message appears.
another item type		an error message appears.

If more than one agent is selected, an error message appears in the tab view.

If the lensed agent is a Virtual Agent, the view shows data for that agent, if it supports that type of selection. You can determine what views are supported for a given item type by selecting an item in the tree, and observing the view tabs that are available.

CA Technologies Prerelease Documentation. Authorized Use Only.

A Virtual Agent is a group of physical agents that are configured to be a single agent, enabling you to see an aggregated view of the metrics reported by several agents.

For information about Virtual Agents, see the *CA APM Configuration and Administration Guide*.

Preconfigured CA APM Dashboards

CA APM is shipped with several Management Modules containing pre-built dashboards. These dashboards provide:

- **Efficient monitoring—**High-level application health and status views of large numbers of applications
- **Rapid notification—**At-a-glance notification of problems in the production application environment
- **Actionable information—**Enables quick identification of what is wrong, what to do, who to call
- **Minimal training—**Pre-defined navigation between high-level and drill-down performance information, reducing the learning curve.
- **Quick resolution—**Operations and Application Support personnel collaborate more effectively to identify and resolve problems

The pre-built dashboards provide an example of how to organize Introscope metrics into a meaningful set of views for Introscope users.

The Enterprise Manager installer places *.jar* files containing these dashboards, with their supporting elements, in the *<EM_Home>/config/modules* directory in a new installation, or in the *<EM_Home>/examples* directory if the installation is an Introscope upgrade.

Sample dashboards

The *SampleManagementModule.jar* file contains:

- the Intro to Introscope dashboard (see page 53)
- the Overview dashboard (see page 54)
- the Problem Analysis dashboard (see page 55)

CA Technologies Pre-Release Documentation. Authorized Use Only

Performance dashboards

The *Collector_1.jar* file contains dashboards which display Collector and MOM performance metrics information. The performance metrics show how fast or slow an Enterprise Manager cluster or a particular Collector or MOM is performing.

- MOM - Infrastructure Overview
- MOM - Infrastructure Performance
- MOM - Cluster Clock Drift
- MOM - Cluster Ping Time
- MOM - Harvest Duration Detail
- MOM - SmartStor Duration Detail
- MOM - Event Handling
- MOM - Collector Query Performance
- MOM - Resources Capacity
- MOM - Collectors Performance
- EM - Collector Performance
- EM - Collector Harvest Duration Detail
- EM - Collector SmartStor Duration Detail
- EM - Collector Event Handling
- EM - Collector Query Performance
- EM - Collector Resources Capacity

See more information about these metrics in Performance Dashboards (see page 58).

Capacity dashboards

The *MOM_Infra_Monitoring_MM.jar* file contains dashboards that monitor capacity:

- MOM - Infrastructure Overview
- MOM - Infrastructure Capacity
- MOM - Metrics Capacity Detail
- MOM - Collectors Capacity
- EM - Collector - Metrics Detail
- EM - Collector - Error Snapshot Events

CA Technologies Prerelease Documentation. Authorized Use Only.

Note: When monitoring an Enterprise Manager cluster, first view the MOM - Infrastructure Overview dashboard.

See more information about these metrics in Capacity Dashboards (see page 62).

Enhanced Dashboards

CA APM includes Management Modules with preconfigured dashboards that enable you to monitor the performance and capacity of an Enterprise Manager cluster consisting of one or more Collectors and a MOM. The Management Modules are:

- Collector_1.jar
- MOM_Infra_Monitoring_MM.jar

To view dashboards in the Workstation Console:

1. Copy Collector_1.jar and MOM_Infra_Monitoring_MM.jar from the <*EM_Home*>\examples directory to the <*EM_Home*>\config\modules directory.
2. Edit the Management Modules contained in these .jar files for your own environment, following the instructions under the heading Configuring Metric Groupings (see page 273).

   Note: When you copy the sample Management Modules Collector_1.jar and MOM_Infra_Monitoring_MM.jar in the <*EM_Home*>\config\modules directory and customize them to edit metrics groupings, you see that the Metrics Expressions pane for some of the metric grouping contains a hard-coded sample metrics expression:

   Delete this sample metrics expression and substitute your own.
3. Configure the Workstation Console to view live data for more than 8 minutes, if necessary. See Viewing live query data in the Console (see page 63) for details.
4. Verify that your applications are instrumented and providing data.
5. Verify that the Enterprise Manager is running.
6. Launch the Workstation by browsing the Start menu.
7. Log in to the Workstation.
8. Navigate to the Workstation Console by clicking Workstation > New Console.

Now you can view the dashboards.

CA Technologies Product Release Documentation
Authorized Use Only

Where the Sample Dashboards are Installed

The sample dashboards described in this chapter are installed when a new Enterprise Manager installation is installed. If you have upgraded from a previous version of CA APM, the old sample dashboards are preserved and the new dashboards are available in the Enterprise Manager's *examples* directory, in the Management Module file named *SampleManagementModule.jar*.

You can Hot Deploy this management module to see the new dashboards in your environment. For more information about the Hot Deploy feature, see the *CA APM Configuration and Administration Guide*.

Users with SAP installations do not see sample dashboards.

Overall Status Indicators on Dashboards

The alert indicators on sample dashboards show the overall state of the environment, and how key performance indicators are affecting the environment:

This indicator	Shows
Overall	How is the overall experience to the application user?
Response Time	How is the response time for the application?
Errors	Are application users experiencing application errors?
Stalls	Is the application experiencing stalls?
CPU	Is the CPU consumption for the application normal?
Thread Pools	Does the application have enough threads available in its thread pool?
JDBC Pools	Does the application have enough JDBC connections in its connection pool?

For more information about the metrics behind these indicators, see the Metrics Reference Appendix (see page 345).

The Sample Intro to Introscope Dashboard

When you open the Sample Management Module you see the "Intro to Introscope" dashboard.

To jump to other dashboards:

- Select another dashboard from the dropdown list at the top of the dashboard, or
- Double-click any of the hyperlinked graphical elements such as the alert indicator.

CA Technologies Pre-Release Documentation Authorized Use Only

The Sample Overview Dashboard

The Overview dashboard is designed for the Application Support team to monitor the key performance indicators of their applications across the entire monitored environment.

Graphs show average response time of monitored applications, their throughput, the CPU utilization, and the connection state of the agents.

Alert indicators appear on each sample dashboard to show the overall state of the environment.

The Overview dashboard includes these graphs:

This graph	Shows
Application Average Response Time and Responses per Interval	The aggregate *Average Response Time* of the monitored applications, and their throughput (*Responses per Interval*). An interval is 15 seconds. If Introscope reports 45 responses per interval for an application, therefore, it translates to a throughput of 3 hits per second.
Backend Average Response Time and Responses per Interval	Average response time and throughput of connected backend systems. Backend systems can be anything that the monitored applications connect to—databases, LDAP servers, and mail servers, for example. Introscope automatically identifies connected systems and monitors their performance. In many cases, poor response time can be directly traced to one of its backend systems.
Key Application Server CPU Utilization	CPU utilization of the .NET and Java processes that CA APM monitors. This graph does not indicate the overall CPU consumption on the machine—it is the CPU consumed by the .NET or Java process itself. Introscope provides data about the CPU consumption of the machine, and you can include them in your custom dashboards.
Connected Agents	Connection state of agents. CA APM reports the state of connected agents as metrics whose value is either 1 or 3: ■ 1 for an agent indicates that agent is connected to the Enterprise Manager. ■ 3 indicates that an agent has disconnected from the Enterprise Manager. The graph shows the top 10 connected agents. Because disconnected agents have a larger value than connected agents, disconnected agents are shown first.

CA Technologies Proprietary & Confidential Information. Authorized Use Only. Documentation.

The Sample Problem Analysis Dashboard

On the Problem Analysis dashboard, overview alert indicators show you the health of the entire environment as you review the details of a particular problem.

The Problem Analysis dashboard also displays graphs that help you locate the cause of a particular problem:

This graph	Shows
Application Average Response Time	The aggregate response time of the monitored applications.
Responses per Interval	The throughput of the monitored applications.
Application Stalls	Shows stalls coming from all components of your application, including backend systems. Stalls are an important metric that can help you determine the cause of many production application problems. Stalls occur when a request has been made of a monitored application, but the application has not responded within thirty seconds. Most stalls in production environments occur because a backend system has stopped responding to an application's requests. Introscope often automatically identifies the backend systems to which the application connects, and monitors those systems for stalls. When Introscope is unable to find a backend system, however, that system remains unmonitored. When an unmonitored backend system stalls, secondary stalls within the application might indicate that a stall is occurring, but Introscope is unable to identify the cause. In this situation, the Top Concurrent Socket Communication graph can help you determine the cause of a problem.
Top Concurrent Socket Communications	Shows results of the Socket Concurrency metric. The two types of socket concurrency metrics are readers and writers. Reader metrics are the number of requests in the application waiting for a backend system to respond with data through a socket. Writer metrics are the number of requests in the application waiting for a backend system to accept data through a socket. If a stall in an application is caused by a backend system that Introscope does not identify, looking at a high level of concurrent socket readers or writers can often identify the offending system.

CA Technologies Proprietary Documentation
Authorized Use Only

Navigation Details

You can drill down through the dashboards in particular sequences to understand specific cluster and Collector performance issues.

To drill down into Collector connection performance:

Drill down from the MOM - Infrastructure Overview dashboard to one of the following dashboards:

- MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Connected
- MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Cluster Clock Drift
- MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Cluster Ping Time

To drill down into MOM performance:

1. Drill down from the MOM - Infrastructure Overview dashboard to MOM - Infrastructure Performance dashboard.
2. Drill down from MOM - Infrastructure Performance dashboard to one of the following dashboards:
   - MOM - Infrastructure Performance > MOM - Harvest Duration Detail
   - MOM - Infrastructure Performance > MOM - SmartStor Duration Detail
   - MOM - Infrastructure Performance > MOM - Event Handling
   - MOM - Infrastructure Performance > MOM - Collector Query Performance
   - MOM - Infrastructure Performance > MOM - Resources Capacity

To drill down into Collector performance:

1. Drill down from one dashboard to the other in the following sequence:
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Performance
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Performance > EM - Collector Performance

CA Technologies Pre-release Documentation. Authorized Use Only

2. Drill down from the EM - Collector Performance dashboard to the following dashboards in the following sequence:
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Performance > EM - Collector Performance > EM - Collector Harvest Duration Detail
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Performance > EM - Collector Performance > EM - Collector SmartStor Duration Detail
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Performance > EM - Collector Performance > EM - Collector Event Handling
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Performance > EM - Collector Performance > EM - Collector Query Performance
   - MOM - Infrastructure Overview > MOM - Infrastructure Performance > MOM - Collectors Performance > EM - Collector Performance > EM - Collector Resources Capacity

You can examine capacity performance by drilling down into specific dashboards in a particular sequence.

To drill down into MOM capacity:

Drill down from the MOM - Infrastructure Overview dashboard to one of the following dashboards in the sequence:

- MOM - Infrastructure Overview > MOM - Infrastructure Capacity
- MOM - Infrastructure Overview > MOM - Infrastructure Capacity > MOM - Metrics Capacity Details

To drill down into Collector capacity:

Drill down from the MOM - Infrastructure Overview dashboard to one of the following dashboards in the sequence:

- MOM - Infrastructure Overview > MOM - Infrastructure Capacity > MOM - Collector Capacity
- MOM - Infrastructure Overview > MOM - Infrastructure Capacity > MOM - Collector Capacity > EM - Collector - Metrics Detail
- MOM - Infrastructure Overview > MOM - Infrastructure Capacity > MOM - Collector Capacity > EM - Collector - Metrics Detail > EM - Collector - Error Snapshot Events

CA Technologies Prerelease Documentation. Authorized Use Only

The following dashboard is available for monitoring the performance and monitoring of Enterprise Manager cluster with Collectors and MOM:

- MOM - Infrastructure Overview (EM-Cluster Monitoring)

This dashboard shows the overall health of Enterprise Manager Cluster Monitoring. It shows performance and capacity alerts. If an alert is red, you navigate to MOM - Infrastructure Performance or MOM - Infrastructure Capacity dashboard.

Performance Dashboards

You can use these dashboards to monitor cluster performance:

MOM - Infrastructure Performance (CA APM Infrastructure Performance)

This dashboard shows the overall health of CA APM Infrastructure Performance. It shows alerts for Collectors Connected, Clock Drift, Ping Time, MOM Performance, Harvest Duration, SmartStor Duration, Number of Event Inserts, Remote Queries, CPU & Memory Resources, and Collectors Performance. If an alert is red, you double-click it and navigate to the relevant dashboard. For example, if the Connectors Connected alert is red, double-click the alert to and see the Collector's connection status.

MOM - Collectors Connected (Collectors Connection Status)

This dashboard shows the overall health of the Collectors Connection Status. It displays an alert for Connection Status for Collector 1.

You can customize this dashboard to display alerts for all Collectors in the cluster. For information about customizing dashboards, see Creating and Editing Dashboards (see page 278).

MOM - Cluster Clock Drift (Collectors Clock Drift)

This dashboard displays the clock skew between Collector's clock and the MOM's clock. Collectors with a clock skew greater than 3 seconds are disconnected by the MOM. The alert (per Collector) is configured with a danger threshold of 3 seconds. This dashboard displays an alert for Clock Skew for Collector 1.

Note: You can customize this dashboard to display alerts for all Collectors in the cluster.

MOM - Cluster Ping Time (Collectors Ping Time)

This dashboard shows the overall health of the Ping Time for all Collectors. It displays an alert for Ping Time for Collector 1.

You can customize this dashboard to display alerts for all Collectors in the cluster.

CA Technologies Prerelease Documentation. Authorized Use Only

MOM - Harvest Duration Detail (MOM - Data Harvesting)

This dashboard shows the overall health of data harvesting. It shows:

- An alert and a graph for Harvest Duration for Harvest Cycle.
- Graphs for Messaging Incoming Threads, Messaging Outgoing Threads, Internal Threads - CPU Time, and Internal Threads - Blocked Time for Threads Performance.
- Graphs for Metrics Evaluated by Alerts, Metrics Evaluated by Calculators, Number of Applications, and Workstations Connected for Metrics Data Subscriber.

MOM - SmartStor Duration Detail (MOM - SmartStor Data Processing)

This dashboard shows the health of SmartStor data processing. It shows:

- An alert and a graph for SmartStor Duration for Processing Cycle.
- Alerts for SmartStor Query Duration, SmartStor Query Per Interval, Data Points Retrieved From Disk Per Interval, Metrics Retrieved From Disk Per Interval, and Metadata Write Duration for SmartStor Performance.
- An alert for Data Points Retrieved From Disk Per Interval for SmartStor Performance.
- Graphs for Number of Live Metrics and Number of Historical Metrics for Number of Metrics.

MOM - Event Handling (MOM - Event Handling)

This dashboard shows the overall health of event handling. It shows:

- Graphs for Numbers of Events Processed, Numbers of Inserts Per Interval, and Number of Dropped Per Interval for Event Storage.
- An alert for Numbers of Events Processed for Event Storage.
- Graphs for Query Duration Per Interval, Number of Queries Per Interval, Insert Duration Per Interval, and Index Insertion Duration Per Interval for Event Query Performance.
- Graph for Number of Events in Database for Total Events.
- Graph for Number of Active Sessions for Active Sessions.

MOM - Collector Query Performance (MOM - Collector Query Performance)

This dashboard shows the overall health of query performance for all Collectors in a cluster. It shows:

- Graphs for Sync Query Duration and Number of Sync Queries Per Interval for Sync Query Performance.
- Graphs for Async Query Duration, Number of Async Queries Per Interval, and Number of Sync Queries by CLW Queries Per Interval for Async Query Performance.

CA Technologies Prerelease Documentation
Authorized Use Only

- Graphs for Data Points Returned Per Interval and Metrics Returned Per Interval for Client Returned Data.
- An alert for Data Points Returned Per Interval for Client Returned Data.

MOM - Resources Capacity (MOM - Resource Capacity)

This dashboard shows the overall health of MOM Resource Capacity. It shows alerts for CPU Usage, GC Duration, and Free Disk Space.

MOM - Collectors Performance (Collectors Performance)

This dashboard shows the overall health of the Collectors Performance. It displays an alert for Collector 1. If an alert is red, you double-click it to open the EM - Collector Performance dashboard to find out the cause of Collector's performance degradation.

Note: You can customize this dashboard to display alerts for all Collectors in the cluster.

EM - Collector Performance (EM - Collector Performance)

This dashboard shows the overall health and status of Collector Performance. It shows alerts for Harvest Duration, SmartStor Duration, Number of Event Inserts, Query Performance, and CPU & Memory Resources.

If an alert is red in this dashboard and you want to know the root cause, you can navigate from this dashboard to the EM - Collector Harvest Duration Detail, EM - Collector SmartStor Duration Detail, EM - Collector Event Handling, EM - Collector Query Performance, and EM - Collector Resources Capacity dashboards.

EM - Collector Harvest Duration Detail (EM - Data Harvesting)

This dashboard shows the health of data harvesting. It shows:

- An alert and a graph for Harvest Duration for Harvest Cycle.
- Graphs for Messaging Incoming Threads, Messaging Outgoing Threads, Internal Threads - CPU Time, Internal Threads - Blocked Time for Threads Performance.
- Graphs for Metrics vs Handled Metrics, Number of Virtual Metrics, Workstations Connected, and Number of Applications for Metric Data Subscribers.

EM - Collector SmartStor Duration Detail (EM - SmartStor Data Processing)

This dashboard shows the health of SmartStor Data Processing. It shows:

- An alert and a graph for SmartStor Duration for Processing Cycle.
- Alerts for SmartStor Query Duration, SmartStor Query Per Interval, Data Points Retrieved From Disk Per Interval, Metrics Retrieved From Disk Per Interval, and Metadata Write Duration for SmartStor Performance.

CA Technologies Confidential Document. Authorized use only.

- An alert for Data Points Retrieved From Disk Per Interval for SmartStor Performance.
- Graphs for Number of Live Metrics and Number of Historical Metrics for Number of Metrics.

EM - Collector Event Handling (EM - Event Handling)

This dashboard shows the overall health of event handling. It shows:

- Graphs for Numbers of Events Processed, Numbers of Inserts Per Interval, and Number of Dropped Per Interval for Event Storage.
- An alert for Numbers of Events Processed for Event Storage.
- Graphs for Query Duration Per Interval, Number of Queries Per Interval, Insert Duration Per Interval, and Index Insertion Duration Per Interval for Event Query Performance.
- Graph for Number of Events in Database for Total Events.
- Graph for Number of Active Sessions for Active Sessions.

EM - Collector Query Performance (EM - Query Performance)

This dashboard shows the overall health and status of Enterprise Manager Query Performance. It shows:

- Graphs for Cache Queries Duration and Cache Queries Per Interval for Cache Query Performance.
- An alert for Cache Queries Duration for Cache Query Performance.
- Graphs for SmartStor Queries Duration and SmartStor Queries Per Interval for Historical Query Performance.
- An alert for SmartStor Queries Duration for Historical Query Performance.
- Graphs for Data Points Returned Per Interval and Metrics Returned Per Interval for Client Returned Data.
- An alert for Data Points Returned Per Interval. The threshold for this alert is the corresponding metric clamp value.

EM - Collector Resources Capacity (EM - Resource Capacity)

This dashboard shows the overall health and status of Enterprise Manager Resource Capacity. It shows alerts for CPU Usage, GC Duration, and Free Disk Space.

CA Technologies Prerelease Documentation. Authorized Use Only

Capacity Dashboards

The following dashboards are available for monitoring cluster capacity:

MOM - Infrastructure Capacity (CA APM Infrastructure Capacity)

This dashboard shows the overall health of CA APM Infrastructure Capacity. It shows alerts for MOM Capacity and Collectors Capacity. If an alert is red, double-click it to open the relevant dashboard. For example, if MOM - Capacity alert is red, double-click it to navigate to MOM - Metrics Capacity Detail dashboard to find out which agent or agents are responsible for the capacity reduction.

MOM - Metrics Capacity Detail (MOM - Capacity)

This dashboard shows the overall health of MOM Capacity. It displays:

- Alerts and graphs for Number of Agents, Number of Live Metrics, Number of Historical Metrics, and Number of Events Processing for MOM Metrics Stats.
- Graphs for Number of Collector Metrics and Collector Metrics Received Per Interval for Connected Collector Metrics.

MOM - Collectors Capacity (Collectors Capacity)

This dashboard shows the overall health of the Collectors capacity. It displays alert for Collector 1. If an alert for a Collector is red, double-click it to navigate to EM - Collector - Metrics Detail to find out which agent or agents are responsible for capacity reduction.

Note: You can customize this dashboard to display alerts for all Collectors in the cluster. For information about customizing dashboards, see Creating and Editing Dashboards (see page 278).

EM - Collector - Metrics Detail (EM - Collector Capacity)

This dashboard shows the overall health and status of Collector capacity. You navigate from this dashboard to the EM - Collector - Error Snapshot Events dashboard by clicking an alert. For example, if an alert is red, click the alert to determine which agent or agents have exceeded metric and event clamps.

This dashboard displays:

- Alerts and graphs for Number of Live Metrics and Number of Historical Metrics for Collector Metrics Stats.
- Graphs for Number of Agents and Agent Connection Status for Collector Agent Stats.
- An alert for Number of Agents for Collector Agent Stats.
- An alert and a graph for Number of Events Processed.
- An alert for Individual Agent Stats for Connected Agents.

CA Technologies Pre-release Documentation
Authorized Use Only

EM - Collector - Error Snapshot Events (EM - Connected Agents Stats)

This dashboard shows the event and metric load on the EM in Top N charts (by agents). You navigate to this dashboard from EM - Collector - Metrics Detail to see which agent or agents have exceeded the metric and event clamps.

This dashboard displays:

- Graphs for Transaction Tracing Events Per Interval, Transaction Tracing Exceeding Limit, Errors Per Interval, and Errors Exceeding Limit for agent events and Numbers of Metrics by Agents, Agents Exceeding Metrics Limit, and Historical Metrics by Agents for Agent Metrics.
- Alerts for Transaction Tracing Exceeding Limit, Errors Exceeding Limit, and Agents Exceeding Metrics Limit.

Customizing Capacity Alerts

You can customize the alerts in capacity dashboards. See Monitoring Performance with Alerts (see page 306).

Live and Historical Data in the Workstation Console

You can view live data in the Console, or select a range of time to view historical data. The default view of data is Live.

You can see whether the Workstation is in Live mode by looking at the drop-down next to the label *Time range.*

Viewing Live Query Data in the Workstation Console

To view live query data and historical data in the Workstation Console with default properties:

- Click Live to enable it and view live data.

The default value for time range is 8 minutes and resolution is 15 seconds. You do not have the option of entering a custom time range and resolution for the live mode in the Console.

Note: Click Live to disable it and select a time range and resolution from the drop-down list to view historical data. You can also enter a custom time range.

CA Technologies Proprietary/Confidential Information. Authorized Use Only.

To view live query data and historical data in the Workstation Console:

1. To view live query data for time range greater than 8 minutes, edit the introscope.enterprisemanager.workstation.extendedLiveQuery property in the IntroscopeEnterpriseManager.properties file in the <*EM_Home*>\config directory as follows:

   introscope.enterprisemanager.workstation.extendedLiveQuery=true

   When this property is set to true, you can use the Time range and Resolution drop-down lists in the Workstation dashboard in the Live mode. This allows you to enter custom time range and resolution for the live mode instead of the default time range of 8 minutes and resolution of 15 seconds.

   You can reset the time range for a greater period than the default time range of 8 minutes.

Note: The maximum time range for which you can view live data is 30 days. If you enter a time range greater than 30 days, the Time Range is set to 8 minutes by default. The number of data points displayed in the dashboard is equal to (Time Range/Resolution). If (Time Range/Resolution) is less than 2, the resolution is set to 15 seconds by default.

Important: Setting the time range to greater than 8 minutes may impact the performance of Enterprise Manager due to the disk I/O operations needed to fetch data from SmartStor.

Enable and Disable Live Mode

In the Workstation Console, live mode is enabled by default. You can enable or disable live mode by clicking the Live button.

Note: When the Console is in the live mode and the resolution set in the Data Options widget of the dashboard is 15 seconds, the resolution shown in the Console (toolbar) is used to display the live data. When the Console is in the live mode and the resolution set in the Data Options widget of the dashboard is a number greater than 15 seconds, the resolution from the widget is used to display the live data. For more information about Data Options, see Adding data to data viewer using the data options dialog (see page 289).

Viewing Historical Data

To view historical data, you select a time range. When you select a time range, Introscope immediately shows the data for that range, sets the end time to the current time, and bases the duration on your time range selection.

To switch from live to historical data:

- Click the Live button.

CA Technologies Prerelease Documentation Authorized Use Only

With live mode disabled, you can select a time range and resolution from the drop-down list, or enter a custom time range, and view historical data.

The time range controls can help you identify the time a problem occurred. For example, you think the problem occurred within the last hour, so you set the time range to an hour and look at the data from the current time backward. If you don't see the problem within that hour range, you can use the controls to move backward or forward to locate the time the problem occurred.

To view historical data:

1. Select the metric or dashboard for which you want to see historical data.
2. Select a time range for the historical view from the Time Range drop-down menu.

   Introscope shows the data for that range, using the duration that you selected from the Time Range drop-down menu and setting the end time to the current time.

   Note: If your historical time range includes a year, a four-digit year is required.

   For example, suppose you select a time range at 4:06:45, with a duration of 8 minutes—the end time for the range is therefore set to 4:06:45, and the start time is 3:59:30.

   Note: When you use the time-range control to view historical data, the range you select is applied to other metrics or dashboards in the same window, and to any new windows that you open.
3. Now you can select a resolution to adjust the granularity of the view, by increasing or decreasing the number of data points that appear.

   Each pre-defined time range is associated with a default resolution. You typically do not need to change this setting. Changing the resolution is useful when you want to see a greater level of detail or granularity in the data than appears by default.

   From here, you can:

   - Select a predefined time range from the drop-down list, or
   - Enter a value into the Resolution field. Enter numeric values, followed by the duration—seconds, minutes, hours, or days. For example, "90 Seconds".
4. After you select a time range you can adjust it, using the time range controls.

Alerts in historical mode do not reflect historical alert state

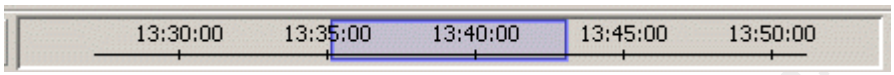
Alert values are not captured in any database, so if a dashboard in historical mode displays alerts, those alerts do *not* reflect the historical state. If data for the alerts is being reported in the present time, the alerts will reflect live, not historical values.

CA Technologies Proprietary. Use pursuant to Company instructions. Authorized Use Only. Confidential Information.

Time Range Controls

You can use time range controls to scroll in increments based on the time range you selected.

Slider



Drag the slider on the time bar to change the time range.

Arrows



Click the arrows to move backward and forward in time.

The single arrows move backward or forward in small increments; the double arrows move backward or forward in time increments that are about equal to the time of the selected time range.

Reset icon



Click the Reset icon to reset the end time of the range to the current time.

Lock Icon



Clicking the lock icon maintains your selected resolution as you select different time ranges by zooming in on data.

Defining a custom time range

To define a custom time range to view historical data:

1. Select the metric or dashboard for which you want to see historical data.
2. Select Custom Range from the Time Range drop-down menu.

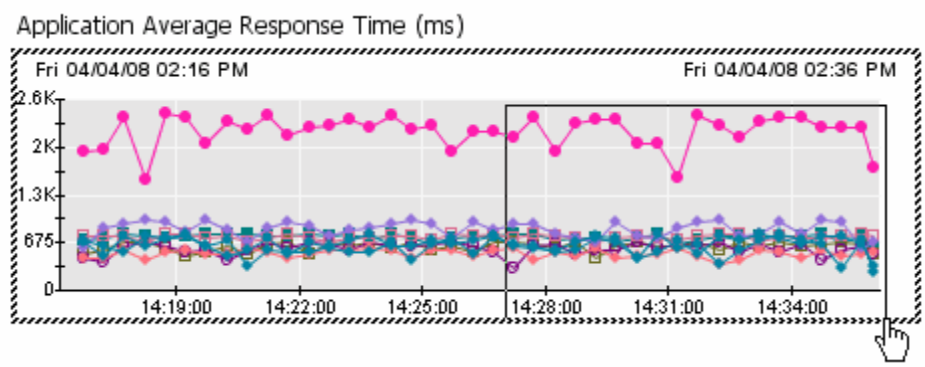
   The Custom Range window opens, showing the current date (Today) highlighted with an outline.

3. Select dates:
   a. Use the calendar controls to select the start and end dates and times.
   b. Use the menu controls at the top of the calendar to select the month and year, choose the date on the calendar, and type in the time in the time field at the bottom of the calendar.
   c. Click OK.

   Workstation shows the data for the custom range.

Zooming into historical data in graphs

When you view historical data in a graph, you can zoom in on data by clicking the mouse pointer on a graph position and dragging, to specify the time range.



Introscope refreshes the data in the viewer based on the new query, and the time range in the viewer shows the new range.

The global time range in the window and the Time Range control do not change automatically when you zoom in on data. For example, if you zoom in on a ten-minute period on a graph with the Time Range set to 1 hour, the graph shows the ten-minute period but the control remains at 1 hour, and the time bar still shows the hour range.

You can override the default zoom actions in these ways:

- Set the global time range and the Time Range control to match the zoomed view: select Viewer > Set Time Range From Zoomed Range, or click the Set Time Range from Zoomed Range icon.
- Lock your selected resolution by clicking the Lock icon.
- Hold down the shift key while you zoom, to constrain zooming to the time axis.

Chapter 3: Using the Workstation Investigator

This chapter describes how to use the Workstation Investigator to view application data.

This section contains the following topics:

High-level Views in the Investigator (see page 69)
The Triage Map Tab (see page 75)
Using the Application Triage Map (see page 87)
The Metric Browser Tab (see page 113)
The APM Status Console (see page 151)
Viewing CA APM Customer Experience Metrics in the Workstation (see page 157)

High-level Views in the Investigator

The Investigator has two top-level tabbed views:

- The triage map tab shows the following views:
  - In the By Frontend view, an application-centric view.
  - In the By Business Service view, a business service/transaction-centric view
  - If you have a Catalyst integration, a locations-centric view.
- The metric browser tab shows an agent-centric view, with detailed metrics on an individual location.

General Investigator Features

Several features are active in the Investigator whether you are looking at the triage map tab or the metric browser tab.

CA Technologies Prerelease Documentation. Authorized Use Only.

Navigation tips

To open an Investigator:

- Select Workstation > New Investigator.

To navigate forward and back:

- Click the Forward or Back arrow buttons in the upper right corner of the Investigator to move forward or backward among previously viewed hierarchical tree items.
- Select from the drop-down lists next to the Forward or Back buttons in the upper right corner of the Investigator.

Investigator panes

The Investigator is generally displayed with two panes:

- A tree hierarchy in a narrow pane on the left side.
- A large viewer pane on the right side.
  - The contents of the viewer pane vary, depending on the type of the item selected in the hierarchical tree.
  - The viewer pane is organized by one or more tabs, each of which displays a different view.

Metric graphs are the most common way to view metrics, though not the only way. For metrics, a view of the metric data appears. Each metric type has a default display in the viewer pane.

Tooltips

Tooltips identify metric paths and values in the hierarchical trees and viewer panes found in both the triage map tab and the metric browser tab

Tooltips in the triage map tab

You access tooltips in the triage map tab by hovering over elements in the application triage map, such as:

- Lines representing connections between map elements
- Threshold lines, if they appear
- Alert indicators, if they appear
- Rectangles representing physical hosts and virtual machines, in some configurations
- The trapezoid shape labeled Resources
- Table cells in the Locations table

CA Technologies Prerelease Documentation
Authorized Use Only.

For more information, see:

- Tooltips in the Frontend View of the application triage map (see page 91)
- Tooltips in the Business Service application triage map (see page 91)

Tooltips in the metric browser tab

You access tooltips in the metric browser tab by hovering over the metric name in the legend area of a Data Viewer.

Workstation displays a variety of information in the tooltip, depending on what Workstation element you are hovering over. This may include the fully qualified metric name, its value and its minimum and maximum values, a count of how many data points were reported in the selected time slice, a timestamp of data value nearest the cursor, or a comparison note -- for example, “Value Too High” when the metric value exceeds a defined threshold.

NOTE: Tooltips are no longer available from nodes in either the triage map tree or the metric browser tree.

Table Views of Data

Many different views include a table on the bottom of the viewer pane. The data contained in this table varies depending on what element you select in either the tree or the viewer. For example:

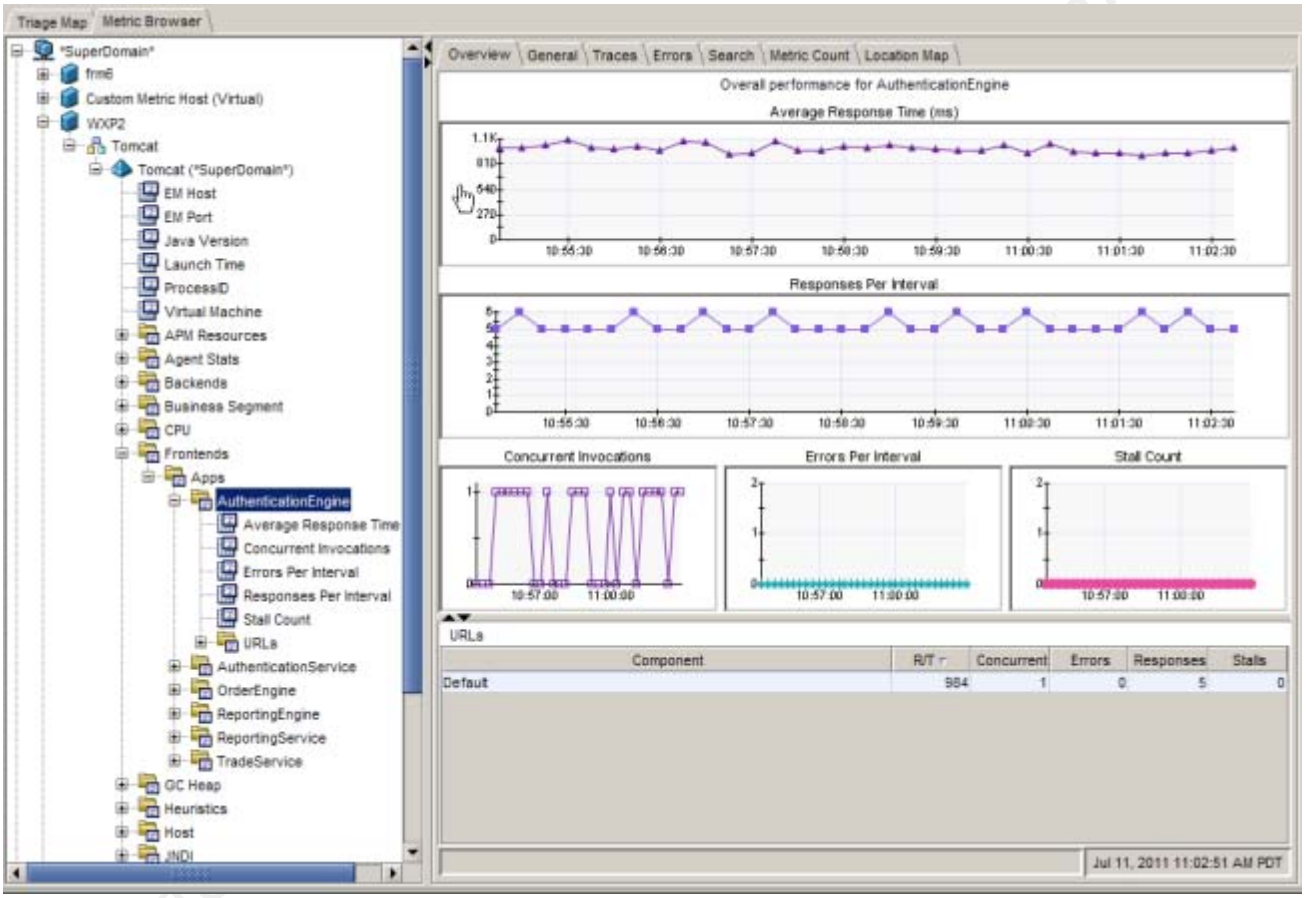
- The illustration in Agent-Centric View (see page 72) shows a table view of the same metric data displayed in graph format above.
- The illustration in Frontend View of the Application Triage Map (see page 90) shows the Locations table, which lists the agent name, the hostname for the selected frontend, and metric data for the frontend in that location.
- The illustration in Resources Element (see page 93) shows resource metrics for the selected frontend resource.

Note: Data shown in the table changes depending on what Workstation element you have selected.

CA Technologies Pre-Release Documentation.
Authorized Use Only.

Agent-Centric View

The metric browser tab allows you to browse a comprehensive list of metrics being reported by a single agent. (An agent is a piece of software installed on a host where an application is deployed; it collects application and environmental metrics and relays them to the Enterprise Manager.) Each application whose data is being reported by an agent appears in a hierarchical tree under a node named Frontends, as shown in the following illustration.

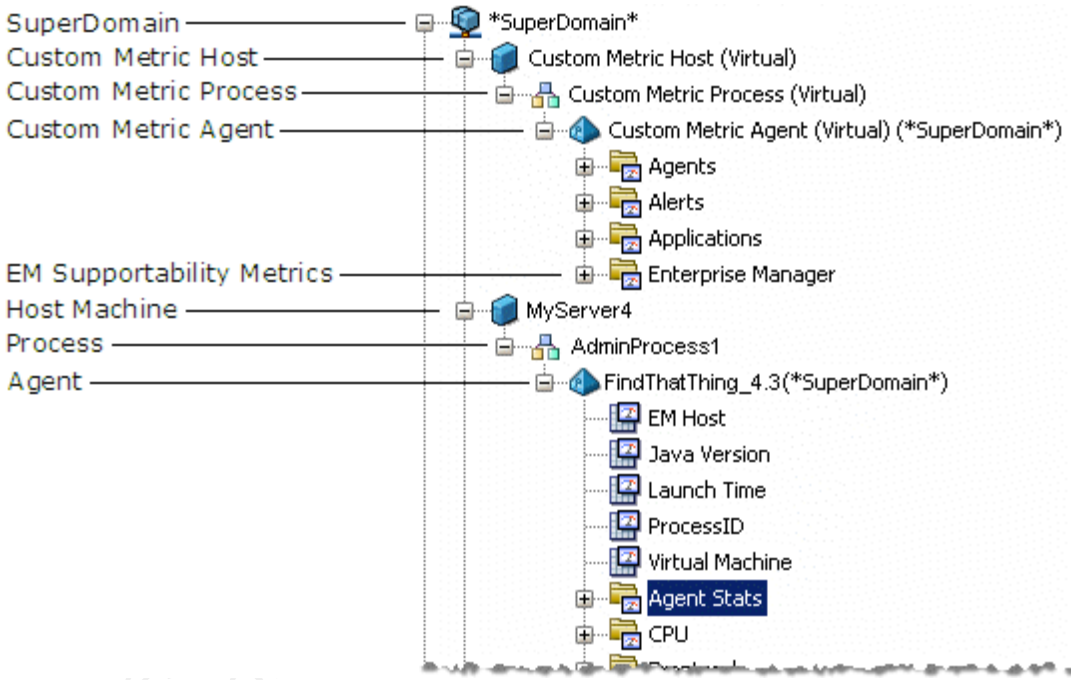


The agent-centric view of the Investigator contains these sections:

- The *agent-centric tree* on the left provides information about each host and application managed by the Enterprise Manager. The metrics that appear in the agent-centric tree are a function of the resources your applications use and the data that your Introscope agents are configured to report.
- The *Viewer pane* on the right presents details, often graphical, for the resource or metric in the tree. You can select View tabs to open different views of data. The tabs that are available vary, depending on the item selected in the tree. For some views, options might be available in the bottom section of the Viewer pane to control the data displayed in the Viewer.

- A table at the bottom of the viewer pane which displays data in a tabular format. The data displayed in the table depends on what you select in the tree or viewer pane.

This illustration shows the agent-centric tree in a Java environment, as seen by a user with read or write permission to the SuperDomain. In this example, the SuperDomain contains no domains, and two agents.



Super Domain node

The SuperDomain node contains metrics for all agents that report to the Enterprise Manager to which the Workstation is connected. Metrics are organized in a Host|Process|Agent hierarchy.

The nodes immediately under the SuperDomain node are virtual and physical hosts.

- **Custom Metric Host (Virtual)—**This node does not correspond to a physical host machine. It is a virtual host that contains metrics that are not reported by a specific, individual agent. For example, if you have configured calculators that create custom metrics, or have configured aggregated agents, they typically appear under the Custom Metric Host.
- **Hosts—**One node for each machine that hosts an agent. Each host node contains a process node for the instance of the application being monitored, which in turn contains an agent node. The agent node contains nodes that correspond to application and system resources, which contain metrics.

CA Technologies Author

Note: The application resources that appear in the agent node differ based on whether the agent type is Java or .NET.

The SuperDomain is that which includes all user-defined domains and agents. The Enterprise Manager administrator can set up the EM to display child domains with separate permissions.

The metrics that appear in the agent-centric tree are a function of the PBDs (ProbeBuilder Directives) used to instrument the application, and the run-time activity of the application itself.

A metric only appears in the tree when the agent starts reporting it. The metric remains visible in the tree, even if the agent stops reporting it.

Note: Metrics might have the same name and appear twice in the Investigator, if the metrics have different metric types. As with all metrics, inactive metrics in this situation are grayed out.

Tools to Monitor Enterprise Manager Health

Supportability metrics

Supportability metrics give information about the state of the Enterprise Manager and the computer it runs on. You can view them under the path SuperDomain|Custom Metric Host|Custom Metric Agent|Enterprise Manager.

The *CA APM Sizing and Performance Guide* contains extensive information about the supportability metrics.

CA APM Status Console

The CA APM Status Console displays graphical and table views of Enterprise Managers, whether stand-alone, in clusters, or in a cross-cluster configuration. See the topic The CA APM Status Console (see page 151).

Domains node

If the agents that report to the Enterprise Manager are organized into domains, the agent-centric tree domain node contains sub-nodes for each domain. Each domain node is structured in the same Host|Process|Agent hierarchy as the SuperDomain, and might also contain a Custom Metric Agent for custom metrics.

CA Technologies Prerelease Documentation
Authorized Use Only.

Application-Centric and Business-Centric Views

With the triage map tab active, the pane on the left displays a hierarchical tree divided into these top-level nodes:

- **By Frontend—**Displays an application-centric view of your applications.
- **By Business Service—**Displays a business-centric view of a business service/process/transaction.

The graphical display in the triage map tab is named the application triage map.

Note: If the Enterprise Manager which the Workstation is logged in to has been configured as a collector on a cluster, it will not display the triage map tab. To use the triage map tab tools on a clustered application, log in to the MOM Enterprise Manager.

For more information about the spplication triage map display, see Navigating in the triage map tab (see page 75).

The Triage Map Tab

The triage map tab in Investigator displays the application-centric view.

On the left side in the pane displaying a hierarchical tree, Workstation displays two high-level nodes:

- By Frontend
- By Business Service

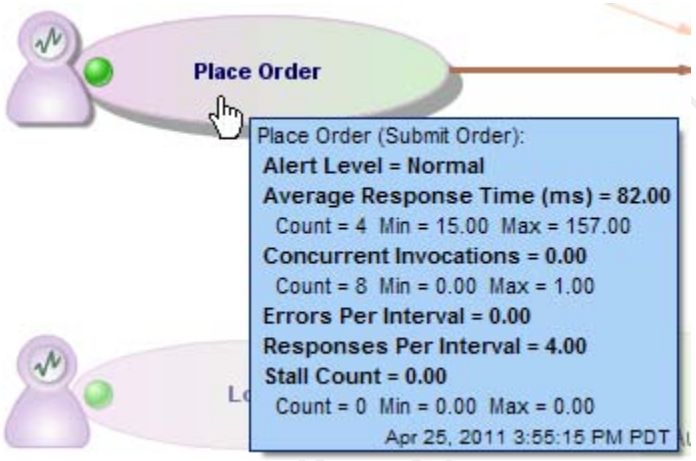
By default, you can see:

- A visual display of an auto-discovered application's components and their dependencies. You view this display in the application triage map view.
- A list of physical infrastructure components hosting the auto-discovered application. You view this list at the bottom of the application triage map view after selecting one of the application nodes in the map.
- Aggregated health metrics for the application. You view these metrics by selecting the Health node under each application listed under By Frontend.

CA Technologies Prerelease Documentation. Authorized Use Only.

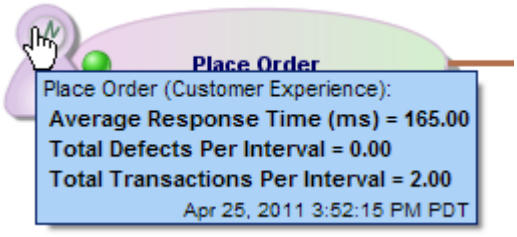
If Business Services have been defined using the Business Service Definition interface in CA APM Customer Experience, you can also see:

- A visual display of the logical dependencies of a Business Transaction (BT). You view this display in the application triage map view when the BT is selected, or when your hover the cursor over a BT oval:



- A table display of the physical locations where application components are executing the BT. You view this list at the bottom of the application triage map view when the BT oval is selected; see List of Physical Locations (see page 86).
- Aggregated health metrics for Business Transaction Components (BTCs). You view these metrics by selecting the node for a BTC. See Business Service and Business Process metrics (see page 95)

If you have a TIM (Transaction Impact Monitor) deployed to monitor web application customer experience metrics and defects, the map displays a Customer Experience (CE) icon next to the BTC oval:



You select, or hover your cursor over, the CE icon to see more information about:

- customer experience metrics (see page 373)
- alert status for customer experience metrics, if an alert has been configured for them.

Permissions

What you see in the triage map tab also depends on your domain and application level permissions. See How user permissions affect what you can view (see page 130).

Navigation in the By Frontend Node

The By Frontend node of the triage map tab allows you to browse:

- Frontends
- Health metrics for frontends
- Metrics for backend calls to each frontend

To see what the application triage map display of these elements looks like, and to understand its various parts, see By Frontend tree and metrics (see page 88) and By Frontend application triage map (see page 90).

Frontends

A frontend is essentially an instance where an application makes socket-client connections to other elements. In the context of the application triage map, these connections are known as backend calls.

As viewed in the By Frontend node of the triage map tab tree, a frontend may represent:

- An application deployed as a .war (web application archive) file. The name displayed is the one configured using the name tag in the .xml file contained in the .war archive, or if there is no name configured, the name of the .war file itself.
- An application using transactions which make socket-client connections using non-SSL, SSL or NIO. The components are named for the socket endpoints.
- An application using EJB connections to backends.
- An application using web services connections, if CA APM for SOA has been configured to record data on such connections.

Note: Users will only see application components for which they have permissions. See How user permissions affect what you can view *(see page 130).*

CA Technologies Proprietary Documentation. Authorized Use Only.

Frontend Sub-Nodes

Each application has two sub-nodes:

- Health metrics—Aggregated metrics across the physical locations:
  - where the selected application is deployed, and
  - where Introscope agents are configured to report metrics for the application.

  Note: To appear in the application triage map, applications must be configured using version 9.0 and later Introscope agents.
- Backend calls metrics—Metrics for calls to other elements supporting the selected application. The components to which these calls are made may include:
  - another frontend.
  - backend systems such as a database.
  - web services called by applications, if you have CA APM for SOA extension installed and configured.
  - unknown elements, which appear as a yellowish shape with a question mark superimposed.

  Metrics on backend calls are taken from measurements on the socket connection to these components.

  See Backend Calls metrics (see page 89) for a list of the available metrics.

When you select a frontend under the By Frontend node in the Map tree, Introscope displays a visual application triage map display of the frontend and its dependencies. Viewing this map, you can:

- Observe alert indicators for the frontend. Frontend dependencies which are themselves frontends may also display alert indicators.
- Hover your cursor over frontends and backend calls to see metrics.
- Right-click map nodes to jump to other frontends or to displays of metrics providing an overview of application health.
- View locations where agents for frontends and dependencies are installed, and jump from the list of locations to metrics in the metric browser tab.

For more information about the By Frontend tree and the application triage map, see Navigation in the By Frontend node (see page 77).

CA Technologies Prerelease Documentation Authorized Use Only.

Navigation in the By Business Service Node

Under the By Business Service node, you see business metrics arranged into the following hierarchy:

- Business Services, which can be any high-level business services carried out by an Introscope-enabled application.
- Business Transactions, defined as individual query-and-response transactions that are children of a business service.
- Business Transaction Components, each of which is the equivalent of a single end-user click or request, and which are children of a business transaction.

You define this hierarchy using the Business Definition interface, as documented in the *CA APM Transaction Definition Guide.*

Depending on what level in this hierarchy you select, you can view:

- A tree hierarchy display of parent-child relationships between Business Services, Business Transactions (BTs), and Business Transaction Components (BTCs).
- A visual display of Business Transactions and their dependencies.
- When a TIM has been deployed to monitor web application customer experience metrics and defects, Customer Experience metrics (see page 373) aggregated across TIMs for that Business Transaction (BT) appear in a Customer Experience node under the Business Transaction node.
- Health metrics for each Business Transaction Component (BTC).
- A list of the physical locations where agents are reporting metrics for the BTC.

Note: Users will only see applications for which they have permissions. See How user permissions affect what you can view (see page 130).

About alert indicators

Colored alert indicators show the aggregated status of the metric or element they decorate. See About Alerts and Alert Indicators (see page 37).

Other Application Triage Map Display Elements

The application triage map display contains several elements to provide more information about applications or business services and their dependencies.

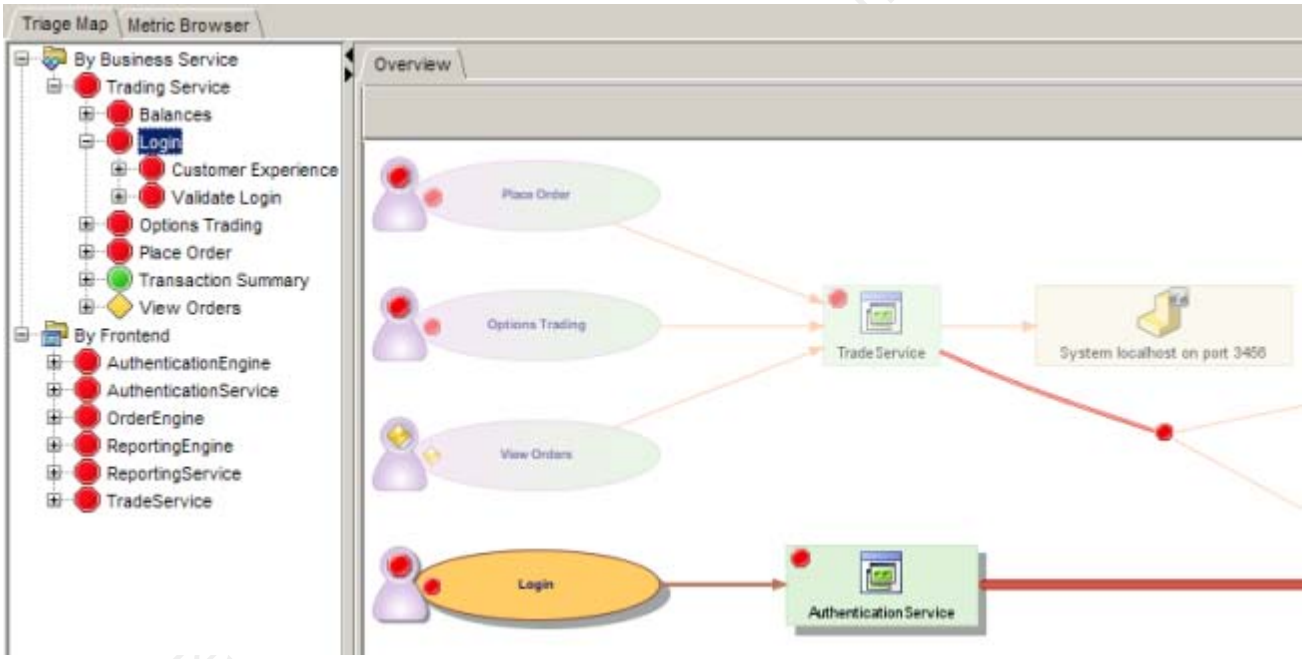
CA Technologies Pre-Release Documentation. Authorized Use Only

Context

Context, as a visual element in graphical displays, refers to the relationship between an element and the application selected in the map tree. It determines several aspects of the appearance of map nodes and connections, as explained further in this table. Context is signaled by shading to the right of and below a map node to indicate that it is either:

- selected in the map tree, or
- a participant in, or actual component of the frontend/business transaction which is selected in the map tree.

For example, notice the difference in appearance between the selected BT oval Login and its dependencies and the other elements in this map. Login and its dependencies are displayed with a shadow:



However, the shadow will generally not appear if the entire map consists of elements that correspond to the primary context. For instance, if a Business Service (BS) is selected in the tree, the map will have no shadow, because everything in the map is associated with that BS. Similarly, a Frontend map will have a shadow only if secondary dependencies are being displayed, and even then, only if they are not involved in transactions originating with the frontend that is selected in the tree.

Connection arrow

The connection arrow between map nodes has four different states:

- **Live—**A brightly colored arrow indicates a live connection.
- **Primary context—**when the map shows both primary and secondary context, and a connection is involved in the primary context, the arrow is a medium brown color (intermediate between the default color and the selected color).
- **Selected—**Select the connection itself by clicking on the arrow, which becomes a darker color, as shown in the illustration. Selecting the connection also highlights the frontend and backend(s) linked by the connection.

  When a backend connection metrics icon appears, you can select the connection arrow itself.
- **Aged—**An uncolored, dimmed arrow indicates an aged connection (see Aged Elements and Connections (see page 83)).

Backend Connection Metrics Icon

The backend connection (or backend call) metrics icon indicates the presence of health metrics on the connection between a frontend and one or more dependencies.

The icon varies according to:

- whether an alert has been configured on the backend connection.
- its position on the map.

Backend connections not configured with alerts

When a backend connection has been configured with an alert, an alert indicator appears in place of the backend connection metrics icon. But when no alert has been configured, the icon takes its default form: a light green disc with a light-colored zigzag emblem:



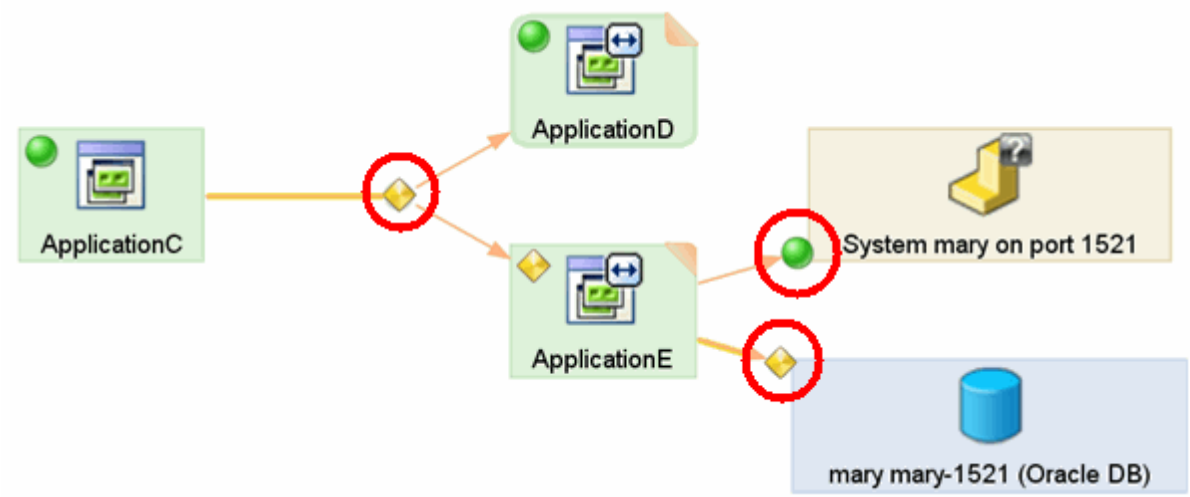
Position of the backend connection icon/alert indicator

The backend connection icon (or the alert indicator which replaces it, if an alert has been configured for that connection) usually appears at the edge of the map element representing the backend itself. In the illustration below, the backend call from ApplicationE to the database mary mary-1521 is displayed with a yellow diamond indicator to show that the alert configured on the backend connection is in Caution state.

CA Technologies Prerelease Documentation. Authorized Use Only.

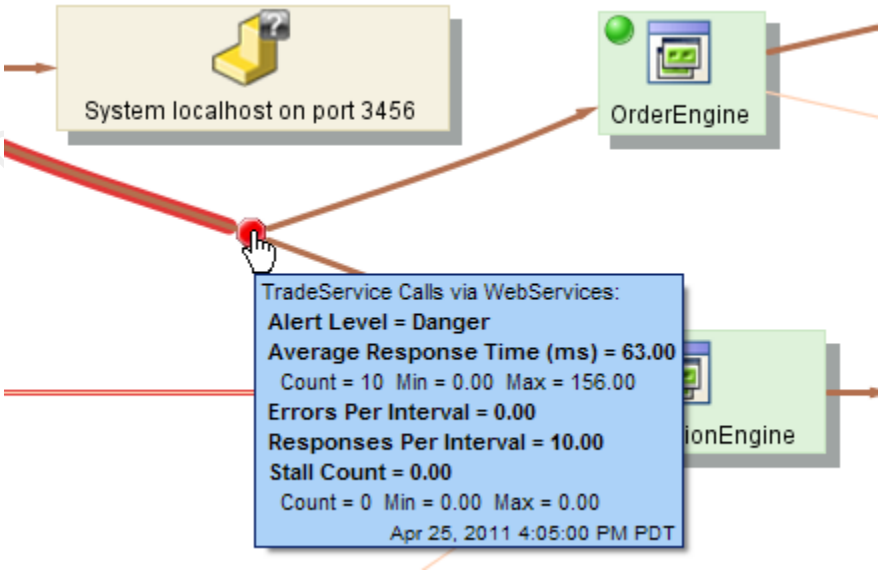
Forked connection

The icon can also appear at the beginning of connection arrows to two or more dependencies. In the illustration below, the icon appears on the connection between ApplicationC and its two dependencies. This forked connection path is displayed only for a web services backend.



Tooltips on backend metrics connection icon

You can hover your cursor over a backend connection metrics icon (or the alert indicator which replaces it, if an alert has been configured for that connection) and evoke a tooltip with metrics for that backend connection.



Dimmed Elements

A map node is still displayed when you do not have permission to view metrics for the element, but:

- It appears faded or dimmed.
- It is displayed in color, not shades of gray.
- No alert indicator will be displayed.

For more information, see How user permissions affect what you can view (see page 130).

Aged elements and connections

A map node is aged when the node no longer participates in the currently displayed map. This can happen when, for example, the name of a database has been changed; the old database will appear as aged, and the new database will appear live. (When the map is in live mode, this period is 24 hours ending now. In historical mode, this period can vary; see Using historical mode in the application triage map (see page 110).)

Generally, an aged element has:

- a gray color
- is dimmed
- no metrics icon
- no alert indicator

However, various conditions can affect how an aged element appears. If you change the primary focus of the map display by selecting a different node in the tree, the state of a particular map node may change because aging is relative to the context of the display.

Dog-ear

You can reveal or "unroll" the dependencies of a map element when it has a "dog-eared" upper right corner, as shown in the illustration below.



CA Technologies Prerelease Documentation. Authorized Use Only

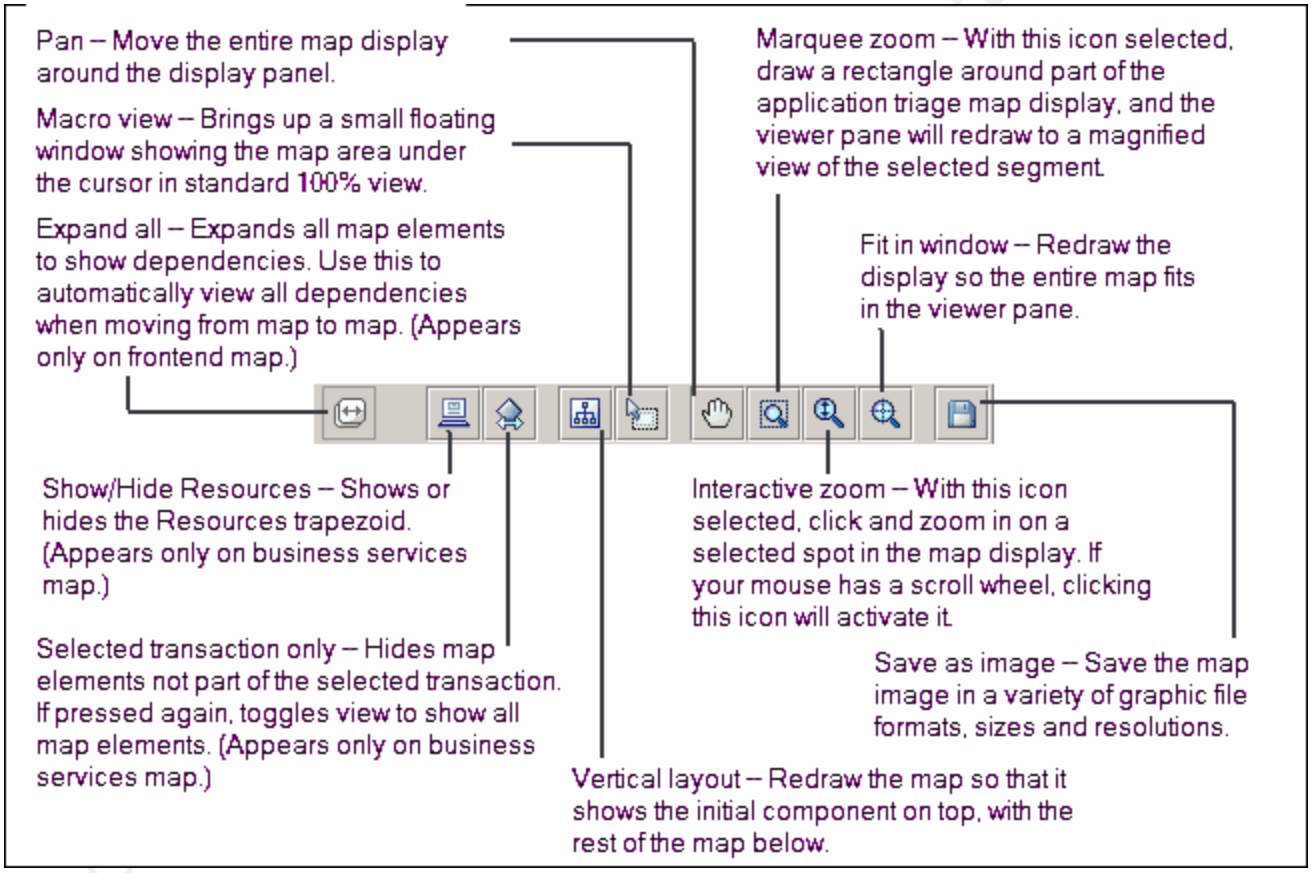
To use the dog-ear control to reveal map element dependencies:

- Double-click the dog-ear.

In the Location Map, elements with dependencies have a similar dog-ear control. See Location Map (see page 119).

Application Triage Map Controls

You can control and customize the graphical application triage map display using the controls at the top of the application triage map viewer.



Application triage map context menu

In addition to the controls on the application triage map menu bar, a menu of commands is available if you right-click any frontend node on the map.

Some of the menu options are available for backend connectors and backend nodes, as noted below.

When you right-click frontend_*A,* you can:

- Show locations for frontend_*A*—Choosing this menu item opens the List of Physical Locations pane in the lower part of the application triage map tab. See List of physical locations (see page 86).

  The *Show locations for...* menu option is also available when right-clicking on the backend metrics icon.

  When the List of Physical Locations table is visible, this menu item reads *Hide Locations for <frontend_A>.*

- View health metrics for frontend_*A*—Choosing this menu item changes the display from the application triage map to a display of multiple health metrics, the same display as if you had selected the Health sub-node for this frontend in the Map tree.

  The View health metrics for... menu option is also available when right-clicking on the backend connection metrics icon.

When you right-click frontend_*B* that is a dependency of frontend_A, you can:

- Show all dependencies for frontend_B -- Choosing this menu item "unrolls" the selected map node to display its dependencies.

  This menu item is available only on dependent frontends; that is, frontends which are not currently selected in the By Frontend tree.

- Display Map For Frontend_B -- Choosing this menu item:
  - selects Frontend_B in the Map tree.
  - changes the application triage map display to a display of frontend_B.

Application triage map refresh

The application triage map displays are based on data sent to the Enterprise Manager by agents on the application servers where an application is deployed. When this data changes, the application triage map will display a control at the top of the map display:



To refresh the map display based on the latest data:

- Select Reload.

You can also enable auto-refresh, to avoid seeing this notification.

CA Technologies Proprietary Documentation. Authorized Use Only

To enable automatic refresh of the map display:

1. Select Enable auto-refresh option.
2. Select Reload.

After this change, the map refreshes automatically without notifying you that data has changed.

To change auto-refresh settings:

1. Select Workstation > User Preferences.
2. Click the Investigator tab.
3. Check or uncheck Auto-refresh underlying map data.
4. Click Apply.

List of Physical Locations

When you double-click a node or a live connection arrow in the visual application triage map display, a table listing physical locations of the selected system element (for example, a frontend or a backend call) appears in the bottom pane of the triage map tab.

The table displays agent locations reporting data for the application you select in the application triage map. In this display, you can:

- See the name of the node you selected immediately above the table.
- Browse the list of locations by scrolling up and down the list.
- Sort the list by clicking any of the table column headings.
- See a tooltip with more information by hovering your pointer over one of the rows in the list of locations, and a tooltip will display the path to the node in the Browse tree where you can see more of the metrics reported by the agent at this location.
- Copy text from the table to your system's clipboard.

Note: Users will only see applications for which they have security permissions. See How user permissions affect what you can view (see page 130).

To jump from the list of locations to the same location in the metric browser tab:

- Double-click the row of the table.

CA Technologies Proprietary & Documentation. Unauthorized Use Only.

When you double-click a row in the physical locations list, the Workstation display will jump to the metric browser tab, opening the tree structure to the location.

For frontends, this is Agent|Frontends|Application|<App_Name> so you can view metrics for that frontend and perform transaction tracing. See Frontend overviews (see page 136) for an illustration of the Frontends display on the metric browser tab.

For backends, this is Agent|Frontends|Application|<App_Name>|URLs|default|Called Backends.

A detailed description of how to combine use of the triage map tab and the metric browser tab displays for application problem triage and diagnosis is contained in the next chapter, Monitoring application performance and problems (see page 163).

Limits on Map Display

Enterprise Manager uses a threshold, set using the property *introscope.apm.query.max.results*, to clamp the amount of data the Workstation attempts to display in the application triage map. When you click on a frontend in the By Frontend tree, or unroll a map component to show its dependencies, and the amount of data to display exceeds this threshold, you will see a message stating "The map is too large to display."

In this case, you may:

- adjust the Enterprise Manager's threshold level upward, to determine whether a higher threshold causes the map to be displayed without impact on Introscope performance. To do this, see the information about application triage map data clamping in the *CA APM Configuration and Administration Guide*.
- configure the *introscope.apm.data.timeWindow* property to a lower value. In some environments, this may lower the number of dependencies enough for the map to be displayed. To do this, see the information about application triage map data collection and aging in the *CA APM Configuration and Administration Guide.*

For more guidelines about configuring your environment for optimal performance, see the *CA APM Sizing and Performance Guide.*

Using the Application Triage Map

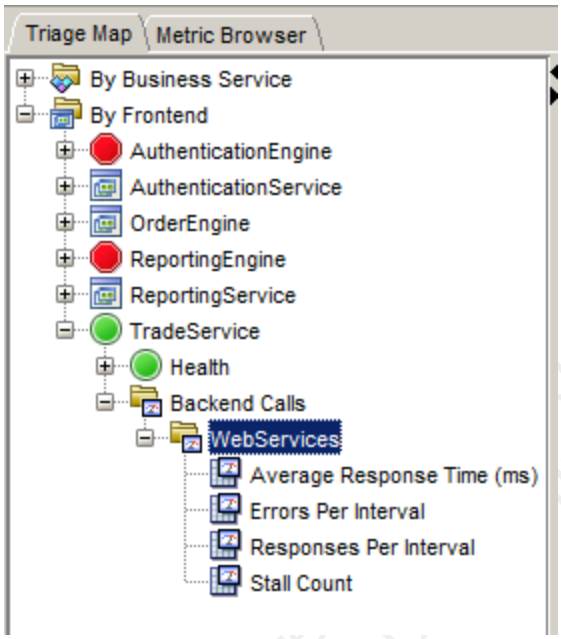
This section discusses the application triage map interface in detail.

For information about using the map to actively monitor system performance, see the chapter Monitoring Application Performance and Problems. (see page 163)

CA Technologies Confidential Document. Unauthorized Use Prohibited.

By Frontend Tree and Metrics

The illustration below shows application health metrics under the application ApplicationA, and health metrics for backend calls made by the application.



Things to notice:

- In the tree structure on the left, each of the nodes immediately under the By Frontend node -- AuthenticationEngine, AuthenticationService, and so on -- represents a frontend. Those which have been configured to display alerts are marked by an alert indicator.
- Where they appear, the alert indicators on both tree nodes and on map elements, showing the aggregated status of the metric or element they decorate. See About Alerts and Alert Indicators (see page 37).
- For each application, you can view:
  - Aggregated frontend health metrics.
  - Metrics for each called backend.

Other things to note about the By Frontend tree, not shown in the illustration above:

- When you click a frontend node, the application triage map tab displays a graphical map of the application and its dependencies in the viewer pane. See Frontend View of the Application Triage Map (see page 90).
- The application triage map displays metrics when you hover your pointer/cursor over some of the map elements. See Tooltip Metrics in the Frontend View of the Application Triage Map (see page 91).

Frontend health metrics

When you select the Health sub-node under a monitored component, the Overview tab displays the five basic Introscope metrics (see The five basic Introscope metrics (see page 349).) The metrics under the Health node are aggregated across all the agents which are reporting data for this application.

To see metrics for an individual agent which is monitoring the application:

1. Select a frontend under the By Frontend node.
2. In the application triage map tab, select the map node representing the same frontend.

   In the lower part of the Map Tab, Workstation displays a list of individual physical locations where an instance of the frontend is installed.

For more information about browsing this list, see List of physical locations (see page 86).

Backend Calls metrics

The Backend Calls node in the Triage Map tree displays a list of backend calls made by the frontend. For each call to a backend, you can see four metrics:

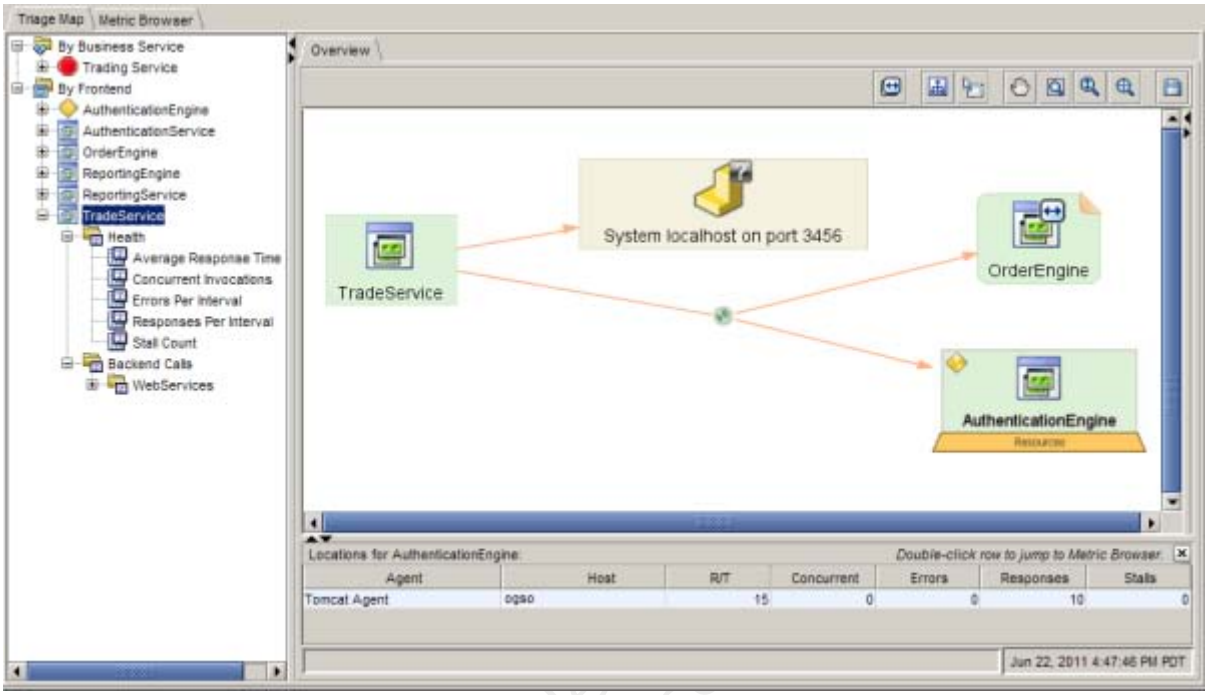
- Average Response Time (ms)
- Errors Per Interval
- Responses Per Interval
- Stall Count

When the backend call is a socket connection, the tree displays an additional metric:

- Concurrent Invocations

CA Technologies Prerelease Documentation. Authorized Use Only.

Frontend View of the Application Triage Map



When you select one of the applications from the tree under the By Frontend node, Workstation displays a graphical map of the selected application in the application triage map tab.

In the illustration above, the user has selected a frontend node, TradeService, in the map tree. The resulting application triage map takes the TradeService frontend as its starting point.

In the map itself, the user has selected the AuthenticationEngine frontend, one of the dependencies of TradeService.

The user has further chosen to display AuthenticationEngine's locations in the **Locations table** under the map. The Locations table shows a list of physical locations where agents are reporting metrics for the selected frontend. You can browse this list to look for metric spikes on individual hosts. See List of Physical Locations (see page 86).

About alert indicators

Colored alert indicators in the tree and the map show the aggregated status of the metric or element they decorate. See About Alerts and Alert Indicators (see page 37).

CA Technologies Proprietary Authorized Use

Showing and Hiding Dependencies

When a supported frontend dependency has additional dependencies available to display, the map element displays a "dog-ear" upper-right corner. The dog-ear looks like an orange turned-down triangle. Both OrderEngine and AuthenticationEngine have dog-ear controls in the illustration above, signifying that both frontends can be expanded to show their dependencies.

To reveal additional dependencies:

- Double-click the "dog-ear" corner of a supported frontend map element.

Note: Users will only see applications for which they have domain security permissions. See How user permissions affect what you can view *(see page 130).*

To reveal *all* dependencies:

You can do either of these:

- Right-click any dependent frontend-component (to use an example seen in the above illustration, the node ApplicationC which is a dependency of ApplicationA) and click "Show All Dependencies of <Frontend_node>".
- Click the Expand All icon in the toolbar at the top of the map. The Expand All icon depicts a double-headed horizontal arrow:



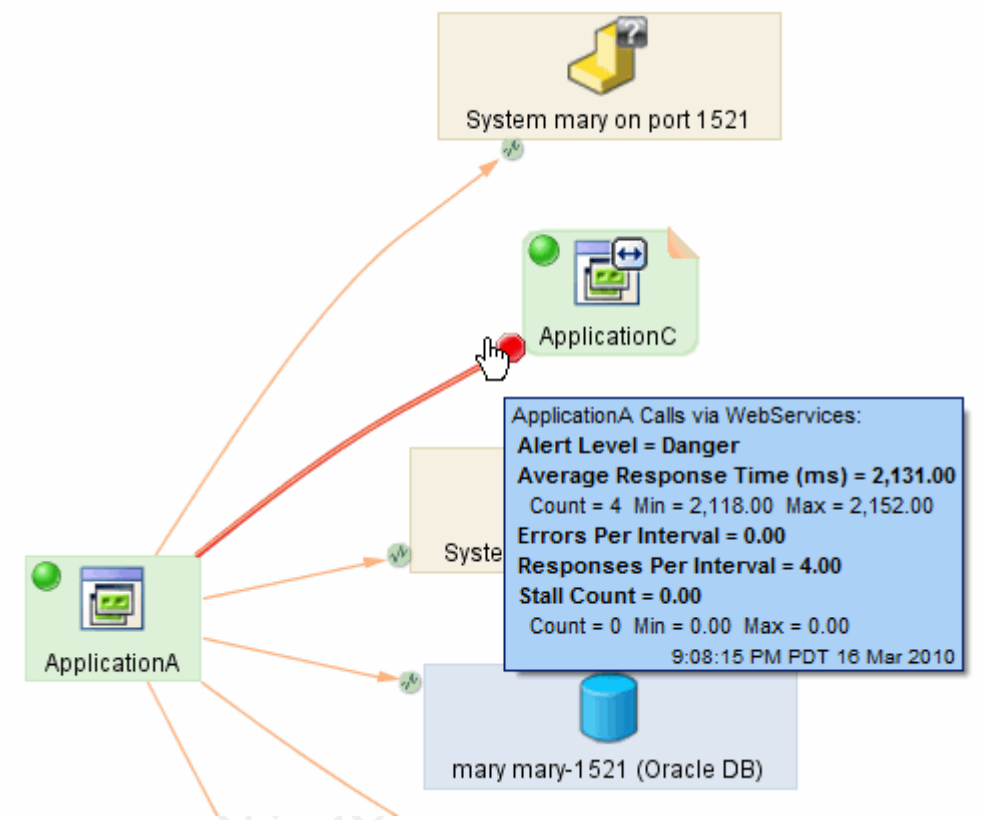
Tooltip Metrics in the Frontend View of the Application Triage Map

When you hover your cursor over certain elements in the application triage map, a tooltip shows aggregated metric information about the corresponding system element.

The metrics are aggregated across all agents reporting metrics for the application. (To see metrics for individual agents, double-click the map node, and a list appears in the bottom of the map tab. See List of Physical Locations (see page 86).)

CA Technologies Prerelease Documentation. Authorized Use Only.

The illustration below shows the tooltip on the backend connection metrics icon:



Each tooltip shows:

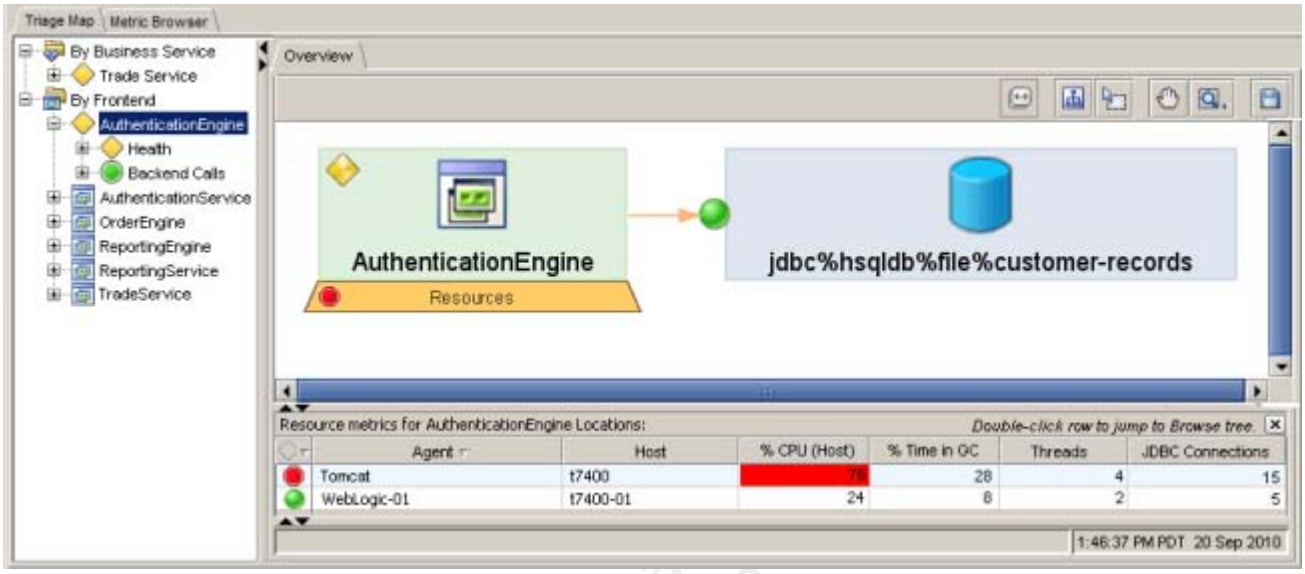
- The same five aggregated health metrics that the Overview tab displays when you select the Health node of an application.
- A timestamp showing when the data was collected.
- If an alert has been configured on the backend call, the tooltip also displays the Alert Level.

For general information about tooltip metrics, see Tooltips (see page 70). For more information about how the icon appears in the application triage map, see Backend connection metrics icon (see page 81).

CA Technologies Proprietary Authorized

Resources Element

The illustration below includes frontends decorated with a trapezoid. This Resources Element allows you to access metrics for that frontend's resources.



Things to notice in the illustration above:

- AuthenticationEngine is selected in the Triage Map tree.
- The Health node in the tree under AuthenticationEngine is decorated with the same alert indicator which appears on the frontend element in the map.
- The Backend Calls node in the tree is decorated with the worst-case alert of any of the backend calls for that frontend. In this example, there is only one backend call, to a database, so the node reflects its state.
- The Resources trapezoid is selected (as shown by its orange color).

  When not selected, the Resources trapezoid is colored green-blue and its label is italicized.
- The user has chosen to display the resource metrics for AuthenticationEngine locations in the table under the map.
  - The Danger status of AuthenticationEngine's resources is a summary of the statuses for each individual location supporting the frontend. Two of these are displayed in the table.
  - Each of the individual agents has its own status. The summary status reflects the worst of these.
  - The offending metric causing the Danger alert is highlighted in the table.

The Resources element can also appear on elements in the By Business Service map. See Resource Element in the By Business Service Map (see page 99).

CA Technologies Proprietary
Authorized Use

For more information about Resource metrics:

- Create and Edit Resource Alerts (see page 109)
- Configure Resource Metric Paths
- Resources Tab View (see page 144)

Understand the Resources Element Display

To see the resource metrics for a frontend:

1. Right-click the Resources trapezoid.
2. Click *Show/Hide Resources for "<Frontend_Name>" Locations.*

   This toggles showing/hiding the metrics table.

To evoke a tooltip for the Resources element:

- Hover your cursor over the Resources trapezoid.

  The tooltip gives element's alert status.

To see resource metrics for one location in the Metric Browser tree:

- Double-click the row in the table whose location you want to inspect.

  The display changes to the Resources Tab View (see page 144) in the Metric Browser tree to show the resource metrics of the selected location.

Note: In the Metric Browser tree they appear under the agent node as follows:

```
<Agent_Name>
|
|--Agent Stats
    |
    |--Resources
       |
       |--<metrics>
```

By Business Service Tree View

The By Business Service tree has two kinds of metrics. Each of them appear in the Business View map display and beneath the Business Transaction nodes in the tree:

- Customer Experience metrics
- Business Transaction Component health metrics

CA Technologies Pre-Release Documentation. Authorized Use Only.

Business Service and Business Transaction Metrics

Under the By Business Service node in the triage map tab, you can view a hierarchy of business services, transactions, and business transaction components, if this hierarchy has been configured using the Customer Experience interface. (For information on how this hierarchy is recorded and configured, see the *CA APM Transaction Definition Guide.)*

For an example of a tree depicting the Business Service - Business Transaction - Business Transaction Component hierarchy, see the topic By Business Service Application Triage Map.

For each Business Transaction Component, you can view the standard Introscope metrics (see page 349).

To see metrics and a list of locations displayed in the Overview tab:

- Select the node for the Business Transaction Component (BTC).

Note: The metrics that Workstation displays for a BTC are aggregated across all hosts where an agent has been configured to report metrics for that BTC.

To see individual metrics and a list of location reporting those metrics:

- Select one of the metrics listed under the BTC node in the tree.

To see an application triage map for a Business Transaction:

- Select a Business Transaction in the tree.

  An application triage map for that Business Transaction appears in the viewer.

Customer Experience Metrics

When a TIM has been deployed to monitor web application customer experience metrics and defects, you can also view customer experience metrics in the By Business Service tree under the Business Transaction node under each Business Transaction. The customer experience metrics are:

- Average Response Time (ms)
- Total Transactions Per Interval
- Total Defects Per Interval

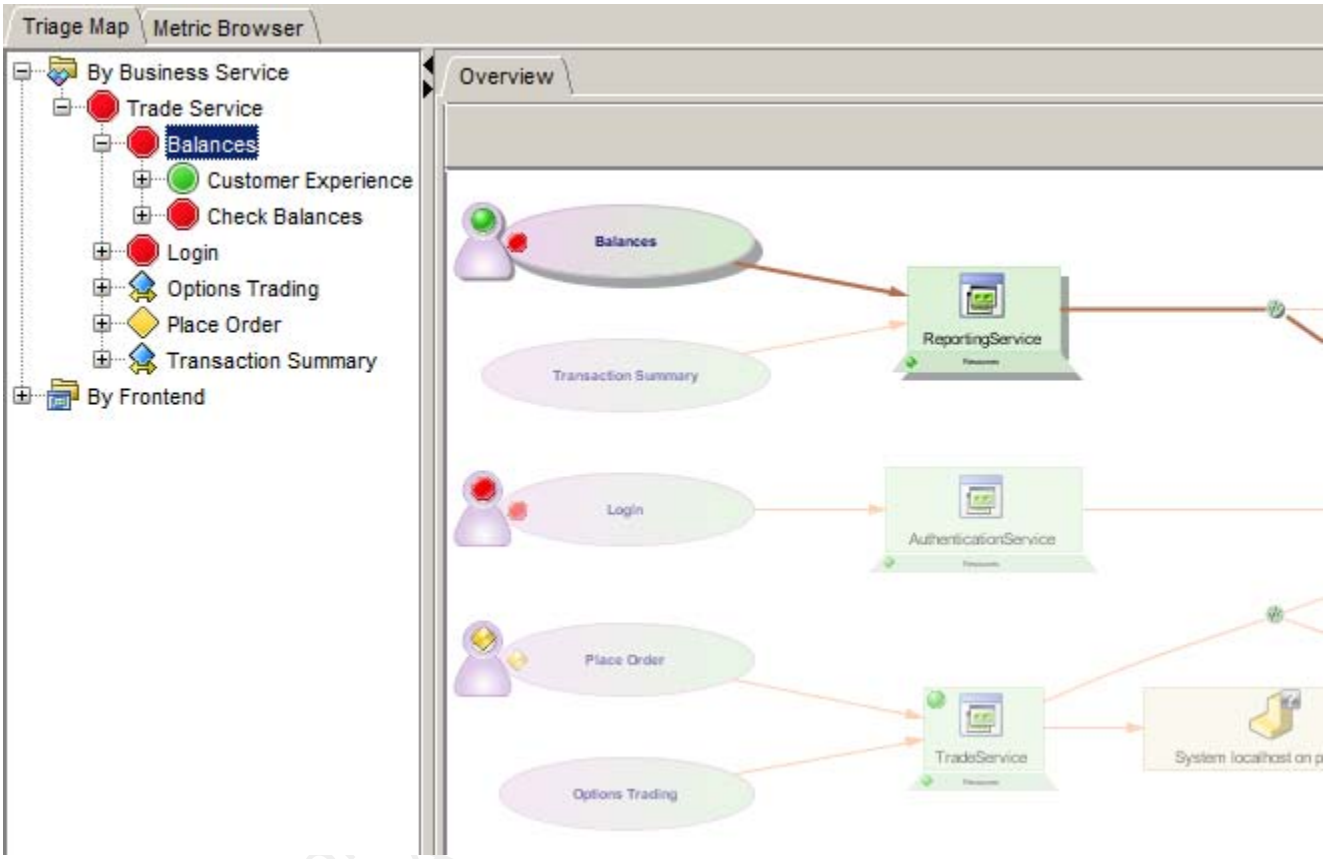
For definitions of these metrics, see Customer Experience Metrics (see page 373) in the Metrics Reference Appendix. For more about TIM components and how they collect and report data, see the *CA APM Configuration and Administration Guide.*

NOTE: If your system does not include a TIM, no Customer Experience metrics are collected, reported to the Enterprise Manager, or displayed in the Workstation.

CA Technologies Prerelease Documentation
Authorized Use Only

By Business Service Application Triage Map

When you select a Business Service or one of its child Business Transactions (BT) from the tree under the By Business Service node, Workstation displays a graphical map of the selected Business Service or BT in the application triage map tab.



In the tree, notice:

- The Trade Service business service has several child transactions: Balances, Login, Options Trading, Place Order, and Transaction Summary.
- The Balances business transaction (BT) is selected in the tree.
- The Balances BT node is expanded in the tree to show:
  - The Customer Experience node (see Customer Experience Metrics (see page 95)).
  - A child Business Transaction Component (BTC), Check Balances.

CA Technologies

In the map, notice:

- The customer experience (CE) icon resembles a chess pawn, and appears next to the BT oval to which it corresponds, when TIMs are available. (A TIM is the Customer Experience transaction processing system. See the *CA APM Transaction Definition Guide* for more information.)
- When you select a node int he tree, the corresponding map element is highlighted with a shadow, and its dependencies appear in full color, while non-participating components appear dimmed.
- Relationships between map components are represented by arrow connectors. Connections between a selected component and its dependencies are emphasized with darkened lines.
- While Balances is a Business Transaction (BT), the alert on the Balances oval corresponds to its child BTC, Check Balances. This is by design.
- Tooltips with metrics and other information appear when you hover your mouse over various map elements. For example, see the illustration in The Triage Map Tab (see page 75).

Comparing alert indicators on the CE icon and BT oval

The CE icon appears on the map next to a Business Transaction (BT) oval element when a TIM component is reporting customer experience metrics. When no CE icon appears, it means that no TIM is deployed and no customer experience metrics are available.

In the illustration above, notice that the alert on the Balances CE icon corresponds to the Customer Experience tree node under Balances, while the Caution alert for the Check Balances tree node corresponds to the BT oval. Even though they reflect the same Business Transaction, the CE icon and the BT oval may display different alert states because their alerts are based on different metrics. Customer experience metrics may include transaction components that Introscope cannot see or ignores; also, the customer experience response time metric includes client-side network time, while the BT response time metric does not.

The alert indicator on the CE icon displays the worst state of the three Customer Experience metrics (see page 373). For example, if two of the metrics are Normal (green) and only one is in Danger (red) state, the CE icon will display a Danger (red) indicator. Similarly, the alert indicator on the BT oval displays the worst state of its Health metrics. (To view these metrics, right-click the element and select View Metrics for <Element_Name>.)

To view health metrics for an element in the map:

1. Right-click the element.
2. Click "View Health Metrics for <*Element*>".

CA Technologies Internal/Restricted Use Only. Please Do Not Distribute Documentation.

Context menu for the CE icon

Right--clicking the CE icon displays a menu with these choices:

- Find Incidents in CEM for "<BT Name>"...

  Launches a browser which opens to the CA APM Customer Experience Incidents page in the CE Console filtered on the current Business Transaction.

  NOTE: Incidents are not directly related to the BT's alert status. The condition that is causing the BT to show a CE alert *may* also be causing incidents to be triggered, but Customer Experience alerts are defined differently from Customer Experience incidents.

- Browse all Customer Experience Metrics for *<BT_Name>*

  Changes the display to the metric browser tab, with the Browse tree expanded to the Business Transaction corresponding to this CE icon, so you can see the hierarchy up to the TESS agent and understand the Business Transaction's context.

- View Metrics for “<BT_Name> Customer Experience”

  Selects Customer Experience under the BT node in the tree, causing the map to be replaced by trend charts for the three customer experience metrics.

For more information, see:

- Business Service and Business Transaction Metrics (see page 95)
- Customer Experience Metrics (see page 95)

Other elements in the By Business Service map

Some called backends (for example, the databases symbolized by the blue cylinder icons) do not have alerts configured, while other backends are "unknown," as symbolized by the yellow puzzle piece decorated with a question mark. When backends such as databases do display alert indicators, those alerts are being imported to Introscope via Catalyst. For more information about data imported by Catalyst, see:

- Catalyst Alert Indicators (see page 38)
- Viewing Data from Catalyst (see page 119)

To show/hide the list of locations:

- Right-click any map node representing a Business Transaction or Business Transaction Component. If the locations list is not visible, you can select "Show Locations for <Selected_element>". If the locations list is already open, you can select "Hide Locations for <Selected_element>".

CA Technologies Product Documentation
Authorized Use Only

Tooltips in the Business Service Application Triage Map

When you hover your cursor over certain elements in the Business Service application triage map, a tooltip displays aggregated metric information about the corresponding system element.

The metrics are aggregated across all agents/TIMs reporting metrics for the application. (To see metrics for individual agents/TIMs, you double-click on the map node, and a location table appears in the bottom of the map tab.)

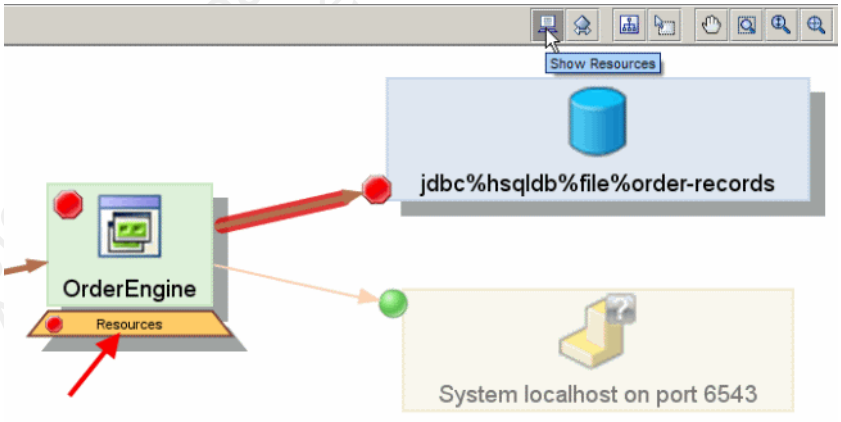
For general information about tooltip metrics, see Tooltips (see page 70). For more information about how the icon appears in the application triage map, see Backend Connection Metrics Icon (see page 81).

Resource Element in the Business Service Map

You can choose to display resource metrics for frontends in the By Business Service triage map.

To display resource metrics in the By Business Service map:

- Click the Show/Hide Resource Metrics button in the toolbar:



The Resources trapezoid (bottom red arrow) will then appear on frontends where resource metrics have been configured.

For more information about the Resources element of the map and resource metrics:

- Resources element (see page 93) main topic
- Create and Edit Resource Alerts (see page 109)
- Resources Tab View (see page 144)
- Resource Metrics (see page 372)

CA Technologies Product Documentation. For Internal Use Only.

Using alerts

Alerts are a powerful CA APM feature, enabling you to set thresholds on metrics and to execute actions when metrics cross thresholds. Many of the objects in the application triage map can display alerts, which reflect the worst-case status of any of the object's baseline metrics.

NOTE: Objects appearing in the map which are imported via [set value for the rose variable at the book level] may display alert indicators, but their alert status is imported with the object, and cannot be manipulated or reconfigured by CA APM users. For more information, see How Catalyst Alert Indicators Appear (see page 38).

Alert Indicators

You can use alerts with most of the elements of the Triage Map tab, in the By Frontend and By Business Service views and with the base metrics that provide the foundation for the displays.

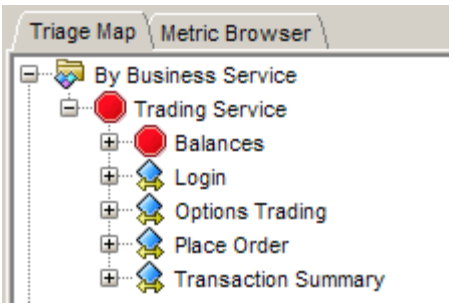
Alert indicators appear on both tree nodes and on map elements when alerts have been configured on those elements. They represent the aggregated status of the element, determined by the alerts configured on that component.

For information about the basic appearance of alert indicators and what they represent, see About Alerts and Alert Indicators (see page 37).

Alert indicators in the Map tree

When you configure alerts on frontends or backend calls in the Triage Map tree, alert indicators appear in place of the standard tree icons. The illustration below shows one of the business transactions in Danger status; the others have no alerts configured.

Notice also that alert states "bubble up." In this example, the Trading Service business service has five business transactions associated with it. One of them is in Danger status, therefore Trading Service is in Danger status. See "Alert Propagation" below.



CA Technologies Pre-release Documentation
Authorized Use Only

Other behavior of nodes in the Triage Map tree

No data: If no data is coming through the frontend or connection, the icon in the tree will revert to a standard tree icon. This may happen at the end of a 15-second interval, or may happen because the metrics have aged out, if aging is enabled.

Downtime: An additional icon, the downtime icon, means that an alert has been configured on this frontend/backend, but the alert is in downtime, as determined by the alert downtime schedule. The downtime icon is a gray octagon with a dark gray bullseye:



For more information, see Working with Alert Downtime Schedules (see page 326).

Alert propagation

In tree displays: Alerts status propagates up the tree, with parent nodes taking the worst alert status of any child nodes. For example, if a called backend anywhere in a dependency chain has a status of Caution, all parent nodes up the chain will have, at best, a status of Caution.

When multiple alerts are defined under the same frontend or backend calls group, a parent node will have the worst status from underlying alerts.

In map displays: Alert indicators in the application triage map display do not propagate in the same way. Alert indicators on map elements reflect only the alert status of the element on which they appear, and not the status of any dependencies of a particular element. This behavior is different from how alert indicators behave in the Map tree.

Base, Contributing and Compound Alerts

When configuring or interpreting Triage Map alerts, it is important to understand the distinction between base, contributing and summary alerts, and how they behave under various circumstances.

It is important to note that triage map alerts are not the same as simple alerts.

Base alert

You configure a base triage map alert by configuring alert thresholds and other attributes on its contributing metrics. The base alert for a frontend or called backend has an actions tab and its own set of properties, including Description.

Contributing metrics

Each base alert has a set of five contributing metrics. You can configure thresholds on these, but not actions.

CA Technologies Prerelease Documentation
Authorized Use Only

Compound alert

Does not have its own contributing metrics, but is a summary of other triage map alerts. You cannot configure thresholds for it, but you can configure actions for it.

Alert Threshold Line Display

Health overview and health metric graphs display danger and caution alert threshold lines by default. Included in this functionality:

- The danger threshold is displayed as a solid red line; the caution threshold is displayed as a dashed yellow line.
- Positioning your cursor over these lines displays a tooltip with information about the alert threshold.



- Alert threshold lines will be displayed for metrics in a downtime state.

NOTE: Alert threshold lines are based on the current alert threshold definitions. Be aware of this when the chart is in historical mode, as the threshold lines will display current threshold values, not "historical" or formerly defined values.

To toggle the alert threshold line display:

- Click the Hide Alert Thresholds button.

NOTE: Clicking the Hide Alert Thresholds button toggles the display and causes the button to appear "pressed down" or "popped back up;" in either case, the button label does not change.

View Alert Details

The Alert Details pane allows you to view a list of all the metrics that currently contribute to the alert status displayed in the Triage Map.



CA Technologies Product Release Documentation. Limited Use Only.

To open the Alert Details pane:

1. Right-click an element in the application triage map which is decorated with an alert.
2. Select *View Alert Details for* "<Object_Name>"

   The pane opens on the right side of the application triage map.

With the Alert Details pane open, selecting a different Triage Map element causes the pane to be updated with the metrics which contribute to the alert status of that element.

If you have SuperDomain permission, you can change the thresholds on any of these contributing alerts.

To change alert thresholds:

1. Right-click any row in the list of alerts in the Alert Details pane.
2. Select *Change Definition...*

The Edit Alert for *<Object_Name>* dialog appears. To use this dialog, see Creating and Editing Application Triage Map Alerts (see page 104).

Create and Edit Application Triage Map Alerts

Alerts are a powerful tool for monitoring and triaging applications. This section tells how to create and edit alerts on application triage map elements.

Note: You must have SuperDomain permission to create and edit application triage map alerts.

Creating Alerts from an Application Triage Map Element

Follow these steps to add an alert on an object in the triage map:

1. Right-click a frontend, backend call, or other alertable element in the map.
2. Select *Edit Alert for <Object_Name>...*
3. In the left pane, identify a metric you want to contribute to alert status.
4. From the Problem drop-down, select the Problem you want to trigger the alert. Available values are:
   - Value Too High -- The alert will trigger when the metric value exceeds the threshold.
   - Specific Bad Value(s) -- The alert will trigger when the metric value is equal to the threshold, and subsequently the threshold will be referred to as "Bad Value" rather than "Threshold Value."
   - Value Too Low -- The alert will trigger when the metric value drops below the threshold.
   - Unexpected Values -- The alert will trigger when the metric value is not equal to the threshold value, which subsequently will be referred to as "Expected Value" rather than "Threshold Value."
5. In the Summary Tab of the Threshold Settings region:
   a. Set the Threshold Values for Danger and Caution alerts.

CA Technologies Pre-release Documentation
Unauthorized Use Only

b. Set the Sensitivity Levels for Danger and Caution alerts. Available values are:

- High -- For Danger threshold, 2 values in 2 samples. For Caution, 1 value in 1 sample.
- Medium -- For Danger, 4 values in 4 samples. For Caution, 2 values in 2 samples.
- Low -- For Danger, 6 values in 6 samples. For Caution, 4 values in 4 samples.
- Custom... -- This allows you to set your own sensitivity levels using the Select Custom Sensitivity Settings dialog.

Helper text is displayed in blue type as you set sensitivity levels, helping you understand your settings.

6. Optional: If you want to set different or less sensitive thresholds on the locations:
   a. Select the Locations tab.
   b. Set different threshold values for the locations which report the metric.
   c. Click *Apply.*
7. Optional: The Properties tab allows you to:
   - Enter a description for the alert.
   - Disable the alert by selecting Disabled (all). This disables all individual contributing alerts AND the object alert as a whole.
   - Reset the Interval
   - Select *Location alerts contribute to overall status.*

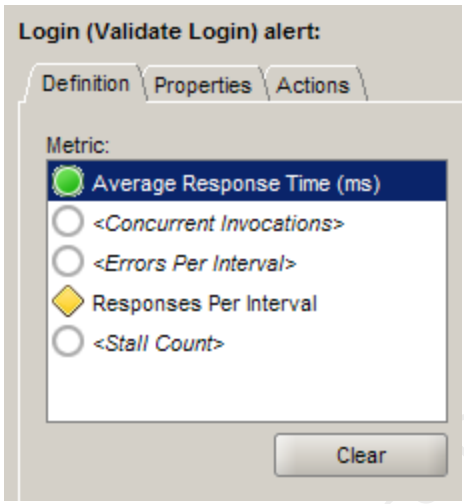
   Note: By default, the map object will display the alert state only if its summary metric is alerting; to see if any individual location is alerting, you must view the Locations table. If selected, the Location alerts contribute to overall status option means that a frontend's alert will include the alerts for each of its locations. Thus, if any single location of the frontend is alerting, the map will show an alert state.

Optional: Use the Actions tab on the Creating and Editing Alerts dialog to add an action to the alert.

CA Technologies Prerelease Documentation. Authorized Use Only.

A Note on Enabling and Disabling Contributing Alerts

Suppose you have a frontend "Login" on which two of the possible five health metrics have been configured. In this scenario, the two live metrics are contributing metrics to the base Health alert.



After being created, each of these metrics are enabled unless and until the Disabled check box on the Definition tab is selected for any of the alerts.

In the case where the contributing alerts are disabled by means of this check box, the parent base alert appears as a gray disc, showing it is getting no data. It will appear in the Delete Unused Alerts dialog after 8 minutes. This happens because the Delete Unused Alerts dialog queries for and displays contributing alerts which have been disabled for the last 8 minutes.

Actions Tab

You can add an alert action to alerts:

- When creating an alert directly from application triage map objects
- On any application triage map object with an existing alert (unless the alert status is imported from a CA Technologies application outside CA APM such as eHealth).

Add an Action to an Alert

To add an action to an application triage map alert:

1. If the Edit Alert dialog is not already open, find the node in the tree corresponding to the alert where you want to add an action. Right-click the tree node and select Set Notifications for <Object_Name>...

   NOTE: This menu item is enabled only if an alert exists for a particular object (i.e., if an alert has been defined for any of its child objects).

2. Select the Actions tab.

CA Technologies Pre-release Documentation. Authorized Use Only.

3. Select an event under the “Trigger Alert” drop-down, one of:
   - When Overall Severity Increases
   - Whenever Overall Severity Changes
   - Each Interval While Problem Persists
4. Optional: Set a delay time setting.

   NOTE: This is disabled if “Each Interval While Problem Persists” is selected in the Trigger Alert.
5. Select from the list of Available Actions.

   The list of available actions is populated with actions already created in the Triage Map Configurations Management Module. If you want to add to the available actions, use the Management Module editor to create new actions in that Management Module.

   NOTE: The alerts you create directly from application triage map objects are also saved in the Triage Map Configurations Management Module. However, they cannot be edited using the Management Module editor; they must be edited using the controls explained in this topic.
6. Select Add to move the selected action(s) to the Actions and Trigger States list.
7. Under “Selected Actions and Trigger States,” select either or both the Danger and Caution states.
8. If you are using the APM Catalyst Connector and want to send the alert status to Catalyst, select the “Broadcast via Catalyst” option. This sends status changes to Catalyst.

   NOTE: Using this option requires you to perform configuration tasks. See the note below on sending alert details to Catalyst.
9. Click OK.

Adding an action to a summary alert

Summary alerts aggregate status from base-level alerts. You use the same steps above to add an action on a summary alert.

A Note on Sending Alert Details to Catalyst

[set value for the rose variable at the book level], which is CA Technologies' unified platform for application development, allows CA applications to exchange data. In the case of application triage map alerts, APM Catalyst Connector allows you to send alert data to the [set value for the rose variable at the book level] platform to be consumed by other CA applications. This is done through the use of an SNMP plugin known as the SNMP Trap plugin.

CA Technologies Pre-release Documentation: Authorized Use Only

Note there are two kinds of alerts in Introscope:

- Application triage map alerts (also known as "entity alerts") as discussed in this topic.
- Traditional (or "legacy") alerts created through the Management Module editor.

If you set up legacy alerts to go to [set value for the rose variable at the book level] via the SNMP Trap plugin, you must enable application triage map alerts to go to [set value for the rose variable at the book level] too. To enable this, you must perform two tasks:

1. Enable the APM Catalyst Connector configuration for application triage map alerts by following the steps in the section "Configure the APM Catalyst Connector" in the *APM Catalyst Connector Guide.* Failure to do this may cause [set value for the rose variable at the book level] to be unaware of changes in alert states that happen when either the Enterprise Manager or the APM Catalyst Connector is down, in which case [set value for the rose variable at the book level] may retain incorrect data about the state of an alert.
2. In the *IntroscopeEnterprisemanager.properties* file in *<EM_Home>*/config, configure the introscope.apm.catalyst.triagemapalert.snmp.destination.host.ip parameter with the IP address of the Enterprise Manager where the APM Catalyst Connector is installed.

   In a cluster environment, you configure this parameter in the *IntroscopeEnterprisemanager.properties* file on the MOM Enterprise Manager.

Note: The Trigger Alertt setting "Whenever severity increases" has slightly different meaning for <wpam> alerts and [set value for the rose variable at the book level] alerts. When setting the trigger on alerts to be sent to [set value for the rose variable at the book level], the trigger refers to the CA APM version of "Whenever severity increases."

Delete Alerts

Triage Map alerts can be deleted by right-clicking an alerting element and using the Edit Alerts dialog; removing all the individual metric alerts contributing to a particular Triage Map alert deletes the alert on the Triage Map element. But when an element becomes inactive, and disappears from the map and the tree, this method is not available. In that case, you use the Delete Alerts dialog.

To delete inactive or disabled alerts:

1. In the Triage Map By Frontend tree, locate the node for the frontend whose alert you want to delete.
2. Select *Delete Unused Alerts…*

   The Delete Unused Alerts dialog appears. The dialog displays all Triage Map alerts that are currently in a gray (no data) state or disabled, sorted by name.
3. Select the alerts to delete.
4. Click the Delete button, then click OK.

CA Technologies Pre-release Secure Offline Authorized Use Only

Create and Edit Resource Metrics and Alerts

Resource metrics are based on configurable metric paths, and alerts on resource metrics are saved as objects in a particular Management Module.

To create resource metrics and then configure alerts on those metrics:

1. Create standardized resource metrics.

   To do this, edit the *ResourceMetricMap.properties* file to configure mappings between the different resource metrics currently found on your agents and the standardized metric paths. See the *APM Configuration Administration Guide* for information about this file and how to edit it.

2. Configure alerts on the standardized metrics.

   To do this, you can:

   - edit the corresponding alert objects in the Triage Map Configurations Management Module, or
   - edit the corresponding metric groupings to eliminate boring or troublesome agents from contributing data to resource metrics.

Note: You must have SuperDomain permissions to edit objects in the Triage Map Configurations Management Module.

About Metric Groupings in the Triage Map Definitions Management Module

You configure alerts on resource metrics from the following default simple alerts, which are included in the Triage Map Configuration Management Module. Corresponding Metrics Groupings for each are also included:

- APM RESOURCES % CPU Utilization (Host) – Default
- APM RESOURCES % Time Spent in GC – Default
- APM RESOURCES Threads in Use – Default
- APM RESOURCES Threads Available – Default
- APM RESOURCES JDBC Connections in Use – Default
- APM RESOURCES JDBC Connections Available – Default

These special simple alerts cannot be deleted, copied, renamed, or moved to another Management Module; nor can their Combination or Metric Grouping selections be modified.

CA Technologies Prerelease Documentation
Authorized Use Only

Possible actions are:

- Activate or deactivate the alert.
- Change the Comparison Operator, Thresholds, Resolution, and sensitivity settings.
- Add or remove actions on the alert.

Edit Resource Metric Groupings

To create or edit an alert on frontend resources displayed in the triage map:

1. Right-click a frontend's Resources element, which appears as a trapezoid under a frontend element.
2. Select *Configure Alert for* and select one of the resource metrics.

   The Management Module editor appears, open to the Triage Map Configuration Management Module, which is where resource metric definitions are stored. The Alerts node will be expanded, and the alert corresponding to the metric you chose will be selected.

   The right pane shows:

   - At top, a metric chart showing the present values of the metric and any existing alert thresholds.
   - At bottom, the alert configuration controls.
3. Set the values and triggers for the alert. For instructions on how to use these controls to configure alerts, see Configuring Simple Alert Settings (see page 311).

For more information:

- For information about the Resource Element, see Resources Element (see page 93).
- For resource metric definitions, see Resource Metrics (see page 372).

Historical Mode in the Application Triage Map

The default view of application triage map data is live. You can switch between live mode and historical mode by using the Time Range drop-down menu. For more information, see Time Range Drop-down.

The application triage map shows either an application centric or business centric display of your solution and its components. What you see in either view— live or historical—depends on data gathered by Introscope agents over a defined time range with a specific end point.

A note on alerts in historical mode

When the application triage map displays historical data, alert indicators in the triage map tree continue to display *current* status, not historical status.

CA Technologies Product Documentation. Authorized Use Only.

Live application triage map views

In live mode, the application triage map's display is based on the last 72 hours of data in the APM database. The application triage map will display applications as live unless they have not been exercised in the context of the selected tree node in the last 24 hours; inactive applications will appear as aged (see Aged elements and connections (see page 83)).

Historical Application Triage Map Views

In historical mode, the application triage map's display changes depending on how you manipulate several controls.

Time Range Dropdown

The time range dropdown menu allows you to select the size of the time window to display. For example, if you select 12 Hours, the time bar slider control will then give you a 12-hour window to manipulate. However, see the note below for limitations.

You can also chose Custom Range to use a time-and-date control to specify the start and end point of the historical time range.

For more information, see Defining a custom time range (see page 112).

Because the application triage map always reflects at least 72 hours of data, if you specify a time range of less than three days, the map will display three days of data.

NOTE: Metrics shown in tooltips will be aggregated over the time frame you select, not over the default 72 hours.

Resolution

The Resolution control allows you to select the data sampling interval.

This control is not meaningful when the application triage map viewer displays only the graphical application triage map. However, when you view the list of locations for a map element such as a frontend node, the metrics displayed for each location reflect the number of data points used to calculate the aggregated metrics.

For example, if you select:

- a time range of 1 Hour
- a resolution of 30 Seconds

... then the metrics shown for each physical location are aggregated from 120 data points.

CA Technologies Prerelease Documentation. Authorized Use Only.

Time bar slider

The time bar slider control allows you to reset the end point of the time window being displayed in the application triage map. The scale shown on the time bar changes according to the time range. If you do not move the slider, and do not specify a Custom Range, the data sampling end point is now.

The time bar slider control is shown in the illustration below:



Application overview metrics in historical mode

Over a historical range, an alert color reflects the worst-case value of the heuristic at any point in the historical range. For example, if at any time during a historical range the User heuristic for an agent was yellow, but never red, the Overview tab for that historical range is yellow.

Defining a custom time range

To define a custom time range to view historical data:

1. Select the metric or map view for which you want to see historical data.
2. Select Custom Range from the Time Range drop-down menu.

   The Custom Range window opens, showing the current date (Today) highlighted with an outline.
3. Use the calendar controls to select the start and end dates and times, and click OK.

Workstation now shows the data for the custom range.

How controls affect historical display

When you use these controls to change the start and end points of the time window being displayed in the application triage map, the display reflects changes in application availability:

- In the lower right of the Investigator window, a status message shows the end point for the data being displayed in the application triage map.
- The start point for the data being displayed will always be 72 hours earlier than now, or the end point minus the selected time range, whichever is greater.

CA Technologies Prerelease Documentation
Authorized Use Only.

- For example, if you choose 2 Days in the time range dropdown, the application triage map will still show 72 hours worth of data. This 72 hours is a constant.
- An node or connection will be displayed as aged if it has not been exercised in the context of the selected frontend or business service/transaction for at least an entire 24-hour period prior to the end of the time range being displayed. See Aged elements and connections (see page 83).

  For example, if you select a time range of seven days, and reset the end point to be three days before now, then an aged frontend component shows that the component did not participate in the current map for at least the 24-hour period that ended three days ago.
- A node or connection will be omitted from the application triage map display if the component has not been active in the selected context -- that is, if there is no contextually relevant data for the component at any time during the entire time range.
- You can still view tooltip data for map components; this data reflects the historical time range and resolution constraints which you specify.

Other things to note about the historical application triage map display:

- No alert indicators will be displayed.

The Metric Browser Tab

The metric browser tab lists metrics and other information in a tree format. The high-level nodes immediately under domain nodes represent agents installed on individual application server hosts or the equivalent.

Among the various components the high-level nodes represent are:

- Components of your J2EE or .NET application, such as servlets, EJBs or ASP pages
- Data imported from Catalyst on applications whose health is being monitored by other CA Technologies products such as eHealth, Spectrum, NetQoS and Insight DB
- System nodes, including the host running your app server and the host computer running CA APM
- Events, defects, leaks and other distinctive occurrences

You can view live data in the Investigator, or select a range of time to view historical data. The default view of data is Live.

More Information:

Viewing Historical Data in the Metric Browser Tab (see page 180)
View Data from External Sources Using the Location Map (see page 119)

CA Technologies Prerelease Documentation
Authorized Use Only.

Metrics in the Metric Browser Tab

The default metrics which the Workstation displays in the metric browser tab vary depending on the node you select in the hierarchical tree.

Standard metrics

For monitored frontend and backend application components, as well as for many other application components, Introscope displays the five standard metrics, sometimes called the Blame metrics:

- Average Response Time (ms)—a measure of application response speed.
- Concurrent Invocations—the number of requests being handled at a given time.
- Errors Per Interval—the number of errors occurring during a specified time slice.
- Responses Per Interval—the number of requests that are completed during a specified time slice.
- Stall Count—the number of stalls, or uncompleted requests, that have not been completed before a specified time threshold.

For more information on what each of these metrics means, and how to affect them, see the Introscope Metrics (see page 345) appendix.

In addition to the five standard metrics, and sometimes instead of them, Introscope collects and displays other metrics relevant to the node. These are also listed and explained in the Introscope Metrics appendix.

Frontends and Backends

By default, Introscope defines a frontend as a .war file or .jsp that first handles an incoming transaction to an application.. In a .NET application, the equivalent would be an ASP page.

A *backend* is an external system that a web application relies on for some portion of its processing. Typically this is a database, but it can be any external system such as a mail server, a transaction processing system (such as IBM CICS or BEA Tuxedo), or a messaging system (such as MQSeries). Introscope automatically identifies databases as backend systems by the name of the database. For other external systems, Introscope analyzes the socket activity of the application and names the backend based on the IP address and port that the application is communicating over.

For information about how Introscope determines frontend and backends, and for instructions for using blame-related tracers to explicitly mark frontends and backends, see the section on Configuring Boundary Blame in the *CA APM Java Agent Implementation Guide* and the *CA APM .NET Agent Impelmentation Guide*.

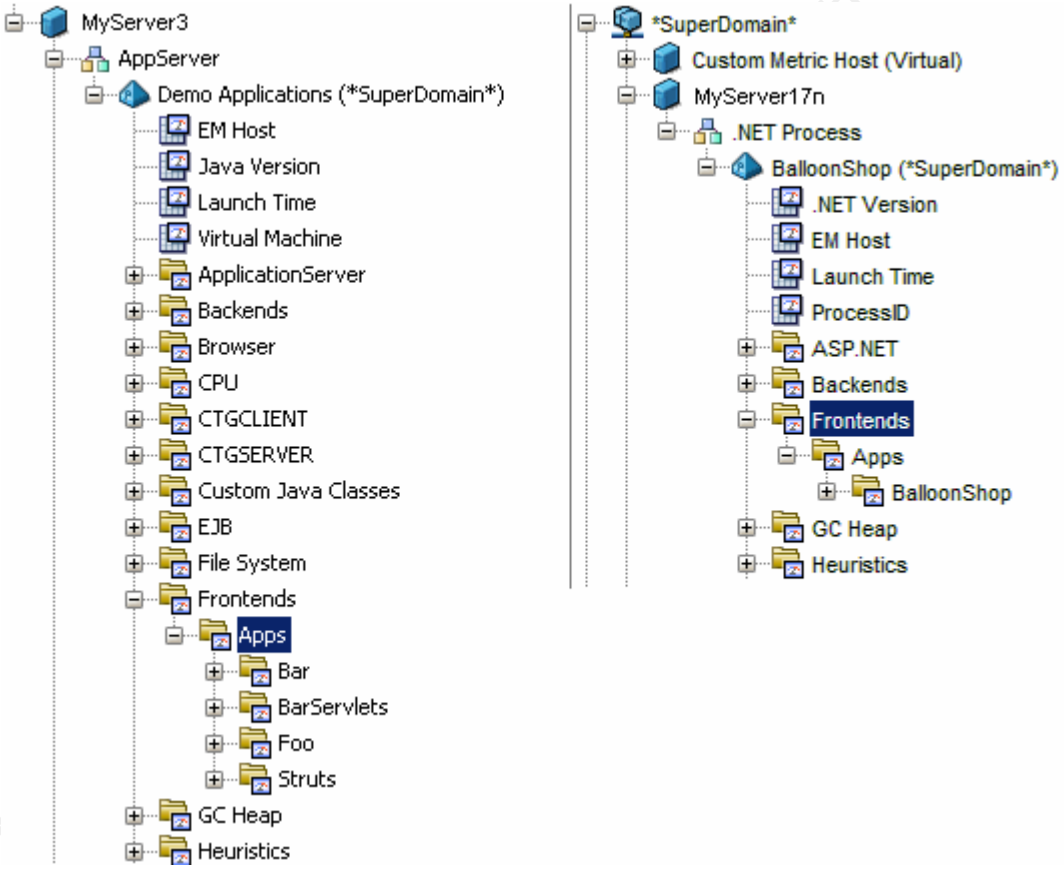
CA Technologies Proprietary Documentation
Authorized Use Only

View Metrics for Frontends in the Investigator

The Frontends node of the metric browser tree contains a node for each frontend, including those automatically detected by Introscope or marked explicitly as a frontend.

The Frontends node organizes frontends by type, and includes an Apps node, under which specific applications are listed.

The illustration below shows the appearance of the Frontends node in the metric browser tab for an application hosted on a Java host (left) and a .NET host (right):



See the illustrations in the section Agent-Centric View (see page 72) for more complete views of the metric browser tab display of Frontends.

Standard frontend metrics

The Investigator displays the five standard metrics (see page 114) for each frontend. See the Metrics Reference Appendix for definitions of these metrics (see page 349).

Other frontend metrics

Other frontend metrics include:

URL metrics

The URLs node under a frontend node shows these metrics for each URL group that is configured for the frontend:

- Average Response Time (ms)
- Concurrent Invocations
- Errors Per Interval
- Responses Per Interval
- Stall Count

URLs that do not match a URL group definition are shown in the Default group. If no URL groups are defined, all URLs belong to the Default group.

Called backend metrics

- Average Response Time (ms)
- Concurrent Invocations
- Errors Per Interval
- Responses Per Interval
- Stall Count

Notice the distinction between *Called Backend metrics* and *Backend metrics.* Called Backend metrics are part of the By frontend node, and display the performance of a backend for a particular URL group.

View Metrics for Backends

The Backends node of the metric browser tree contains a node for each backend, including those automatically detected by Introscope, or marked explicitly as a backend during ProbeBuilding.

Backends are most commonly a database, but may be any external system such as a mail server, a transaction processing system (such as IBM CICS or BEA Tuxedo), or a messaging system (such as MQSeries).

CA Technologies Prerelease Documentation. Authorized Use Only

Database backend metrics

When the backend system is a database, these metrics reflect the activity and performance of the backend across all applications it serves:

- Average Response Time (ms)
- Concurrent Invocations
- Errors Per Interval
- Connection Count—The number of connections to the database during a particular interval.
- Responses Per Interval
- Stall Count

See the Metrics Reference Appendix (see page 345) for definitions of these.

Database backend naming format

This section explains the Introscope naming convention for database backends.

Oracle

The backend name is a concatenation of the Oracle SID string, the database host and port delimited by a hyphen, and the string *(Oracle DB)*.

For example:

```
PRODORCL3 sfoprod6.globex.com-1521 (Oracle DB)
```

DB/2

The backend name is a concatenation of the DBName string and the string *(DB/2 DB)*.

For example:

```
Inventory4 (DB/2 DB)
```

Microsoft SQL Server

The backend name can be a concatenation of the database name, instance name, the database host and port delimited by a hyphen, and the string *(MS SQL Server DB)*, depending on the configuration of the database driver.

If the driver has a database name and an instance name, the backend name in Investigator would look like this:

```
PRODORCL3 (instance Mx22) on prod6.globex.com-1521 (MS SQL Server
DB)
```

CA Technologies Prerelease Documentation. Authorized Use Only.

If the driver has no database name, the backend name in Investigator would look like this:

```
SQLServer on prod6.globex.com-1521 (MS SQL Server DB)
```

If the driver has a database name and no instance name, the backend name in Investigator would look like this:

```
PRODORCL3 on prod6.globex.com-1521 (MS SQL Server DB)
```

If the driver has an instance name and no database name, the backend name in Investigator would look like this:

```
(instance Mx22) on prod6.globex.com-1521 (MS SQL Server DB)
```

Defaults and fallbacks

In cases where the database driver does not support querying for the database name, the name of the database defaults to the JDBC URL, with colon characters (:) replaced by percent characters (%). In some cases even this fallback value is not available, so the database name defaults to the class name of the database driver. Exact behavior depends on the vendor and version of the database driver.

Other backends metrics

Each backend system can also be configured to report the following metrics:

- Commits
- Rollbacks
- SQL

Alert metrics in the agent-centric tree

Each alert color has a metric value:

- Gray—0, no data is available
- Green—1, OK
- Yellow—2, Caution
- Red—3, Danger

There are some special cases to be aware of.

- During the first minute of baseline calculation, the baseliner always reports that the metric is normal. The baseliner is learning during this time, but it will not report problems, to reduce false positives.
- Another special case is in the calculation of baselines for average response time. If an application component is idle, and the average response time metric has a count of zero, the baseliner ignores this value in its learning. It does not learn that 0ms was normal for that time period. Instead, it assumes that the calculated baseline was expected during that time.

CA Technologies Prerelease Documentation
Authorized Use Only

The following table shows how metrics drive alert values in the Overview Tab.

Metric Type	What a yellow indicator means	What a red indicator means
User	Frontend errors are abnormal Frontend response time is abnormal Server execute threads in use are abnormal (for WebLogic Server only) Stall count is abnormal	Frontend errors are *very* abnormal Server execute threads in use are *very* abnormal (for WebLogic Server only) Stall count is *very* abnormal
VM	Aggregate CPU utilization is abnormal and greater than 30 percent JDBC connection pool utilization is abnormal	Aggregate CPU utilization is *very* abnormal and greater than 50 percent JDBC connection pool utilization is *very* abnormal
Backend Summary	Backend response time is abnormal Backend error count is abnormal Backend stalls are abnormal	Backend error count is *very* abnormal Backend stalls are *very* abnormal

You can view the alert metrics by selecting the User, VM, and Backends|*BackendName* metrics, below the Heuristics node in the Investigator.

The underlying metrics that drive the alert metrics appear in the User, VM, and Backends|*BackendName* folders in the tree.

View Data from External Sources Using the Location Map

From the Browse tree, the Location Map tab displays frontends and called backends in the context of their locations on physical or virtual machines.

Compatibility notes

The Location Map feature is backward compatible with all 9.x agents, but requires version 9.1x of the Enterprise Manager and Workstation.

The Location Map feature does not support metrics from URL groupings. Frontend and called backend elements which have been defined using URL groupings may not appear on the Location Map.

CA Technologies Internal Use Only. Authorized Release Documentation.

Open the Location Map

To view the Location Map from the Browse tree:

1. In the Browse tree, expand the host node.
2. Expand the node corresponding to the frontend application being monitored.
3. Select the agent node.
4. Select the Location Map tab.

   The viewer pane displays a map of the agent's infrastructure, decorated with alerts.

To view the Location Map starting from the Application Triage Map:

1. Select any frontend or called backend in the map.
2. Right-click the element and select "Show Locations for <Element_Name>..."
3. In the Locations table in the lower pane, which lists each of the agents reporting data for that element, select one of the agents and right-click its table row.
4. Select "View This Location."

   The display changes to show the Location Map tab in the Browse tree.

   Note: the first time you open the Location Map, a delay may cause only the Overview tab to be displayed. In this case, manually click the Location Map tab.

With the Location Map tab selected, the viewer pane displays a map of the agent's infrastructure, decorated with alerts (once they have been configured).

Note: The Location Map is not available in Historical Mode.

Understand Location Map Views

The Location Map's appearance changes depending on whether the Enterprise Manager is receiving data from sources outside the Enterprise Manager through APM Catalyst Connector.

Note: The alert indicators appearing on the Location Map are not the same as Triage Map alerts:

- It is not possible to set sensitivity for these alerts.
- They signify only whether the summary metric for a map element has exceeded a threshold.

CA Technologies Pre-Release Documentation. Authorized Use Only.

Two-tiered view

The Location Map displays two tiers when one of these is true:

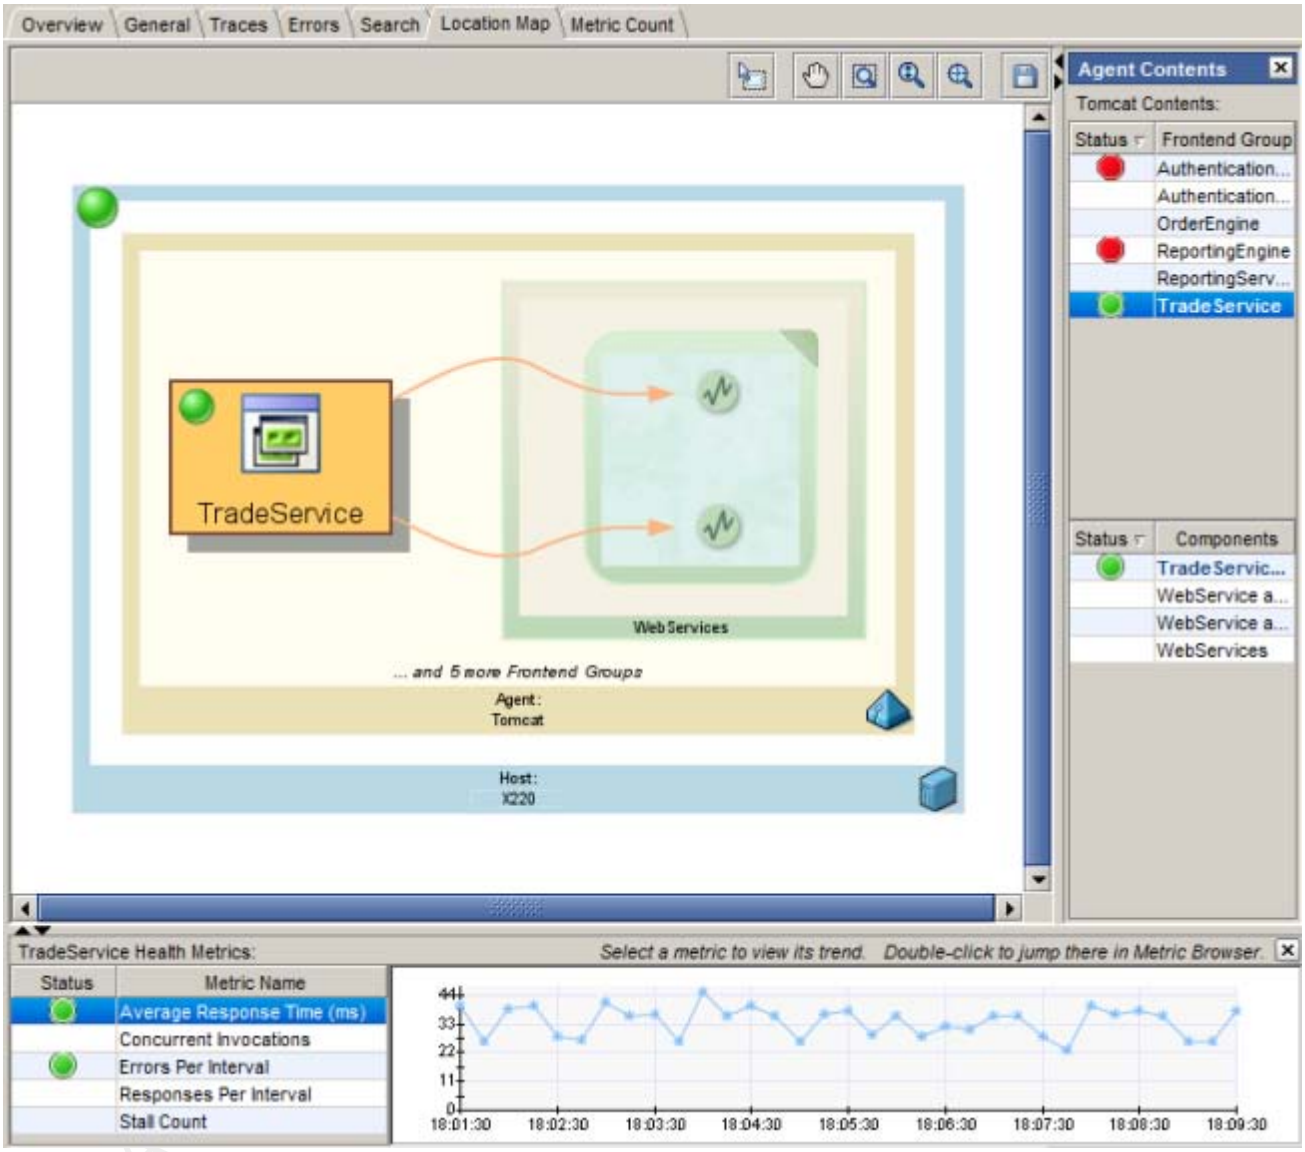
- There is no APM Catalyst Connector present.
- The application is not deployed in a virtual environment (such as a JVM).

The two tiers are:

- agent
- host where the agent is deployed

CA Technologies Prerelease Documentation. Authorized Use Only.

The illustration below shows a simple Location Map for the TradeService frontend as it appears when the frontend is selected through the Browse tree.



The map shows TradeService and the web service backends in the context of the agent, named Tomcat, which is on a host named X220.

Things to notice:

- The Location Map is displayed on the Location Map tab.
- Each rectangle is labeled on its bottom side with the name of the container.

- Introscope alert indicators for the agent and host appear on the upper-left corner of the rectangles. These alerts are summary alerts for the container's resources; in other words, the alert shows a worst-case status of the resource alerts that have been configured on that container. These statuses are distinct from the alert status of frontends or called backends.
- Where alerts are configured on frontends and called backends, the alerts appear in the map display. Otherwise, plain icons denote frontends and called backends.
- Where an agent contains more than one frontend, the legend "... and *n* more Frontend groups" tells how many other frontends reside on the agent.
- With the Agent rectangle selected, its contents are displayed in the details pane on the right.
- Health metrics for the agent appear in the lower pane.

To see the true source of the metric:

- Hover your cursor over the metric name in the table in the bottom pane.

  You do this to see the original "source" of the metric, as the names appearing in the table are standardized.

Three-tiered view

The Location Map displays three tiers when both of the following are true:

- The application is hosted on a virtual machine.
- You have the APM Catalyst Connector importing data from outside the APM system.

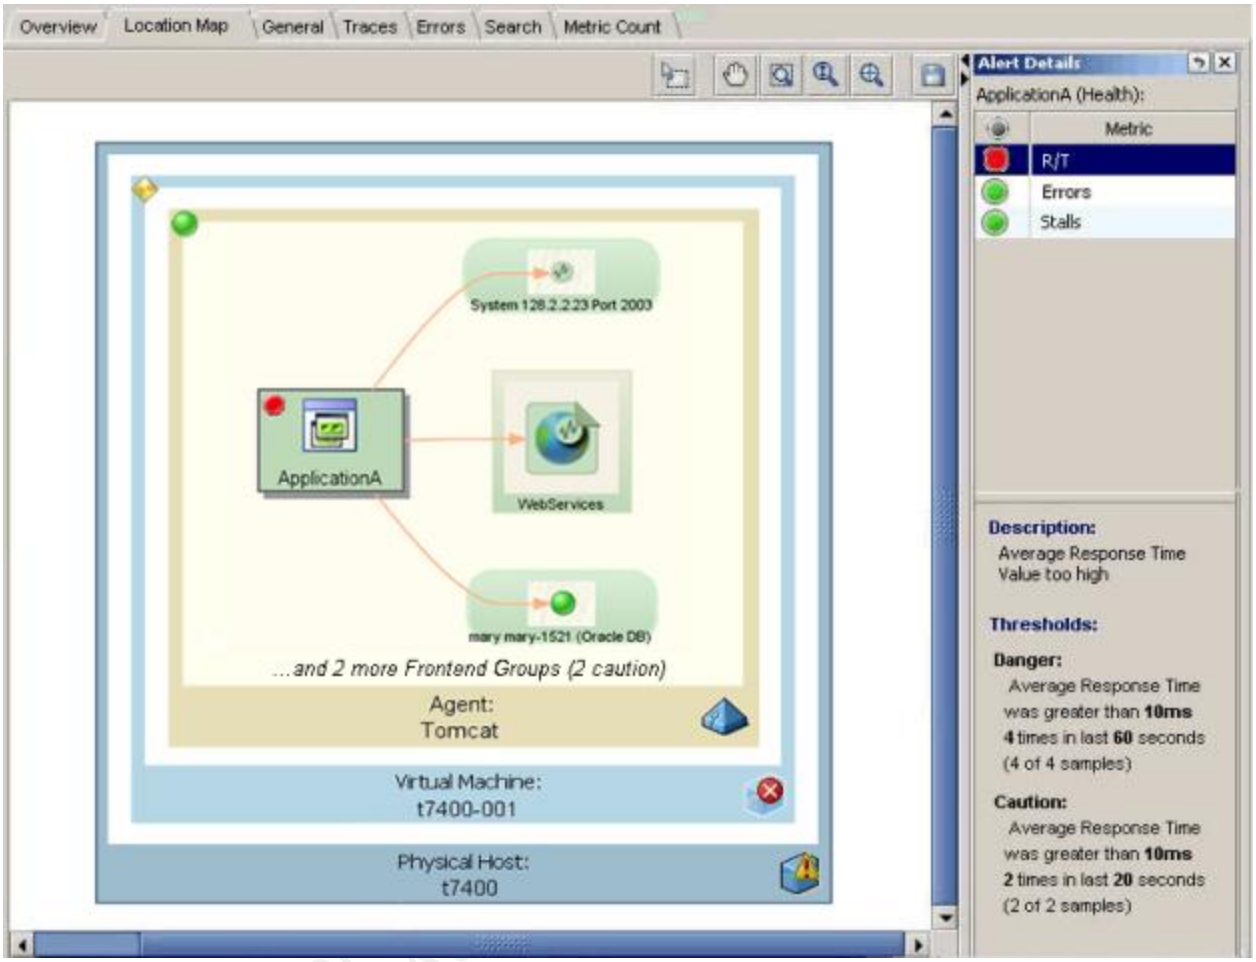
  The APM Catalyst Connector imports data from applications whose health is being monitored by other CA Technologies products such as eHealth, Spectrum, NetQOS and Insight DB. To install the APM Catalyst Connector and configure CA APM and the Connector to communicate, download the APM Catalyst Connector from the CA Technologies software download site and follow the instructions in the accompanying documentation.

The three tiers are:

- agent
- virtual machine
- physical host

CA Technologies Prerelease Documentation
Authorized Use Only

The next illustration shows the same application pictured above, only in the three-tiered view.



Things to notice:

- The frontend ApplicationA is selected in the Metric Browser tree, and is also the focus of the Location Map display, where it appears in relation to its called backends.

- Alert indicators:
  - Where alerts have been configured on frontends and called backends, they appear in the display of dependencies.

    However, because you are viewing the element in the context of only one location, an alert (such as the Danger alert on ApplicationA in the illustration) is not identical to a similar-appearing alert in the Triage Map. In the Location Map, these alerts are configured on the fly using thresholds and current metric values.
  - Alert status imported from [set value for the rose variable at the book level] appears as indicators in the lower-right corners of the Virtual Machine and Physical Host rectangles.

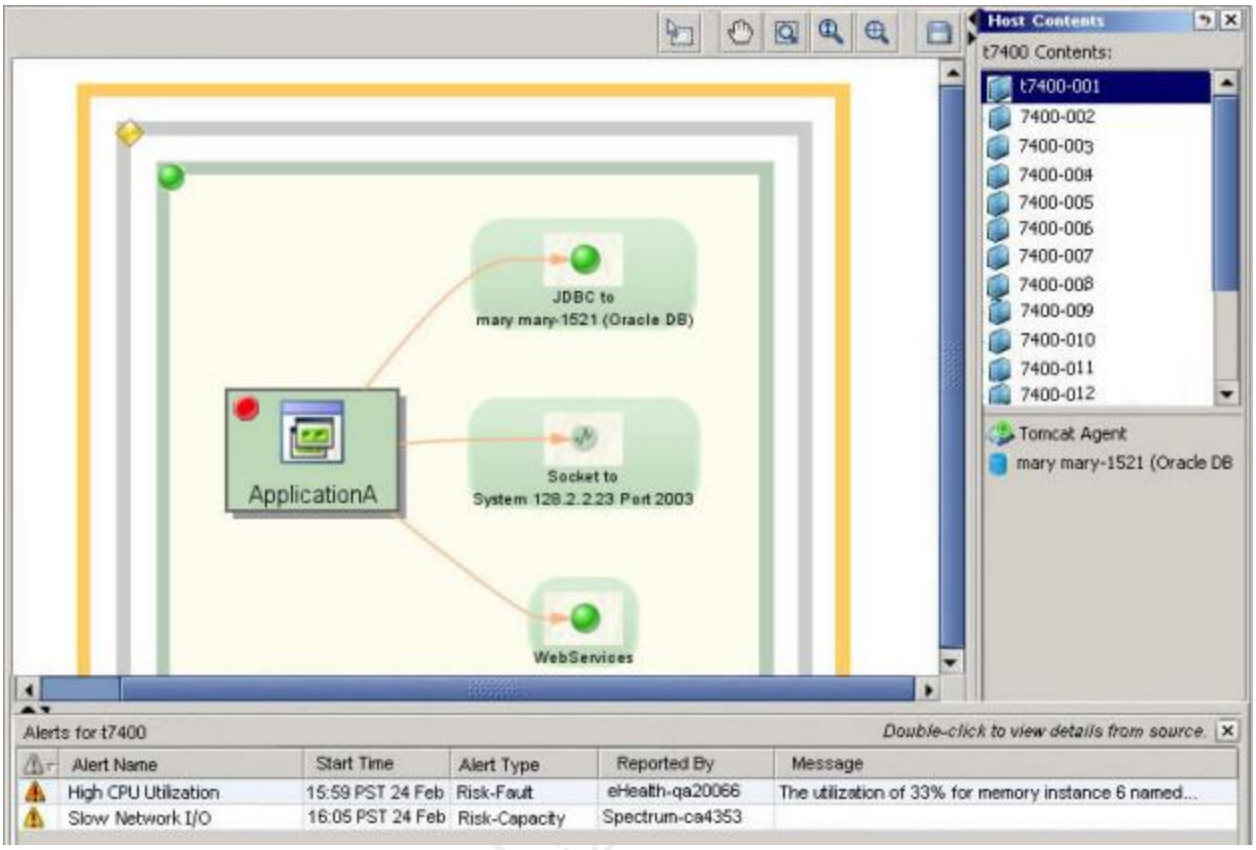
    NOTE: [set value for the rose variable at the book level] alert indicators are different from Introscope alert indicators. See [set value for the rose variable at the book level] Alert Indicators (see page 38).
- Web services:
  - All web services backend calls are aggregated into a single summary element.
  - The web services element, as shown in the above illustration, is decorated with a dog-ear corner, indicating it can be expanded to show individual web service calls.

    To view individual calls contributing to the summary web services called backend:

    – Double-click the dog-ear corner.

CA Technologies Prerelease Documentation. Authorized Use Only

The next illustration shows the same three-tiered view with the host rectangle selected.



Notice that with the host rectangle selected:

- The Physical Host Contents pane lists the Virtual Machines found on that Physical Host .
- Alerts for the host are displayed beneath the map. Notice they are reported by systems outside CA APM.

See More Information About Location Map Elements

You can see more information about one of the rectangles or one of the map nodes.

To see details about any of the elements in the Location map:

- Hover your cursor over the element.

  A tooltip shows more information.
- Right-click the element to see a context menu with several choices:
  - View the triage map for the selected frontend.
  - View the Location Map for the selected item.

- Show metric or alert information for the selected element.
- Jump to the corresponding node in the Metric Browser tab, where relevant.
- Double-click the element to open the Contents pane (see below)..

The Contents pane

To see details about an alerting element:

- Double click the element to open the Contents pane.

For hosts, the pane displays a "Contents" list of Agents, databases, and/or Virtual Machines that reside on that host. For example, if you select the Physical Host rectangle, the contents pane shows a list of virtual machines and databases that reside on that physical host.

If you selected a frontend or backend call, the Contents pane reverts to showing its default content (Agent Contents).

To see the components contributing to a frontend group alert:

- Click the alert in the Alert Details pane and the lower pane will display its contributing components.

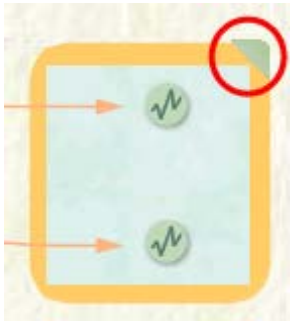
Alert indicators on a Location Map element

The presence of an alert indicator indicates the availability of more information about the health of that element. You can:

- Hover over the alert indicator to see a tooltip with more information.
- Right-click the alerting element to see a context menu with several choices.

To see the individual WebService calls that contribute to the WebServices element:

- Double-click the "dog-ear" control, circled in the illustration below:



Note: Individual WebService calls inside the WebServices container will not have their own alerts. Only the summary WebServices element will have an alert (because it derives from the WebServices backend call alert in Triage Map).

CA Technologies Pre-release Documentation. Authorized Use Only.

Other Details About the Location Map Display

Aging in the Location Map

When agents stop reporting metrics, a map which previously showed live data displays its elements and alert indicators in a gray color, and the corresponding elements in the metric browser tree are labeled in a gray font.

Caching of some Location Map data

Data in the Location Map reflects the content obtained from [set value for the rose variable at the book level] via Restful interfaces. While the most critical data (alerts, virtual hosts, relation between virtual host where application is running and its physical host) are always updated directly from [set value for the rose variable at the book level], some data that do not change frequently are cached by the Enterprise Manager. The size and expiration frequency of the cache is set in the following properties, which are found in the file *<EM_Home>*/config/Catalyst.properties:

- catalyst.entity.cache.size
- catalyst.entity.cache.expirationsec

The Location Map display is updated from the cache whenever the map is opened.

If Catalyst becomes unreachable

If a current [set value for the rose variable at the book level] Service connection has been lost by the Enterprise Manager and cannot be restored automatically, a message at the top of the map will notify you: "Infrastructure alerts are temporarily unavailable; display may be out of date."

When you dismiss the banner, the Workstation will attempt to reconnect; if reconnection fails, you will *not* be notified further. The message appears only when an existing connection is lost.

Administering agent connections from the Workstation

You can issue commands directly from the Workstation to unmount or shut off agents or individual metrics.

When an agent is deployed on an application server, it automatically starts when the app server starts, and appears in the Metric Browser tree under the Enterprise Manager to which it reports metric data. When the agent appears in the tree, it is said to be *mounted.*

CA Technologies Pre-release Documentation. Authorized Use Only.

When an app server goes down, the agent automatically stops reporting data to the Enterprise Manager. This agent is said to be *disconnected,* and will appear in the Metric Browser tree as grey and dimmed rather than colored.

A disconnected agent will still appear mounted in the Metric Browser tree, and you can still browse the metrics it reported before it disconnected. If you want to remove it from the Metric Browser tree, you must *unmount* the agent.

To unmount an agent:

1. Right-click a disconnected agent.
2. Choose Unmount <*Agent_Name*>.

   The Agent will disappear from the Browse tree.

If you want to view the historical data stored in the SmartStor database for an agent that has been unmounted, you can remount the agent so it appears again in the Metric Browser tree.

To remount a disconnected agent:

1. Choose Manager > Mount Agent.

   An Agent Chooser dialog appears.
2. From the list, select an agent to remount.
3. Click OK.

The Metric Browser tree displays the disconnected agents, and you can browse the data stored in the SmartStor database.

In the case where you want the Enterprise Manager to stop storing data from an agent that is still running, you can stop data collection without stopping the app server by choosing the Shut Off command.

Note: The Shut Off command does not actually shut off the agent; it shuts off the connection between a running agent and the Enterprise Manager.

To shut off the connection to a running agent:

1. Right-click a connected agent.
2. Choose Shut Off "<*Agent_Name*>".

The agent continues to run on the app server as long as the app server runs, but the Enterprise Manager no longer is connected to it and no longer stores metric data for it.

CA Technologies Prerelease Documentation. Authorized Use Only.

After you shut off the connection to an agent, you can turn on the connection again.

To turn on the connection to a shut off agent:

1. Right-click a shut off agent.
2. Choose Turn On All Agent Components.

The connection between the agent and the Enterprise Manager will be reopened, and the agent will begin reporting data to the Enterprise Manager. Note that you will have to wait up to 30 to 45 seconds for data to begin appearing in the Workstation.

Note: The Turn On All Agent Components command works only if you have previously shut off the agent connection through the workstation using the Shut Off "*<Agent_Name>*" command.

How User Permissions Affect What You Can View

This section contains information about how permissions affect what you can view under the triage map tab's application-centric view and the metric browser tab's agent-centric view.

What each Workstation user sees depends on the permissions they have been assigned by the CA APM administrator. The permissions are available only when an administrator has configured them using Embedded Entitlements Manager. For more information about Embedded Entitlements Manager, see the *CA APM Security Guide.*

In addition, to appear in the application triage map, applications must be configured using version 9.0 and later agents.

The following notes apply to the behavior of all application triage map views:

- Users with admin or SuperDomain privileges have permission to see all Frontends, Business Services, and metrics.
- If an administrator has changed user permissions to view applications or parts of applications, these changes will not be reflected in the application triage map until the user logs out and logs back into Workstation.

Triage Map Tab Viewing Permissions

The following displays in the triage map tab are based on your permissions:

- Application triage map displays of Frontends and their dependencies.
- Application triage map displays of Business Services, Business Transactions, and Business Transaction Components.
- Contents of the application-centric tree in the By Frontend node.
- Contents of the tree in the By Business Service node.

CA Technologies Pre-release Documentation
Authorized Use Only

In some cases, the application triage map will display a dependent element, but if you do not have permission to view that element, you will not be able to select it in the map or view any data reported by that element. For example, if AppA calls AppB, and you only have permission to view metrics on AppA, then you will see a node for AppB, but if you do not have permission to view data for that object, you will see the following message when you hover your mouse over it: "Access to this object requires additional permissions."

If you don't have permission to view data sent from a certain agent (that is, from a certain physical location where an application is running), the agent will not be included in the list of physical locations (see List of physical locations (see page 86)).

However, the aggregated application metrics include data reported across the application, even if you do not have permission to view metrics on some of the contributing agents. You will be able to access the aggregated data in application triage map tooltips, for example (see Tooltips (see page 70)).

Metric Browser Tab Viewing Permissions

Contents of the metric browser tab are based on user domain permissions:

- Users with SuperDomain permission (at least read permission) see all domains for that Enterprise Manager in the agent-centric tree.
- Users with permissions for multiple domains see domain information for those domains in the agent-centric tree.
- Users with permissions for only one domain do not see domain information in the agent-centric tree; they only see the folders for metrics and Management Modules.

Views in the Metric Browser Tab

With the metric browser tab selected in the left pane of the Investigator, the views that appear in the right pane vary depending on the resource or metric selected in the metric browser tab's tree. Depending on the type of node selected, you see tabs for one or more of these views:

- General tab (see page 132)
- Overview tabs (see page 132)
- Search tab (see page 140)
- Traces tab (see page 141)
- Errors tab (see page 142)
- Metric Count tab (see page 143)
- Thread Dumps tab (see page 144)

Note: Investigator tab views are available only in the Workstation, not WebView.

CA Technologies Confidential Document - Authorized Use Only.

General Tab

When you select a metric, the General tab shows a graphic view of the metric—either for live data, or for a selected historical period. See Viewing Historical Data in the Metric Browser tab (see page 180) for an explanation of how to select ranges of historical data to view.

For some nodes in the tree, the General tab shows the path to that node object in the Investigator hierarchy. For example, when the Frontends node is selected, the General tab shows this path:
SuperDomain|*HostName*|*ProcessName*|*AgentName*|Frontends

For some other nodes in the tree, the General tab shows the Slowest 10 view of the selected node. For example, when the EJB node is selected, the General tab shows the response times of the top ten called components of the selected EJB node.

Ten slowest or worst metrics

When you select certain resources in the Investigator, the General tab of the Viewer pane shows the ten slowest/worst metrics for the selected resource. Java resources include servlets, JSP, EJBs, and JDBC; for .NET, resources include ASP.NET, ADO.NET, and serviced components.

These metrics appear in a bar chart in the Investigator viewer pane. An example is shown in Bar chart on page 29.

You can also view the response times of the top-ten called components of a selected Servlet, EJB, or JSP for Java, or ASP.NET, ADO.NET, and serviced components for .NET.

If you see fewer than ten bars in the bar chart, it is because there are fewer than ten monitored components under that resource. If the metrics don't contain data, you might see the metric names in the viewer pane but no data bars.

Overview tabs

The Investigator summarizes information in an Overview tab for:

- the overall Application—see Application Overview (see page 133)
- the health of the EM—see EM overview (see page 135).
- data from ASP .NET pages—see ASP. NET overview (see page 135).
- data from EJBs—see EJB overview (see page 135).
- data from application frontends—see Frontend overviews (see page 136).
- data from application backend systems—see Backend overview (see page 136).

CA Technologies Pre-release Documentation
Authorized Use Only

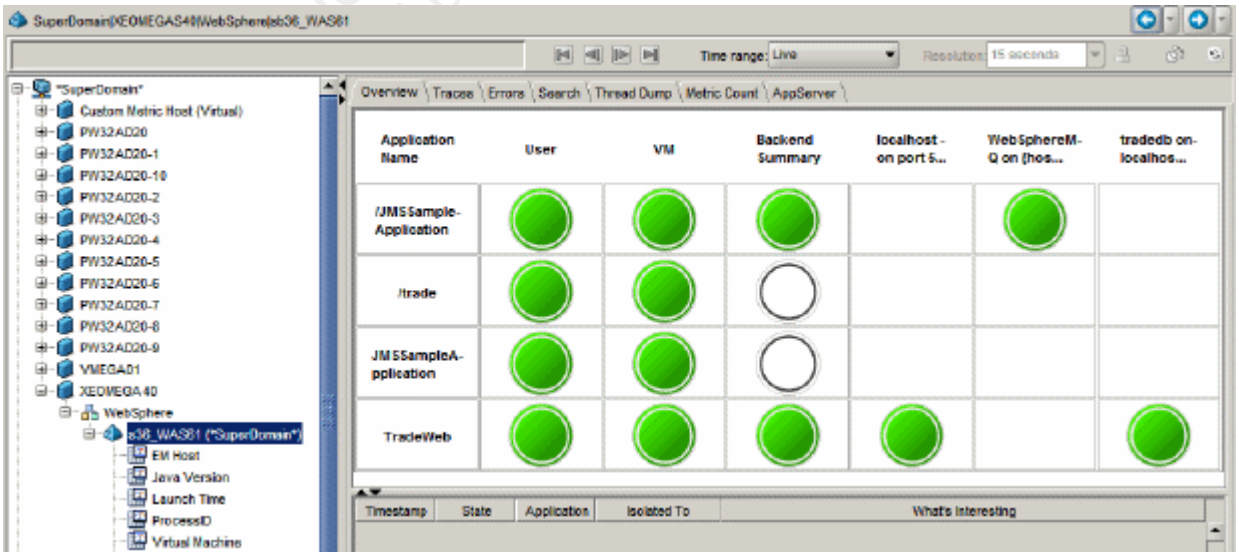
- the garbage collection (GC) heap—see GC heap overview (see page 136).
- instance counts of Java classes instantiated on the JVM—see Instance Counts (see page 137).
- data from JavaNIO—see JavaNIO overview (see page 137).
- data from JTA components—see JTA overview (see page 137).
- data from servlets—see Servlet overview (see page 138).
- socket connections—see Socket overview (see page 138).
- data from struts—see Struts overview (see page 138).
- data on running threads—see Threads overview (see page 138).
- data from XML components—see XML overview (see page 138).
- data from the Leak Hunter extension—see Leak Hunter metrics (see page 145).

Application Overview

The Application Overview is available when you select an agent in the agent-centric tree, and enables application monitoring and triage. It shows high-level health indicators, and a log of related events and historical metric information.

The Overview shows a row of indicators for each application managed by the currently selected agent. Introscope presents this data for each application it discovers—when a servlet executes, Introscope makes a call to *getServletContextName()* of the *ServletContext* interface to determine the name of the application. After the application starts, the Overview tab automatically updates to display a row of indicators for it.

The illustration below shows the Overview tab for an agent on a WebSphere app server named *s36_WAS61:*



This illustration shows four applications -- one in each row of the table -- managed by this agent. For this application, you can view alerts showing the state of:

User: Indicates how satisfactory the end-users' interactions with the application are likely to be. Satisfaction is a function of response time, waits, stalls, and errors.

- Green -- normal, satisfactory user interactions with the application.
- Yellow -- an attempt to use the application is likely to yield unsatisfactory results, for instance poor response time or errors.
- Red -- indicates a serious availability issue and that an attempt to use the application will probably fail.

VM: Indicates the health and availability of server resources, such as resource pools and CPU.

- Green -- normal health of server resources.
- Yellow -- resource limitations or outages
- Red -- serious resource limitations or outages.

Backend summary: Indicates the worst health and availability across all backends accessed by the application. For example, if one of three backends has a serious resource limitation or outage, the All Backends indicator is red. The purpose of the All Backends indicator is to allow the user, with minimal scrolling, to quickly assess whether any of the backends have problems that require investigation.

- Green -- normal backend health and availability across all backends accessed by the application.
- Yellow -- at least one backend accessed by the application is experiencing errors or stalls, or poorer than expected response times.
- Red -- at least one backend accessed by the application is experiencing serious resource limitations or outages.

Backends: Any indicators to the right of the Backend Summary indicator correspond to the individual backends. For information about how Introscope identifies backends, see Viewing metrics for Backends in the Investigator (see page 116).

- Green -- normal backend health and availability.
- Yellow -- backend errors or stalls, or poorer than expected response times.
- Red -- serious backend resource limitations or outages.

The indicators refresh every 15 seconds. The rows are sorted first by color—rows with red indicators precede those with yellow, which precede rows with all green—to reduce scrolling needed to identify potential problems. Within a color category, rows are alphabetized by application name.

CA Technologies Pre-GA Use Documentation. Authorized Use Only.

Using alerts to drill down for more data

You can double-click an alert from the overview tab to display the underlying data for that application tier. For example, if you double-click the User alert, the Workstation will display the URLs node for that agent.

Application overview metrics in historical mode

Over a historical range, an alert color reflects the worst-case value of the heuristic at any point in the historical range. For example, if at any time during a historical range the User heuristic for an agent was yellow, but never red, the Overview tab for that historical range is yellow.

Application overview metrics for a virtual agent

For Virtual Agents, heuristics are evaluated on the basis of Virtual Agent metrics. For this reason, the Overview tab for a Virtual Agent might indicate a different value than for the physical agents in the Virtual Agent.

For example, the Overview tab for a Virtual Agent could display a green User alert, even though the Overview tab for one of the agents in that Virtual Agent shows a yellow User alert.

Heuristic metrics are only generated if the metrics they analyze exist. So, for example, if the Virtual Agent is configured not to include CPU, JMX, or WebSphere PMI metrics, there is no VM folder and the VM alert remains gray.

For information about configuring Virtual Agents, see the *APM Installation and Upgrade Guide*.

EM overview

You can view a variety of metrics on the Enterprise Manager itself by selecting the EM node under Custom Metric Agent.

ASP. NET overview

In environments where Introscope is monitoring a .NET application, an ASP .NET node on the agent-centric tree allows you to monitor metrics for application components.

EJB overview

The EJB (Enterprise Java Beans) overview shows statistics for Entity beans, Session beans, and Message Driven beans.

CA Technologies Pre-Release Documentation
Authorized Use Only.

Frontend overviews

Overviews for Frontend nodes show graphed application metrics, and statistics related to transactions in the application:

The programs which the Investigator displays under the Frontends node represent the component of an application that first handles an incoming transaction. For more information, see Frontends and backends (see page 114).

Backend overview

Overviews for Backend nodes show graph views of database metrics and a table view of SQL below the node.

GC heap overview

The garbage collection (GC) heap overview shows heap use.

GC Monitor Overview tab

Clicking the GC Monitor node in the Metric Browser tree causes the GC Monitor Overview tab to be displayed in the viewer pane. The Overview tab displays three panes:

- Top: an alert indicator on the Percentage of Java Heap Used metric for the JVM.
- Middle: a tabular view of garbage collectors on the JVM
- Bottom: a tabular view of memory pools on the JVM

NOTE: The alert indicator in the top pane of the Overview tab, and colored shading that appears in table cells in the middle and bottom panes, are based on preset caution and danger thresholds. Users cannot reset these thresholds.

When you select any of the individual Garbage Collector or Memory Pool nodes, graphs display the same metrics shown in the Overview tab.

For more information:

- See definitions and thresholds for each of the GC Monitor metrics (see page 359)
- Understand how to use the GC Monitor metrics (see page 164) to tune your JVM's memory allocation

CA Technologies Prerelease Documentation. Authorized Use Only.

Enable/Disable GC Monitor

GC Monitor metrics are enabled by default.

To disable GC Monitor metrics:

1. Open the *IntroscopeAgent.profile* file.
2. Edit the value of the property introscope.agent.gcmonitor.enable from true to false.
3. Save and close the file.

NOTE: This is a hot-configurable property; changes do not require restarting the Enterprise Manager.

For more information about editing IntroscopeAgent.profile, see the *CA APM Agent Implementation Guide* or the *CA APM .NET Agent Implementation Guide*.

Instance Counts

The Instance Counts overview tab shows the classes instantiated on the JVM.

JavaNIO overview

The NIO overview shows tables for datagrams and channels, including client and server metrics. With the JavaNIO node selected, the Overview tab displays general information about the selected node, including all ports with NIO activity.

NIO Channels overview

The Channels node Overview tab displays server and client information for datagrams and sockets.

NIO Sockets overview

The Sockets node Overview tab displays graphs for input and output bandwidth data and concurrent readers and writers data, as well as server and client information for sockets

NIO Datagrams overview

The Datagrams node Overview tab displays graphs for input and output bandwidth data and concurrent readers and writers data, as well as server and client information for datagrams.

JTA overview

The JTA Overview tab displays data about JTA components.

CA Technologies Pre-release Documentation. Authorized use Only.

Servlet overview

The Servlet overview shows a table of servlets in the node. When you select a servlet, the Investigator shows its statistics in a graph.

Select an individual servlet to see its Overview summary tab:

Socket overview

The sockets overview (not to be confused with the NIO Sockets overview (see page 137)) shows tables for client and server sockets, and socket information for each port. With the Socket node selected in the agent-centric tree, the viewer pane on the right side displays all the ports with active sockets. Selecting a port in the Server table at the top of the viewer pane displays that server's client ports in the Client table at the bottom. Selecting a port in the agent-centric tree displays metric graphs about events and load.

Struts overview

The Struts Overview tab shows an overview of Struts components, with a display of the average response time for all components.

Selecting one of the component nodes shows an overview of the metrics for that node.

Threads overview

The Threads overview shows all active threads being processed through an agent.

XML overview

The Overview tab for the XML node displays metrics for XML components.

Heuristics and Metric Baselines

CA APM determines the color of an alert indicator in the Overview tab evaluating current metrics against a baseline for those metrics. With an agent node selected in the agent-centric tree, the Heuristics node shows the metric values related to these indicators.

Baselines are calculated using a statistical algorithm that has been successfully applied in domains such as sales forecasting and weather forecasting. For a given metric, the baseliner algorithm determines the next expected value and the expected deviation from that value. If the actual deviation exceeds (2x), or significantly exceeds (4x) of that expected deviation, the baseliner indicates a moderate or severe violation, and an associated heuristic turns yellow or red.

CA Technologies Prerelease Documentation. Authorized Use Only.

Internally, the baseliner evaluates the slope of the time series, and determines the expected value of the slope. Recent data is given more weight than older data.

Note: While the Enterprise Manager polls for metric data every 15 seconds, the baseliner logic runs only every 60 seconds. This means that during the 60-second interval, Enterprise Manager will poll for heuristic data and report an unchanged heuristic value which can only be updated at the end of the 60-second interval.

The baseliner has a notion of periodic seasons, time intervals during which we expect environmental conditions to repeat. During the first week that a baseliner is active, current values are compared against measurements taken on previous days, with weekdays and weekends distinguished from each other.

Example

Suppose the Enterprise Manager is started on Thursday at noon. During the first 24 hours the baseliner compares current values against data from all 24 hours, with more recent data more heavily weighted. Starting Friday at noon, current data is compared against data measured during the same 30 minute period on previous weekdays. So, on Tuesday at 3:15PM, current data is compared against data on Thursday, Friday, and Monday between 3:00PM and 3:30PM.

Weekend data is only compared against itself. On Saturday the baseliner learns from scratch, and on Sunday current data is compared against data from Saturday.

After the first week we switch from a daily season to a weekly season. So, in our example, starting on Thursday at noon we begin comparing current values against 30 minute periods from the same time in previous weeks. Over time, an increasing amount of historical data improves the quality of the baseline data and the analytics.

Some benefits of seasonality

Scheduled downtime is not supported in baselines, but baseline seasonality compensates for this in cases where the scheduled downtime occurs regularly. For example, if scheduled downtime occurs from 2 a.m. to 3 a.m. on Sunday morning, the baseliner learns to expect strange values during this time, but those values are not expected at other times of the week.

Abnormal data might pollute the baseline temporarily—the baseliner could slowly learn that abnormal data is typical. However, abnormal data would need to be sustained for a long time, and in seasonal mode (after the first day) the baselines are even more robust against this. The baseliner looks at expected values over 30 minute periods in previous seasons, so unless a problem persists for many days or weeks the baseliner expects good, normal activity.

CA Technologies Proprietary Documentation
Authorized Use Only

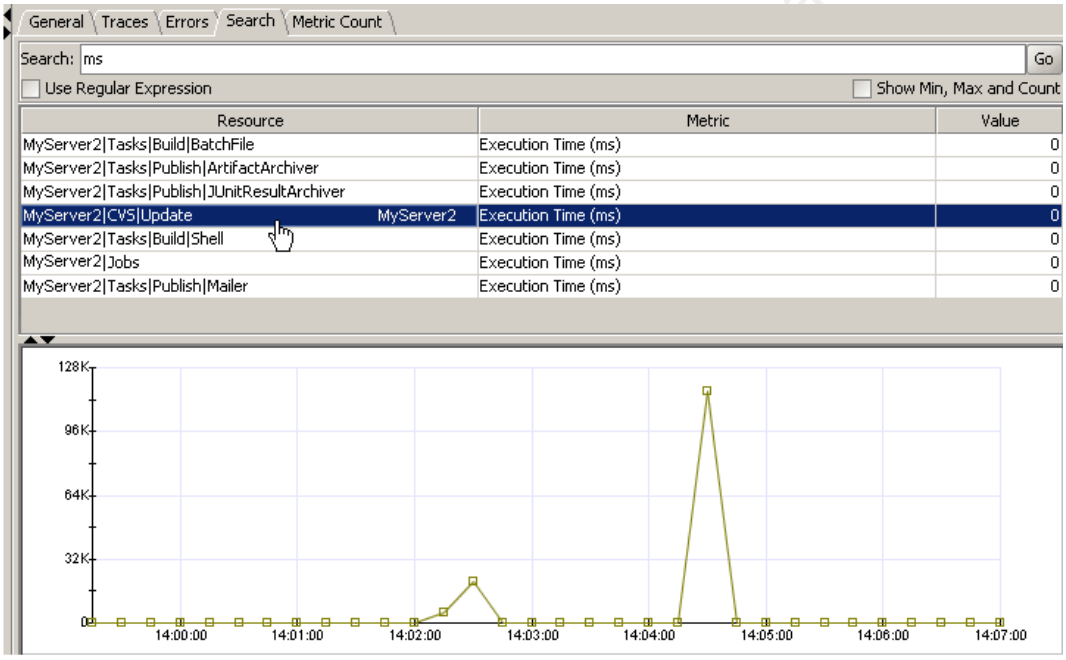
Other tabs

In addition to the Overview tabs, other tabs include the Search, Traces, Errors and Metric Count tabs.

Search tab

The Search tab is available when you select a node in the agent-centric tree that contains metrics. It enables you to quickly find metrics.

The illustration below shows how the Search tab appears in the viewer pane.



Things to notice:

- The node selected in the agent-centric tree sets the scope of a search. For example, if you select Frontends in the tree, search will search only the resources under that node.
- You can enter either a string or a regular expression in the Search field.
- If you enter a regular expression, check the Use Regular Expression box.

Note: Regular expressions cannot filter by agent, so it is not possible to search for agent name.

The right pane lists the resources with metrics that match the search argument, and the value for each. To display Min, Max, and Count columns, click Show Min, Max and Count.

If you click a metric in the list, a view appears in the bottom of the right pane.

If you click on a different node that contains metrics, the search argument used in the previous search remains active, and is applied to the newly selected node.

For information on how to use Search, see Using search (see page 181).

What's Interesting events

The lower half of the Overview lists *What's Interesting* events, which Introscope generates automatically when the color of an alert indicator changes to yellow or red. In Live mode, the previous 20 minutes of events appear.

For each selected item, you can see:

- Timestamp—time at which alert indicator changed to yellow or red
- State—The state of the alert, identified by color
- Application—the application for which the alert displays status
- Isolated to—the tier associated with the state change
- What's Interesting—a description of what drove the state change—for example:
  `The number of errors in /pipeorgan's User tier is unusual. The current value is 28, while the typical value is 4.`

  This information also appears in the What's Interesting View tab, as shown.

Also notice the tooltip that appears when you mouse over one of the alerts in the What's Interesting table.

Traces Tab

The Traces tab, available when a resource or component is selected in the agent-centric tree, is similar to the Transaction Tracer (see Using the Transaction Tracer (see page 209)). The Traces tab lists the recorded Transaction Trace events for the selected resource or component.

Note: The default time range for traces in live mode is 20 minutes. Traces older than 20 minutes are not displayed in live mode; they will be aged out (not shown) after they are more than 20 minutes old.

CA Technologies Preliminary Documentation. Authorized Use Only

Setting the Duration Unit

By default, the Traces tab shows the duration of transactions and transaction components in milliseconds (ms), thousandths of a second.

You can change this unit to:

- seconds
- microseconds (ms)

To change the unit for the Duration column in the Traces tab:

1. Right-click the Duration (ms) column heading.
2. From the drop-down menu, select one of:
   - seconds
   - milliseconds (default)
   - microseconds

The Traces tab displays the new unit in the column heading, and renders the duration using the new unit in all transaction views (including in the Transaction Trace Viewer—see Using the Transaction Trace Viewer (see page 209)).

Errors Tab

The Errors tab, available when a resource or component is selected in the agent-centric tree, lists errors and error details for the selected item. The Errors tab allows support personnel to detect and diagnose the cause of serious errors as they occur, determine the frequency and nature of the errors which can prevent end users from completing web transactions, and deliver specific information about the root cause to developers.

Note: You must have ErrorDetector enabled to see the Errors tab. For information about enabling ErrorDetector, see the *CA APM Java Agent Implementation Guide.*

The top half of the Errors tab lists the time, description, and type of each error. The lower half of the tab shows detailed information for each component involved in the error selected in the list above.

More information:

ErrorDetector (see page 406)
Viewing errors with Transaction Tracer (see page 216)

CA Technologies Pre-release Documentation. Authorized Use Only.

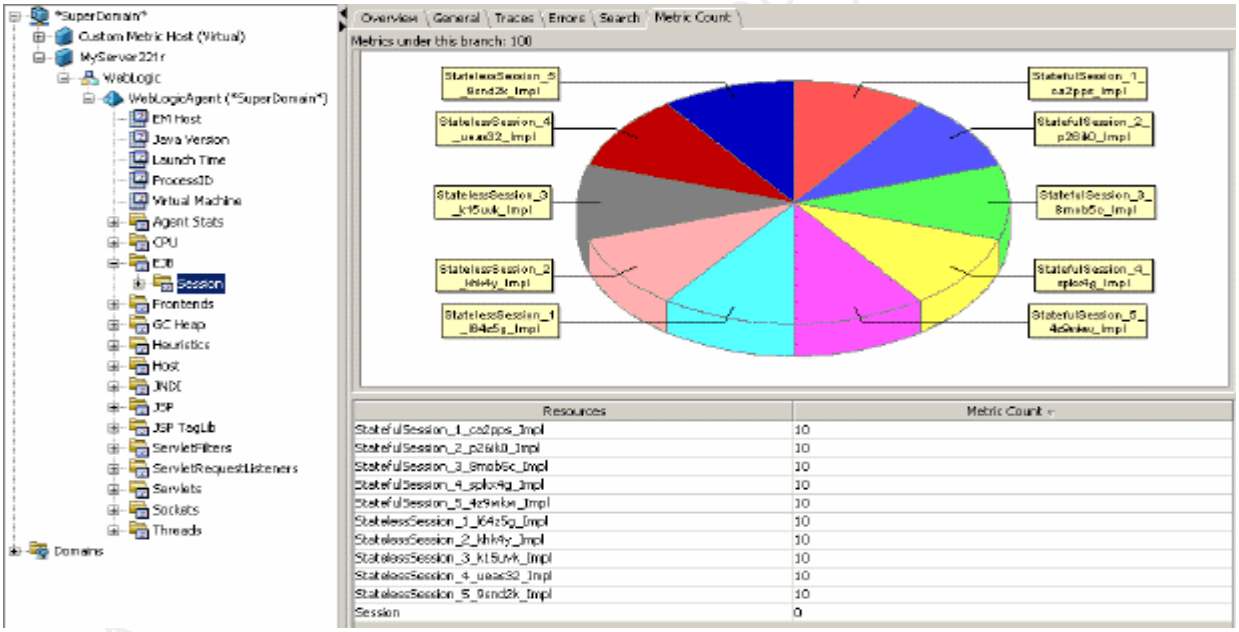
Metric Count tab

Many of the nodes in the agent-centric tree have a Metric Count tab, which displays a pie chart of the metric distribution for the node.

The illustration below shows the pie chart, with a table display of the same data beneath it.

The pie chart displays a maximum of 50 slices. When there are more than 50 resources in the selected node:

- The pie displays the resources reporting the 50 highest values.
- In addition to the slices representing the 50 highest values, an additional slice will be labelled "All Other Metrics" to show the proportion of metrics with data outside the top 50 reported.
- The status bar displays the message "Displaying the top 50 resources. Remaining resources grouped in "All Other Metrics"."



Mousing over an area of the pie chart displays a tooltip with count and percentage.

Long labels will be truncated, but when you select a slice of the chart, the fully qualified name of the resource will appear in the table beneath the chart.

The Metric Count tab is available in both live and historical modes.

Resources Tab View

The Resources tab shows graphs of the Resource metrics (see page 372). The Resources tab is viewable in the Metric Browser tree, when the agent is selected.

Note:

- Though graphs for all of the Resource metrics will appear in the Resources tab view, graphs will display no data when those metric classes are unavailable for the agent.
- The original source of the metric that appears in the Threads in Use and JDBC Connections in Use graphs will vary from agent to agent, depending on the agent's type (WebLogic, Tomcat, .NET, or others).

In the Metric Browser tree they appear under the agent node as follows:

```
<Agent_Name>
|
|--Agent Stats
   |
   |--Resources
      |
      |--<metrics>
```

-

Thread Dumps tab

Each agent node on the metric browser tree has a Thread Dumps tab. This tab allows you to collect Java thread dumps (thread dumps) and display current and historical thread dump data. A thread dump provides information about all the threads running inside a JVM at one point in time. For each thread, a thread dump provides the thread name and ID; state; and a stack trace, which lists all the methods called.

The Thread Dumps tab includes these parts:

- The header displays the time of the thread dump.
- The thread dump summary bar displays the total number of threads and number of threads that are waiting, blocked, and running.
- The search pane allows you to search for a specific string within all the thread dump information. The results display in the thread information table.
- The threads state drop-down list filters the thread information table by thread state. When you select a state, the thread information table updates.
- The thread information table displays a list of all threads. For each thread, it provides the thread ID, name, state, and the last method called by the thread immediately before the thread dump.
- The thread stack trace table displays all the methods in the order called.

CA Technologies Prerelease Documentation. Authorized Use Only.

- The % Threads by State pie chart displays the threads in these states: deadlocked, blocked, running, or waiting.
  - Hovering over an area displays a tooltip with the number and percent of threads in each state.

The Thread Dumps tab is viewable in the metric browser tree when you have selected an agent node.

Note: If you are triaging agent issues, view the <Agent name> | Threads | Deadlock Count metric in the metric browser tree. This metric indicates whether there are deadlocked threads affecting the agent. Introscope configuration is required to enable the Deadlock Count metric. For more information, see the *CA APM Java Agent Implementation Guide.*

You can click the:

- Collect New button to collect a thread dump.
- Save as Text button to save current thread dump to a text file.
- Load Previous button to load a single previously collected thread dump, and to see the time stamp and the associated data.

No thread dump data displays until a thread dump is collected or after an Enterprise Manager restart.

The Thread Dumps tab is available in live mode; no historical thread dump data displays in Historical mode.

LeakHunter metrics

CA APM shows LeakHunter metrics under a LeakHunter resource in the agent-centric tree. The LeakHunter overview shows statistics graphically and in a table. Leak tabs appear for nodes under LeakHunter, and show details of the leak and a graph of the number of collections over time.

LeakHunter produces these metrics:

- LeakHunter:Potential Leak Count
- LeakHunter:Tracked Collection Allocated Rate—Number of new collections allocated per second. This metric is reported to the Enterprise Manager every 15-second interval.
- LeakHunter:Tracked Collection Count—Number of collections examined in your instrumented application since LeakHunter started running until the timeout period expires. After the timeout period expires, LeakHunter stops looking for potential leaks in newly allocated collections, but continues to monitor collections previously identified as potential leaks.

CA Technologies Prerelease Documentation
Authorized Use Only

A resource is created for every potential leak, where <CollectionID> is the assigned unique ID.

Collection IDs provide uniquely identify each individual potential and real leak across sampling periods, providing an easy way to access historical information about specific leaks over time.

Under this resource, these metrics are reported if they are available:

- LeakHunter|<CollectionID>:Allocation Method
- LeakHunter|<CollectionID>:Field Name
- LeakHunter|<CollectionID>:Field Name #<sequential number>

Note: If there are multiple Field Name Metrics, each one is named in sequence, such as Field Name #2, Field Name #3, etc.

- LeakHunter|<CollectionID>:Collection Class
- LeakHunter|<CollectionID>:Allocation Time
- LeakHunter|<CollectionID>:Size
- LeakHunter|<CollectionID>:Currently Leaking
- LeakHunter|<CollectionID>:Allocation Stack Trace

Note: The Allocation Stack Trace Metric is only provided if the *introscope.agent.leakhunter.collectAllocationStackTraces* property in the *IntroscopeAgent.profile* is set to true. See information about configuring LeakHunter in the *CA APM Java Agent Guide.*

Using tooltips to view metric names and values in a Data Viewer

In a Data Viewer, you can hover your cursor over a point on a graph to open a tooltip.

To open a tooltip:

- Mouse over any element in the Workstation metrics tree or in a Data Viewer, such as a point on a graph.

The illustration below displays information about a particular data point in the graph, showing:

- Metric name
- Exact value of the metric
- Min/max values for the metric across the period represented by the data point. (See How time range affects data points, below (see page 147).) Instead of rounding to a value using K for thousand or M for million, tooltips show exact values.

CA Technologies Prerelease Documentation. Author Use Only.

- The count of 15-second intervals represented by the data point.
- The date and time for the data point in the graph.

Pressing F2 while a tooltip is active allows you to click on the hyperlinked text. When you do this, an Investigator window opens with the tree expanded to the metric shown in the tooltip.

Note: For information on tooltips used in the Transaction Trace window, see Tooltips (see page 70).

How time range affects data points

Each data point on a graph represents an equal division of the time covered by the graph. If the time range is set to Live (as in the illustration above), each data point represents a 15-second interval.

If the time range is set to another value, the interval represented by each data point will be different. If the time range is set to two hours, for instance:

- Each data point represents a two-minute interval, or eight 15-second intervals.
- Since there are eight 15-second intervals in two minutes, the count of each data point is 8.

Manipulating the contents of Data Viewers

Data Viewers in the Investigator viewer pane or in a dashboard show data from an instrumented application in a visual form. Data appears in a Data Viewer based on the type of data—for example, metrics appear as graphs, and alerts appear as colored indicators. Data Viewers can display data from a metric, a resource, or an element, such as an alert.

In Data Viewers, you can:

- Display minimum/maximum metric values in a graph (see page 148)
- Show or hide metric data in a graph (see page 148)
- Change the scale of graphs (see page 148)
- Move metrics to the front or back in graphs (see page 150)
- Export data (see page 151)

CA Technologies Prerelease Documentation: Authorized Use Only

Displaying minimum/maximum metric values in a graph

You can configure a graph to show minimum and maximum values.

To show the minimum and maximum values of metrics and metric groupings in a graph:

1. Click the graph in the Console to select it.
2. Show the minimum and maximum values in one of two ways:
   - Right-click the Data Viewer and select Show Minimum and Maximum.
   - Select Properties menu, and select Show Minimum and Maximum.

Note: This change remains in effect only while you view the current dashboard. If you open a new Console or switch to a different dashboard, this setting reverts to the default, which does not show minimum and maximum metric values. To show minimum and maximum metric values by default in a Graph, turn on this option while editing a dashboard with the Dashboard Editor.

Showing/hiding metric data in a graph

If you are viewing the data from multiple metrics in one graph, you can show or hide individual metric data.

To show or hide a metric in a graph:

1. Display a graph in the dashboard in the Console.
2. You can:
   - Show the metric by clicking its check box.
   - Hide the metric by unchecking its check box.

Note: Show/hide metric options are not available when you view graphs or bar charts that are displaying sorted or filtered data.

Changing the scale of graph charts

You can change the scale of graph charts while viewing live data in Workstation, to provide a more readable view. You change the scale of a chart by setting a minimum and maximum value for the chart's data axis.

The chart scaling feature is available only for graph charts in Live mode. It is not available for any other viewer type such as bar chart, top ten, or string viewer.

Note: Scale changes that you make to a chart are temporary—the settings are not saved with the dashboard. When you select a new dashboard or close the Console window, Introscope discards the settings and returns to the scale options that were applied when the dashboard was created.

CA Technologies Prerelease Documentation. Authorized Use Only.

To view the scale of a graph chart:

- Click on a chart to select it, and then:
- Select Viewer > Scale Options, or
- Right-click the chart and select Scale Options from the context menu.

  The Data Options dialog box opens.

Setting the Auto Scale Minimum and Maximum default values provides a more readable view of charts in Live mode.

To rescale using min and max values:

1. Click on a chart to select it, and then:
   - Select Viewer > Scale Options, or
   - Right-click the chart and select Scale Options from the context menu.
2. Enter the minimum and maximum values for the data axis of the graph.
3. Click OK.

   For example, if the chart data values lie primarily between 350 and 550 but the chart value axis shows 0-1000, it might be helpful to set the scale Min value to 300 and Max value to 600 for a better view of the relevant data.

To force minimum and maximum values:

1. Click on a chart to select it.
2. Select Viewer > Scale Options.
3. Select Pin at on both the Minimum and Maximum sides of the dialog, and enter a value for the minimum and maximum points of the data access.
4. Click OK.

Setting Min and Max values for a chart showing live data is risky, however, if there is a chance the data may exceed the values you set.

To avoid this problem, use the Auto Scale option to automatically set the graph to change its scale according to the data it displays.

To rescale using Auto Scale:

1. Click on a chart to select it.
2. Select Viewer > Scale Options.
3. Select AutoScale on both the Minimum and Maximum sides of the dialog.
4. Click OK.

CA Technologies Prerelease Documentation
Authorized Use Only

The resulting chart's data axis is reset based on the data in the chart, as shown in the illustration below. This often results in sharper valleys and peaks in the graph display

You can also set the scaling options to Auto Expand. This option uses 0 as the bottom of the data axis and automatically expands and scales the data axis to display all data for the time range.

To rescale using Auto Expand:

1. Click on a chart to select it.
2. Select Viewer > Scale Options.
3. Choose Auto Expand on both the Minimum and Maximum side of the dialog.
4. Click OK.

Moving metrics to front/back in graph

When a graph contains multiple metrics, it is possible for data points to overlay each other. You can use the Bring to Front or Send to Back options to choose which metric appears at the top of the list of metrics.

Note: The Bring to Front/Send to Back options are not available when viewing graphs displaying sorted or filtered data.

To change the overlap order of metrics in a graph:

1. Open the Console and display a graph in a dashboard.
2. Right-click the label of the metric to change, and choose an option from the menu:
   - Bring to Front (moves selected metric to the top of the metrics listed)
   - Send to Back (moves selected metric to the bottom of the metrics listed)

   The metric moves to the chosen position.

Copying a Data Viewer to the clipboard

You can copy a snapshot of the data in a Data Viewer to the clipboard as a bit-mapped image. You can then paste the image into an email or other document, or any application that accepts bit-mapped images. This is a handy tool if, for example, you want to show data in a Data Viewer to a colleague, or perhaps use it in a presentation.

To copy a Data Viewer to the clipboard:

1. Open a Console and select a Data Viewer
2. Select Viewer > Copy to Clipboard as Image.

Note: You cannot copy multiple Data Viewers.

CA Technologies Prerelease Documentation. Authorized Use Only

Exporting data from Data Viewers

You can take a snapshot of current data in a Data Viewer and export it to a comma-separated values (.csv) file. You can export data from all Data Viewer types except the alert.

To export data from a Data Viewer:

1. In the Console, select a Data Viewer.
2. Select Viewer > Export Data.
3. In the Save dialog box, choose a location to save the *.csv* file and click Save.

The APM Status Console

The APM Status Console allows you to view important status and events for a stand-alone or clustered Enterprise Manager. Its purpose is to allow administrators to monitor and address issues in the health of the Enterprise Managers they administer.

This functionality provides out-of-the-box monitoring capabilities which would otherwise require the administrator to configure alerts on Enterprise Manager supportability metrics.

Note: Only users with apm_status_console_control permission can see the APM Status alert icon and use the APM Status Console. For more information, see the *CA APM Security Guide.*

The following sections explain details about the APM Status Console interface. For more information about how to use the APM Status Console to monitor Enterprise Manager health, see Monitor performance with the APM Status Console (see page 172).

APM Status Console Interface

This section explains the APM Status Console interface. To see how to use the tool to actively monitor the health of your Enterprise Manager, see Monitor Performance with the APM Status Console (see page 172).

CA Technologies Prerelease Documentation. Authorized Use Only

The APM Status alert icon

If you have enabled the APM Status Console, an icon appears in the Workstation Investigator.

To enable the APM Status Console:

- See the topic on configuring APM Status Console clamps and Important Events in the *CA APM Configuration Administration Guide.*
- See the topic on using server.xml to configure Enterprise Manager server permissions in the *CA APM Security Guide.*

This icon has two states, as shown in the illustration below.



On the left is the icon in normal state. On the right is the icon in alerted state.

When the icon appears in the alerted state, one or more of of these events has happened on one or more of the nodes in the Enterprise Manager map:

- A clamp has been activated .
- An important Enterprise Manager event has occurred.
- An agent has been denied access to the node.

Open the APM Status Console

When the APM Status icon appears in the alerted state, you open the APM Status Console to investigate the cause of the alert.

To open the APM Status Console:

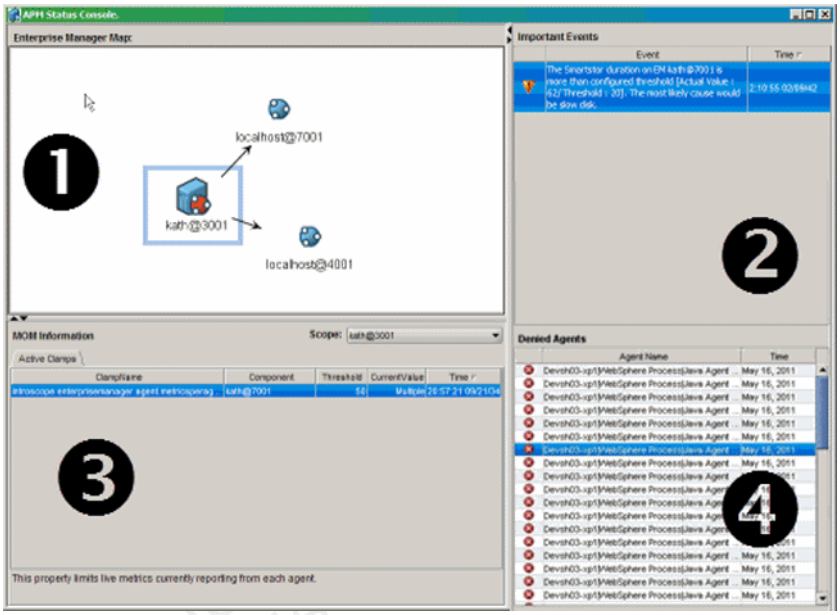
- Click the APM Status icon.

Note: you can also open the APM Status Console by selecting Workstation > New APM Status Console.

CA Technologies Proprietary Documentation. Authorized Use Only.

Overview of the APM Status Console interface

The APM Status Console consists of four panes, marked on the illustration below. Each pane is described in more detail after the list of panes.



1: Enterprise Manager Map

The Enterprise Manager Map displays one of three views, depending on what computer the Workstation is connected to.

- If the Workstation is connected to a stand-alone (non-clustered) Enterprise Manager, that computer alone is displayed.
- If the Workstation is connected to an Enterprise Manager MOM, the map shows a diagram of the MOM and its collector Enterprise Managers.
- If the Workstation is connected to a Cross-Cluster Data Viewer (CDV) Enterprise Manager, the map shows a diagram of the CDV and its connected clusters.

2: Important Events table

See "Use the Important Events table" below for more information.

CA Technologies Proprietary
Authorized Use

3: Information table

The Information table is labeled depending on node you have selected in the map. For example, "APM System Information" or "Collector Information." The pane consists of one tab:

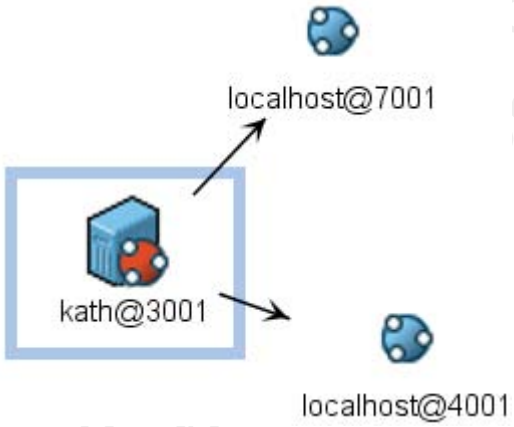
- Active Clamps

See Use the List of Active Clamps (see page 155).

4: Denied Agents list

See The Denied Agents List (see page 156).

Use the APM Cluster Map



The Enterprise Maanger Map displays a stand-alone Enterprise Manager, a MOM cluster, or a CDV cluster. In any of these views, when an Enterprise Manager experiences problems, the map displays that Enterprise Manager's icon in an alerted state.

In the illustration above:

- A MOM Enterprise Manager kath@3001 is shown in alerted state.
- The MOM Enterprise Manager connects to two Collector Enterprise Managers.

To get more information about an alerting Enterprise Manager:

- Select the Enterprise Manager.

  Depending on the problem causing the alert, other panes in the APM Status Console are populated with more information about the problem.

Use the Important Events Table

The Important Events table lists the following events when they pass the thresholds listed here:

- Slow collector response -- The connection is still live; but the EM is not responding because it is busy.
- High CPU consumption -- CPU use is over 80%
- High memory use -- JVM memory is over 90%
- Long SmartStor duration -- Time it takes SmartStor to write data is over 3.5 seconds.
- Long harvest duration -- Time to harvest data is over 3.5 seconds.
- Excessive queued queries -- The number of queued queries is 3 or more.

These thresholds are set in <EM_Home>/config/apm-events-thresholds-config.xml.

Note: The threshold values in apm-events-thresholds-config.xml reflect CA Technologies recommended settings, and changing them is not recommended. For more information, see the *APM Configuration and Administration Guide.*

Things to notice:

- Events are listed with the most recent at the top.
- Events are listed only while they are active. When the metric that caused the event goes below the threshold, the even disappears from the table in the next 15-second interval.
- The events listed in the Important Events table are not logged or saved in history.
- The following actions are not supported:
  - Users cannot change the list of events that appear in the Important Events table.
  - Users cannot change the thresholds for these events.

Use the List of Active Clamps

The Information pane contains the Active Clamps tab. When the user selects an Enterprise Manager or CDV in the Enterprise Manager Map, the Active Clamps tab lists the following information about any clamps that are active on that computer:

- Clamp name
- Component -- Name and port of the Enterprise Manager where the clamp is located

CA Technologies Prerelease Documentation. Authorized Use Only

- Clamp threshold value
- Current value of the data controlled by the clamp
- Time clamp was activated

Things to notice:

- In the case where multiple per-agent clamps have exceeded threshold values, the label Multiple appears instead of a value. Clicking the row opens a table with clamp details. See "Clamp Details table" below.
- Scope dropdown list -- use this to switch the contents of the Active Clamps tab to reflect the active clamps on another Enterprise Manager.

To see more information about the clamp:

- Double-click a row in the table.

  An Investigator window opens to the Metric Browser tree, expanded to show the metric graph for the clamp. This allows you to see the metric trend over time.

  Note: You must have SuperDomain or appropriate domain permission to see metric information for the clamp in the Metric Browser tree.

Clamp Details table

If the clamp the user selects is active on more than one component, the word "Multiple" appears. Clicking this opens the Clamp Details table, which displays the following information.

Affected Component

Name of the agent or other component controlled by the clamp.

Threshold

Threshold value for this clamp.

Current

Current value of the metric or other component being clamped.

Clamped Time

Time the clamp was activated.

The Denied Agents List

The Denied Agents list displays agents that have been denied connection to one of these:

- The Enterprise Manager selected in the APM Cluster Map.
- The cluster, if the MOM Enterprise Manager is selected.

CA Technologies Pre-release Documentation. Authorized Use Only.

Agents are denied connection under one of two conditions:

- The agent is no longer listed in the file loadbalancing.xml on the selected Enterprise Manager.
- The connection to the agents has been lost or intentionally disconnected.

Use of this list is documented in full in the *CA APM Configuration and Administration Guide* in the chapter on Configuring the Agent to Enterprise Manager Network Topology.

Viewing CA APM Customer Experience Metrics in the Workstation

You can use the Workstation to view metrics from CA APM Customer Experience. CA APM Customer Experience TIMs collect real-time transaction metrics and pass them to the Enterprise Manager. Metrics are visible in the Introscope Investigator and Console.

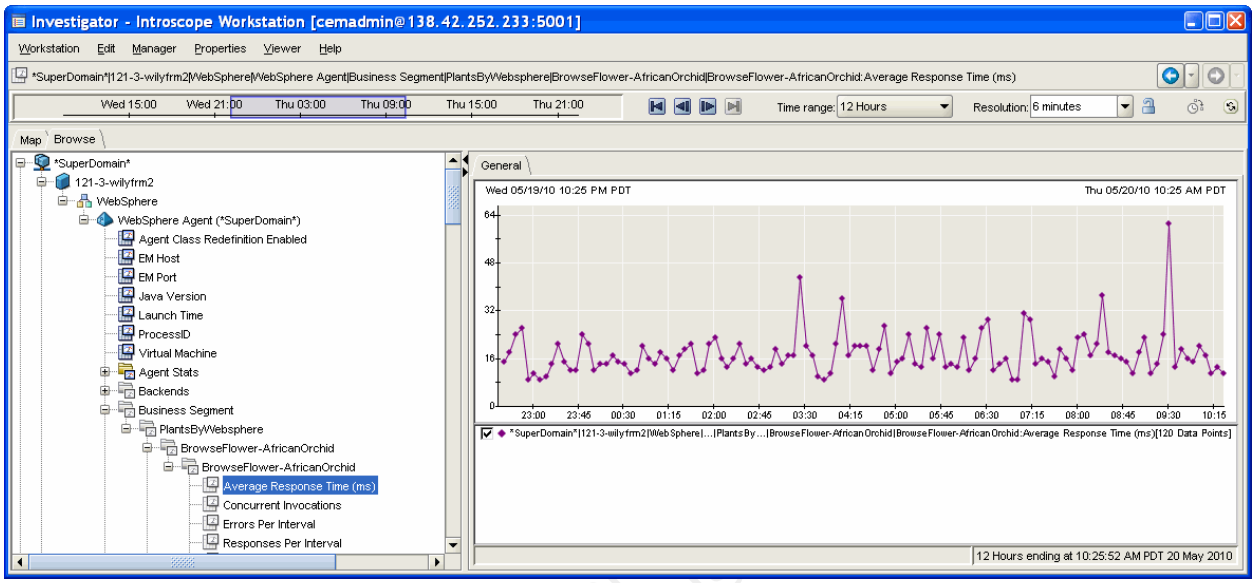
For each aggregated business transaction, the following metrics appear:

- Average Response Time (ms)
- Total Defects Per Interval
- Total Transactions Per Interval

CA Technologies Prerelease Documentation. Authorized Use Only.

Viewing CA APM Customer Experience Metrics in the Investigator

When you select a node corresponding to data imported from CA APM Customer Experience such as the Business Segment node in the Metric Browser tree, the right pane displays the key metrics for that transaction.



You can also navigate to a specific TIM to see detailed metrics.

Viewing CA APM Customer Experience Metrics in the Console

The Console provides you with graphical dashboards so you can display at-a-glance information on real-time transaction health.

CEM Overview Dashboard

The CEM Overview dashboard provides a view into the top 10 business transactions with the highest total defect-to-transaction ratio.

Use the Dashboard drop-down at the top of the Console to select the metric view you need. Double-click on alerts to view the related dashboards.

About metric data values greater than 100 percent

In the Metric Browser tree and the Console, the metric data for business transactions can appear as greater than 100 percent. This can occur when:

CA Technologies Proprietary Authorized Use

- Multiple transaction defects in CA APM Customer Experience are summed under the same business transaction name in Introscope. For example, 3 slow time defects:
  - login.do/login.do
  - login.do/logged in
  - login.do/search.do
- The same business transaction has more than one defect type defined and the same transaction triggers more than one defect. For example, for login.do/login.do, 3 defect types:
  - Slow Time
  - Large Size
  - Low Throughput

Using Wily CEM real-time transaction dashboards

The real-time transaction metrics feature includes default dashboards for you to get started.

- Management Module — CA APM Customer Experience business transaction statistics
- Javascript calculators — tree aggregation and defect percentages
- Investigator typeviews — the default dashboards

Several default dashboards are available to get started, including:

This default dashboard ...	Provides ...
CEM Overview	Top 10 business transactions with the highest total defect-to-transaction ratio.
AVG Response time	Top 10 business transactions with highest average response time (ms).
Defect Count	Top 10 business transactions with the highest total defect count.

CA Technologies Prerelease Documentation
Authorized Use Only.

This default dashboard ...	Provides ...
Defect Percent	Top 10 defect types with highest defect percentage, with detailed dashboards showing top 10 highest defect percentage for: ■ Client request error ■ Content error ■ Fast time ■ High throughput ■ Large size ■ Low throughput ■ Missing component ■ Missing response ■ Missing transaction ■ Partial response ■ Server response error ■ Slow time ■ Small size ■ Unauthorized access
Response Count	Top 10 business transactions with the highest transaction count.

Creating Your Own Dashboards

Once you understand the basics of the real-time transaction metrics and the default dashboards, then you can create custom dashboards that answer the business questions you have. For example:

- How many visitors per hour have been on the new web site?
- What is the real-time status of new orders in terms of ...
  - frequency of orders?
  - errors during transactions?
  - response time during transactions?
- How many orders in the past 10 minutes?
- How many new users are logging in right now?
- What is the transaction throughput trend during the past 24 hours?

CA Technologies Prerelease Documentation. Authorized Use Only.

See Create and Edit Dashboards (see page 278) for information about how you can create custom dashboards of CA APM Customer Experience metrics.

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 4: Monitoring System Performance and Problems

CA APM power users understand that CA APM is best used not only in investigating application problems, but also to monitor nominal application performance. Once you understand what nominal performance looks like for your application, you're better equipped to understand the signs of application performance issues and breakdowns.

This chapter tells you how to monitor nominal application performance, how to respond to notifications that there's something wrong, and how to use the tools in Introscope Workstation to find the cause of the problem.

This section contains the following topics:

Understanding nominal performance (see page 163)
Reading and understanding notifications (see page 173)
Respond to a Notification (see page 176)
Diagnose the Problem with the Metric Browser Tab (see page 179)
Diagnosing Problems with CA APM Customer Experience Transactions (see page 185)

CA Technologies Proprietary Documentation
Authorized Use Only

Understanding nominal performance

Understanding what normal application performance looks like builds familiarity with your system and gives you a broader context to understand the inevitable problems. It also builds your familiarity with Introscope tools and utilities. When something goes wrong, you'll have the background knowledge to dig in and find the problem.

Three different nodes in the Introscope Workstation Investigator tree are especially helpful in enabling you to monitor application performance. These nodes—GC Heap, Frontends, and Backends—might be thought of as your application's vital signs.

Monitoring performance with the GC Heap metrics

Garbage Collection is the process of freeing memory taken up by objects no longer in use; once memory is freed up it is usable by other objects. The GC Heap (Garbage Collection heap) metrics provide a good tool for monitoring and understanding application performance.

GC Heap|Bytes In Use

GC Heap|Bytes In Use reports the amount of memory being currently used by objects.

GC Heap|Bytes Total

GC Heap|Bytes Total reports the total amount of memory allocated by the JVM.

As explained at greater length in the *APM Sizing and Performance Guide,* allocating either too small or too large an amount of memory to the JVM can lead to performance problems. In brief, you can use these guidelines:

- Allocating too small an amount of memory means the GC process will run more frequently, leading to short but frequent performance degradation problems.
- Allocating too large an amount of memory means that when the GC process runs, it will take a relatively long time and cause a performance degradation for that time.

Therefore, the application administrator can use these metrics to help determine the correct size of the memory heap.

Once the correct size has been determined, you can watch these metrics over time to understand what nominal performance looks like. The Bytes In Use metric should show periodic increases and decreases which, over time, form a repeated pattern and show no evidence of a memory leak.

Monitor Performance with the GC Monitor Metrics

GC Monitor provides a new set of metrics to give you a view to the internals of the JVM, including memory allocation and heap growth rate. It allows you to allocate the heap inside the JVM by verifying that all garbage collectors and their memory pools are allocated properly. In this way, you can detect GC issues that are adversely affecting performance.

Supported JVMs

GC Monitor supports only the following JVMs:

- Sun JVM, version 1.5.0 and higher, both 32-and 64-bit
- IBM, version 1.5.0 and higher, both 32-and 64-bit

To use GC Monitor metrics to tune memory allocation:

1. In the Browse tree, browse to the agent node on the host whose GC activity you want to monitor, and expand the GC Monitor node.
2. Monitor the metrics and memory pool use of each garbage collector. See specific definitions of each GC Monitor metric (see page 359) in the metrics appendix.
3. Based on the metrics, reallocate the size of memory pools to increase GC efficiency.

If you need guidelines to help you reallocate the size of memory pools, see documentation appropriate for your JVM.

CA Technologies Internal Use Only Authorized Use Only Pre-Release Documentation

Monitoring Status with the Application Triage Map

You can use the application triage map to monitor the status of your application.

Make use of both:

- the By Business Service map, which allows you to monitor the status of Business Services and Transactions.
- the By Frontend map, which allows you to monitor application status.

By Business Service Map

To use the By Business Service map to monitor the status of Business Services and Transactions:

1. Open the Workstation Investigator and click the Triage Map tab.
2. Expand the By Business Service tree and select one of the business services.

   If you are running TIM and have configured business transactions and services for monitoring and reporting through the Enterprise Manager, the By Business Service map displays a map of the business transactions that make up your business service, along with the frontends and called backends where the transactions happen.

   For information about configuring business services and transactions for monitoring, see the *APM Transaction Definition Guide*.

   For information about understanding the By Business Service map, see By Business Service Application Triage Map.
3. Set alerts on one ore more of the following. See Create and Edit Application Triage Map Alerts (see page 104).
   - Business Transactions.
   - Customer Experience.
4. When problems arise, they become visible through alert indicators. See Responding to Alerts in the By Business Service Map (see page 166).

By Frontend Map

To use the By Frontend map to monitor application status:

1. Open the Workstation Investigator and click the Triage Map tab.
2. Expand the By Frontend tree and select one of the frontends.

   The By Frontend map automatically displays the frontends and called backends of applications you have set up agents to monitor. See By Frontend View of the Application Triage Map (see page 90).

CA Technologies pre-release documentation. Authorized use only.

3. Set alerts on:
   - one or more of the frontends and called backends. See Create and Edit Application Triage Map Alerts (see page 104).
   - the resources where the frontends are deployed. See Create and Edit Resource Metrics and Alerts (see page 109).

   When you set alerts on the **frontends** themselves, you set alerts on one or more of the health metrics for that frontend; when you set alerts on the **resources** element of a frontend, you set alerts on health metrics for the application server(s) on which that frontend is deployed.
4. When problems arise, they become visible through alert indicators on the By Frontend map, and through any actions you have configured for the alerts. See Responding to Alerts in the By Frontend Map (see page 167).

Responding to Alerts in the By Business Service Map

When you see a Caution or Danger alert on the Customer Experience icon:

1. Hover your cursor over the Customer Experience (CE) icon and look at the CE health metrics.
2. Right-click the CE icon and choose “Find Incidents…”

   A browser window opens and displays the Customer Experience Incidents page, filtered on the current Business Transaction and sorted by the Last Defect Reported timestamp.
3. Return to the Workstation's By Business Service Triage Map display.
4. Right-click the alerting Customer Experience (CE) icon and select “Browse All Customer Experience Metrics.”

   The Investigator display changes to the Metric Browser tab, expanded to the Business Transaction node under the TESS Agent. From here you can explore the full range of available customer experience metrics, including the per-TIM breakdowns. See Diagnose the Problem with the Metric Browser Tab (see page 179).

When you see a Caution or Danger alert on the Business Transaction (BT) oval:

1. Right-click the BT oval and select “Show Alert Details.”

   The Alert Details pane opens to the right of the map and shows all the alerts that have been defined for this BT, along with their current state. Note the abnormal alert.
2. Double-click the BT to open the Locations table in the bottom pane, where you can see which Location or Locations are over threshold.

   Tip: If only one of the locations is over threshold, this could point to a denial of service attack, or perhaps resource problems on that particular machine.

CA Technologies Pre-Release Documentation
Authorized Use Only

3. Select the alert in the Alert Details pane to see how it was defined.

   Tip: You may want to adjust the alert's sensitivity.

4. Double-click a row in the Locations table to switch the Investigator display to the node corresponding to that BT under that location.

   For example, if you are examining the Login BT and double-click the row in the Locations table for MyServer22r, the Metric Browser tab will open with the tree expanded under the MyServer22r agent down to the level of Business Segment | <Business_Service_Name> | <BT_Name> | <BTC_Name>.

   From here you can diagnose the problem at that Location. See Diagnose the Problem with the Metric Browser Tab (see page 179).

Responding to Alerts in the By Frontend Map

When you see a Caution or Danger alert on an element in the By Frontend map:

1. Hover your cursor over the element and view health metrics in a tooltip.
2. Right-click the element and select select “Show Alert Details.”

   The Alert Details pane opens to the right of the map and shows all the alerts that have been defined for this element, along with their current state. Note the abnormal alert.

3. Double-click the element to open the Locations table in the bottom pane, where you can see which Location or Locations are over threshold.

   Tip: If only one of the locations is over threshold, this could point to a denial of service attack, or perhaps resource problems on that particular machine.

4. Select the alert in the Alert Details pane to see how it was defined.

   Tip: You may want to adjust the alert's sensitivity.

5. Examine any possible problems with the Location by doing one or both of these procedures:

   To look into problems with application performance at a single location:

   - Double-click a row in the Locations table to switch the Investigator display to the node corresponding to that frontend or backend under that location. From here you can diagnose the problem at that Location. See Diagnose the Problem with the Metric Browser Tab (see page 179).

   To look into problems with infrastructure performance at a single location:

   - Right-click a row in the Locations table and select "View this location..." to switch the Investigator display to the Location Map. See Monitor Performance with the Location Map (see page 168).

CA Technologies Product Documentation
Authorized use only.

Monitor Performance with the Location Map

You can use the Location Map to monitor performance of your infrastructure, in implementations where you have only CA APM monitoring as well as implementations where other CA Technologies applications are part of your system.

On the Location Map, "locations" correspond to infrastructure layers in two- and three-tier arrangements. These layers correspond to agents, physical hosts, virtual hosts, or external systems whose data is imported by [set value for the rose variable at the book level]. The layers are represented in the map by "container" rectangles.

For more information on understanding the Location Map, see View Data from External Sources Using the Location Map (see page 119).

To open the Location Map:

From the Triage Map display:

1. Right-click a map element and select "View Locations for <Element_Name>".

   The Locations table opens in the bottom pane of the triage map.

2. Right-click a row in the Locations table and select View This Location.

   The display changes to the Metric Browser tree.

3. If necessary, click the Location Map tab.

From the Metric Browser tab:

1. Expand the tree to the Frontends | Apps node.
2. Select one of the Frontends.
3. Click the Location Map tab.

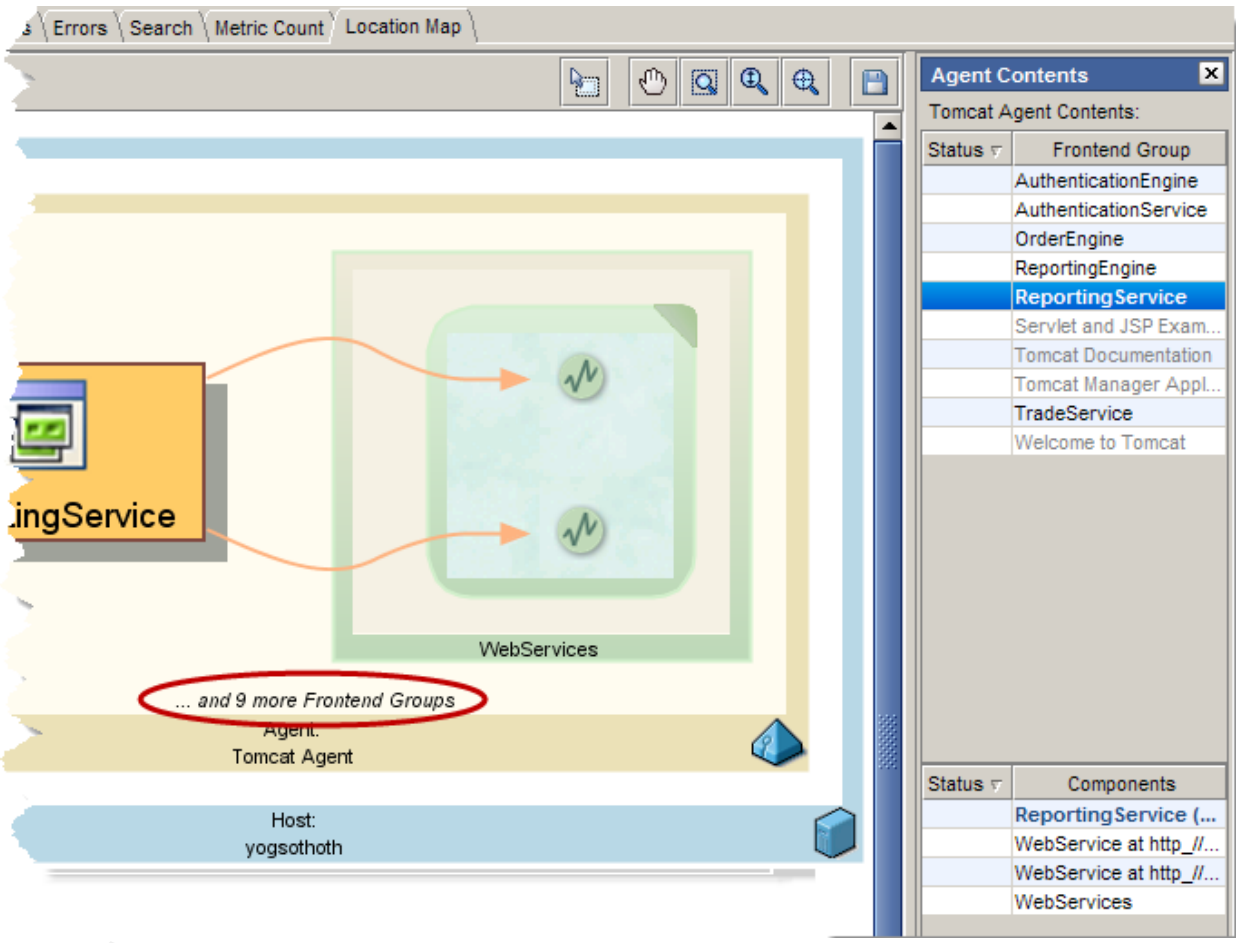
CA Technologies Pre-release Documentation
Authorized Use Only

On an APM-only system

To use the Location Map to monitor performance:

1. Notice the message about additional frontend groups.

   The Location Map automatically limits its display to one or two frontend groups per container. When there are additional frontend groups, the legend "... and *n* more frontend groups" will be displayed, as shown in this illustration:



2. Right-click the frontend element and select "Show Contents."

   The Agent Contents pane opens on the right of the map. Notice that the number of "more Frontend Groups" plus the selected frontend equal the number of Frontend Groups listed in the Agent Contents pane.

3. Double-click the alerting host container to get more information.

4. Right-click a displayed frontend or called backend and select View Health Metrics.

   Health metrics open in a pane at the bottom of the map. This pane allows you to select an available metric and view it as a graph.

5. Double-click the agent container and repeat the steps above.
6. Right-click the agent container and select "Go to Location Map for <Agent_Name>" to concentrate on a particular agent.

On a system with information imported from [set value for the rose variable at the book level]

Integration with [set value for the rose variable at the book level] allows CA APM to display information from other CA Technologies applications.

Use the same widgets and commands described above to monitor performance on a Catalyst-enabled infrastructure. Some containers will be different, referring to virtual hosts or applications outside the APM implementation.

Monitor Performance with Frontends Metrics

You can use the Frontends node in two ways to monitor general application performance: by monitoring the standard metrics, and by noting the top URLs.

Monitoring standard frontends metrics

Introscope displays the five basic Introscope metrics (see page 349) for each frontend in the Metric Browser tree, under Frontends|Apps|<*Frontend_Name*>.The same metrics appear in the Triage Map tree under By Frontend|<*Frontend_Name*>|Health.

Good performance

An application is performing well which shows a high volume of requests ("Responses Per Interval") being handled, with a corresponding low latency (i.e., a low Average Response Time). A good rule of thumb is approximately one second per transaction.

Problematic performance

Concurrent methods are methods that started during an interval without finishing during the same interval. Because you want methods to complete quickly, an unusually high number of concurrent invocations is undesirable. You may see temporary spikes in concurrent invocations, but the metric should return to zero each time. If it does not, this may indicate a "bottleneck" of threads, number of database connections, or some other shared resource.

For more information about using Frontends metrics to monitor performance, see the *CA APM Sizing and Performance Guide.*

CA Technologies Product Documentation: Authorized Use Only

Recognizing worst-performing transactions

Another good way to monitor performance is to be aware of which transactions are consistently slow. You can configure a data viewer to display the slowest transactions as a bar chart.

One of the best ways to do this is to configure URL groups as the base for your metrics grouping. See Defining type and number of metrics shown in Filtered View (see page 291).

Monitoring performance with backends metrics

The Backends node of the Investigator tree shows the five standard metrics for each connected backend system.

Two different metrics under the Backends node help you recognize nominal performance.

SQL statement frequency

Staying aware of which SQL statements are processed most is a good way to familiarize yourself with application performance.

To look at SQL statement frequency as a measure of performance:

1. Under the Backends node, open the node for the application you want to monitor.
2. Under the application, open the SQL node.

   The Overview tab displays a list of queries and other SQL statements running against the database resource.
3. In the Queries section of the right pane, click the Responses column heading to sort the table by number of responses.
4. Note the SQL queries that are sent most.

Database connection patterns

You should also stay aware of your application's database connection patterns, and be aware when the pattern is broken.

The way your application establishes and maintains database connections depends on the platform.

CA Technologies Prerelease Documentation. Unauthorized Use Only.

Monitor performance with the APM Status Console

You can use the APM Status Console to monitor the runtime characteristics of your Enterprise Manager:

- Activated clamps
- Important Events
- Denied agent connections

The APM Status icon is displayed in the Workstation and changes in appearance when one of these events happens. This icon has two states, as shown in the illustration below.



On the left is the icon in normal state. On the right is the icon in alerted state.

When the APM Status icon is red

When the icon appears in its alerted state, click it.

The APM Status Console opens.

Respond to APM Status Console Events

Follow these steps:

1. In the Enterprise Manager Map, select the nodes with problems, as identified by a red color on the map elements representing Enterprise Mangers or Cross-Cluster Viewers.

   The lists of clamps, important events and denied agents display the problems with the selected Enterprise Manager.

   Note: If you select the MOM Enterprise Manager, the lists in other panes display *all* problems on the cluster.

2. Read the lists of clamps, important events and denied agents to see problems.
3. Double-click an item in the list of clamps to browse to the Workstation's Metric Browser display of the metric associated with that clamp.

   For clamps with no associated metric, a message will appear: "Unable to browse to the selected metric."

4. Clamps, important events, and denied agents also appear in the log file. Consult IntroscopeEnterpriseManage.log to see more details.

CA Technologies Prerelease Documentation. Authorized Use Only.

Note: The APM Status Console does not provide functionality for directly resolving the problems it displays. As an applicaiton administrator, you take action based ont he information it provides.

Reading and understanding notifications

Notifications of application performance problems and breakdowns can take several forms.

Alert notifications in dashboards

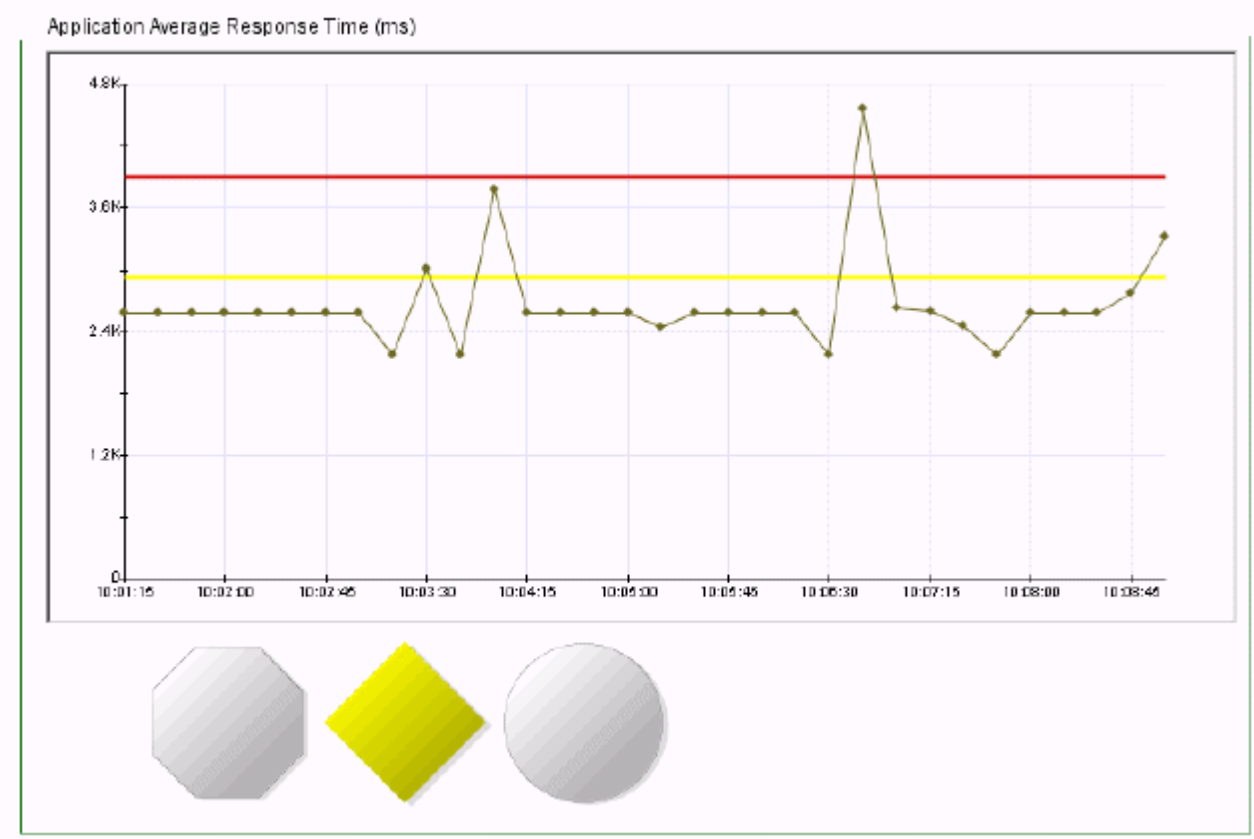
The most obvious form is a visual notification on a Console dashboard. The illustration below shows a dashboard with a single graph which has been configured with a yellow line showing the Caution threshold at 3000 ms and the Danger threshold at 4000 ms.

The graph shows:

- The Average Response Time crossed the Caution threshold several times in the last several minutes.
- The Average Response Time crossed the Danger threshold once, about two minutes ago.

CA Technologies Prerelease Documentation. Authorized Use Only.

- According to the most recent measurement, application performance is currently in a Caution state.



The indicators under the graph show another way of indicating alert status. The way your dashboards appear depends on how you, or your administrator, have configured them.

When you see a dashboard showing a Caution or Danger condition, the alert indicator is usually rolling up metrics from several sources. Your task should be to find out what underlying metrics are causing the condition.

To change an alert view:

- Display an alert in the Investigator Preview pane and select Properties > Alert View.

Alert messages

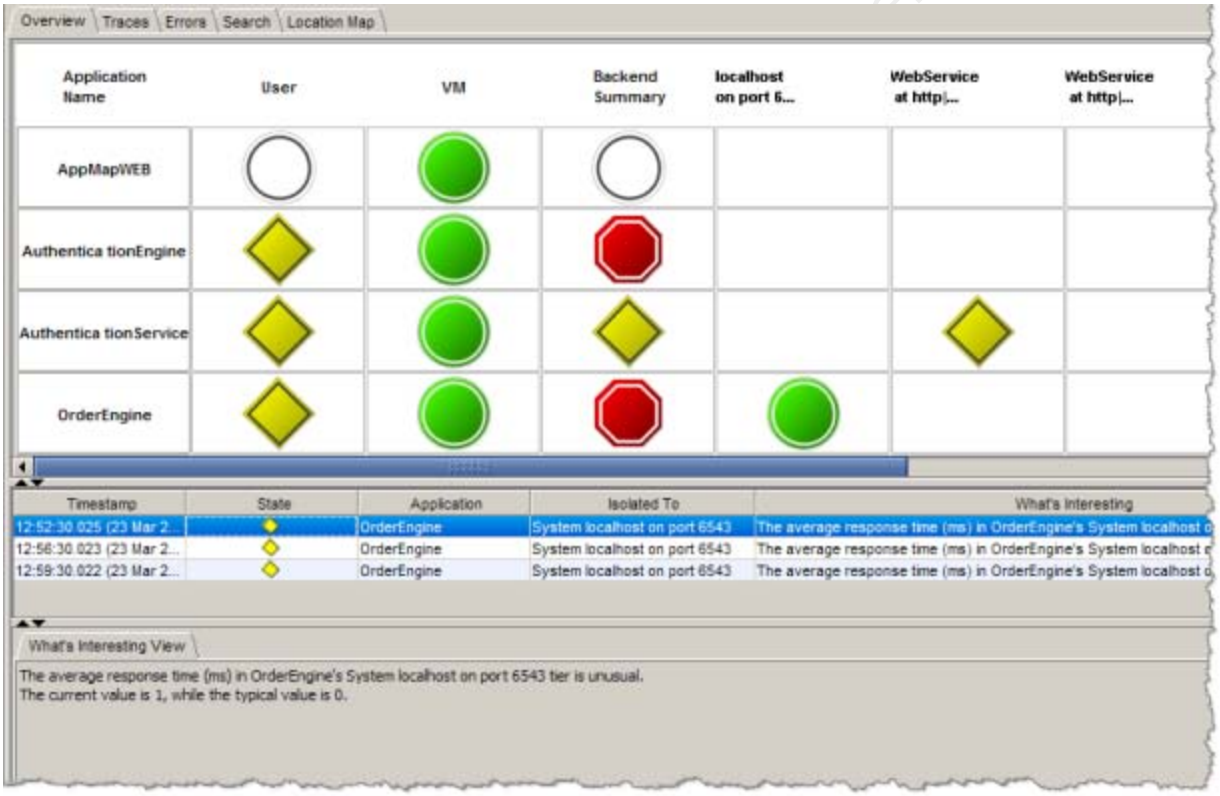
Alert messages are triggered by an action associated with an alert status. These alerts appear automatically. You can also view alert messages by selecting Workstation > Show Alert Messages.

Alert notifications in What's Interesting events

The Investigator displays a list of What's Interesting events in the Overview tab view when the agent is selected.

In the illustration below, alert indicators in the application grid show anomalous conditions for three out of four frontends.

Whenever alert indicators in this grid are in a caution or danger state, more information appears below in the What's Interesting table. You can use the information in the tabular pane to see the application, the tier that the problem affects, and a description of the anomaly. Here the user has selected the first row of the table, and more information is displayed in the bottom pane.



Other Kinds of Notifications

Introscope alerts may be configured to take various actions when metrics cross a caution or danger threshold, including:

- sending a call to a pager
- sending email

For information about configuring actions like these, see Monitoring performance with alerts (see page 306).

When you receive one of these notifications, you first determine the source of the alert. Doing so is easiest when the alert has been configured with a meaningful name that provides information about the source. If you recognize the alert by its name, based on your familiarity with your applications or with the alerts themselves:

- Look at the Console dashboard you are responsible for monitoring.
- Look at the application triage map for the application or business process you are responsible for monitoring.

Respond to a Notification

When you receive a notification of a problem, the general procedure to follow is as follows. (This assumes you have access to both of the main CA APM user interface utilities, Console and Investigator.)

1. Confirm the problem (see page 176).
2. Use links (see page 178) from the application triage map or the Console to jump to the metric browser tab.
3. Use Investigator tools to uncover and isolate the problem -- see:
   - Using the Metric Browser Tab to Diagnose the Problem (see page 179)
   - Using Search (see page 181)
   - Using Transaction Trace (see page 182)

Confirm the problem

If you received the notification through a pager, email, or other way, you should find the Workstation display that gives you more information about the problem.

Use the Application Triage Map to Determine the Source of an Alert

When you receive an alert notification, usually by email, you first determine its source.

To determine the source of an alert:

1. Identify the alert source by the name in the notification.

   Alert notifications contain only the alert name and the fact that a threshold has been exceeded. Ideally, that alert name will give you clues to its source. If you cannot tell the source of the alert from the name, see Tracing unfamiliar alerts (see page 177).

CA Technologies Prerelease Documentation. Authorized Use Only.

2. In the Workstation Map tab, do one of these:
   - Expand the By Frontends tree to find the Frontend that is the source of the alert.
   - Expand the By Business Services tree to find the Business Transaction that is the source of the alert.
3. In the graphical map display, double-click the corresponding map node to display the Locations table.

   Because alerts on elements in the application triage map are always aggregated across all agents reporting metrics for a given frontend, the only meaningful investigation path is to determine which agent is reporting the excessive metrics. To do this, you look in the Locations table, which shows all the physical locations hosting agents that supply metrics for the frontend or business transaction. For more information about the Locations table, see List of Physical Locations (see page 86).
4. In the Locations table, click column headings representing metric names to sort table data by that metric.

   As you sort the table by each of the metrics being reported by each agent, you should be able to recognize anomalous readings.
5. When you have tentatively identified the metrics that triggered the alert, double-click the table row containing the anomalous readings. The display will jump to the Metric Browser tree.

From here, go to Using the Metric Browser Tab to Diagnose the Problem (see page 179).

Unfamiliar Alerts

Ideally, you will be able to recognize the source of an alert notification because the alert has been configured with a meaningful name. In some cases, however, you will have to search manually through the applications you're responsible for to find the source of an alert notification.

To find the source of an alert you don't recognize:

1. In the triage map tab, do one or more of these:
   - Expand the By Frontends tree and look for a node in Caution or Danger state, as signified by a yellow diamond or red octagon respectively.
   - Expand the By Frontends tree and locate Frontends for applications you are responsible for.
     - Expand the Health node for each frontend, and browse through health metrics looking for instances of metrics exceeding thresholds.

CA Technologies Proprietary Documentation. Authorized Use Only.

- Select individual frontends and examine the map displays. Unroll successive called frontends using the "dog-ear" control (see page 83) and examine the alert indicators for each.
- Examine the alert indicators on connections to other called backends, as shown in Backend connection metrics icon (see page 81).

2. In the metric browser tab:
    1. Expand the metric browser tree to the Frontends > Apps node and examine the metrics for each Frontend to find anomalous readings.
    2. Use these clues to search further in the metrics contained in the metric browser tab.

From here, go to Using the Metric Browser Tab to Diagnose the Problem (see page 179).

Problems Displayed in Console Dashboards

The dashboard often contains links to specific metrics in the Investigator.

If the problem is not currently happening, use Investigator tools to look at metric performance over a historical time range. See Viewing Historical Data in the Metric Browser tab (see page 180).

Using Hyperlinks to Find More Information

You can use various kinds of hyperlinks to find more information about alerts and anomalies.

Hyperlinked dashboard elements

Two kinds of dashboard hyperlinks allow you to navigate between Introscope dashboards and the Investigator:

- **Automatic hyperlinks**—Introscope automatically links a data viewer to the metric grouping it is based upon—the Links menu for the viewer contains a link to the underlying metric grouping definition in the Management Module Editor. Similarly, dashboards that contain Data Viewers based on the same metric grouping are automatically linked, and you can navigate between them using the Links menu.
- **Custom hyperlinks**—You can define custom links for dashboard items, to link to other dashboards or to web pages. You can define custom links if you have dashboard editing permission.

Note: Some out-of-the-box Console dashboards—for example, EM Capacity—do not automatically contain links to underlying data. Edit these default dashboards or create new dashboards with links. For information about creating and editing custom links, see Creating and Managing Custom Hyperlinks (see page 302).

CA Technologies Prerelease Documentation. Authorized Use Only.

To follow dashboard links:

1. For linked objects, the cursor changes to a hand when you hover your cursor over the object.
2. Double-click the object to follow the link to its default target.

To see a list of available links:

1. Select a dashboard object and select Properties > Links.
2. Right-click the dashboard object and select Links from the context menu.

To see the target of a hyperlink in a new window:

- Press Shift and click the object.

If no links are available for an object, the Links menu is disabled.

Hyperlinks in tooltips

In a Data Viewer, you can use the built-in links in tooltips to go to the underlying metric for the data being displayed.

To use a tooltip hyperlink:

1. Hover your cursor over any element in the Workstation metrics tree or in a Data Viewer, such as a point on a graph.
2. While the tooltip is active, press F2.

   This "focuses" the tooltip, which otherwise would disappear when you moused off the data point.
3. Click on the hyperlink.

Diagnose the Problem with the Metric Browser Tab

The following tools in the metric browser tab help you find more information about an issue:

- Historical metrics (see page 180)
- Search (see page 181)
- Transaction Tracer (see page 182)
- Thread dumps (see page 183)

CA Technologies Prerelease Documentation. Authorized Use Only.

Using Live and Historical Metrics

By default, the views in Workstation are of live metrics, with the data refreshing every 15 seconds. Data which is not displayed in a live chart is saved by Enterprise Manager as historical data. To diagnose a problem which may have begun some minutes or hours ago, you view historical data.

Viewing Historical Data in the Metric Browser Tab

To view historical data, you select a time range—using a time range can help you quickly identify the time a problem occurred. For example, if you think the problem occurred within the last hour, you could set the time range to an hour and look at the data from the current time backward. If you don't see the problem within that hour range, you can use the controls to move backward or forward to locate the time the problem occurred.

To view historical data:

1. Select the metric or dashboard for which you want to see historical data.
2. Select a time range for the historical view from the Time Range drop-down menu.

   Introscope shows the data for that range, using the duration that you selected from the Time Range drop-down menu and setting the end time to the current time.

   Note: When you use the time-range control to view historical data, the range you select is applied to other metrics or dashboards in the same window, and to any new windows that you open.
3. To select a resolution to adjust the granularity of the view, increase or decrease the number of data points that appear.

   Each pre-defined time range is associated with a default resolution. Typically, you will not need to change this. Changing the resolution is generally useful when you need to see a greater level of detail or granularity in the data than is displayed by default.
4. After selecting a time range you can adjust it, using the controls to scroll in increments based on the time range you selected:
   - Drag the slider on the time bar to change the time range.
   - Click the arrows to move backward and forward in time.

     The single arrows move backward or forward in small increments; the double arrows move backward or forward in time increments that are about the time of the selected time range.
   - Click the Reset icon to reset the end time of the range to the current time.

For more information about these tools, see Historical application triage map views (see page 111).

CA Technologies Product Documentation. Authorized Use Only.

Using Zoom on Historical Data in Graphs

When you are viewing historical data in a graph, you can zoom in on data.

To zoom in on data in a chart:

- Do one of these:
  - Click the mouse pointer on a graph position and drag to specify the time range.
  - Right-click on the graph and click Zoom to fit data.

Introscope refreshes the data in the viewer based on the new query, and the time range in the viewer shows the new range.

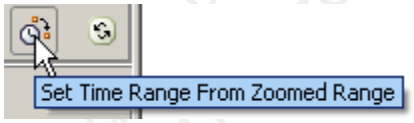
To zoom back out:

1. Right-click the zoomed-in chart.
2. Click Zoom Out or Zoom All the Way Out.

The global time range in the window and the Time Range control do not change automatically when you zoom in on data. For example, if you zoom in on a ten-minute period on a graph with the Time Range set to 1 hour, the graph shows the ten-minute period but the control remains at 1 hour, and the time bar still shows the hour range.

To set the global time range and the Time Range control to match the zoomed view:

- Click the Set Time Range From Zoomed Range button:



Using Search

The Search tab (see Search tab (see page 140)) is active for every node of the Investigator tree. Using this tab, you search for any of the metrics under a particular node.

To find the Search tab:

1. Select a node in the Introscope tree.
2. Select the Search tab.

CA Technologies Product Documentation. Preliminary. Authorized Use Only.

To use plain text search:

1. Enter a string in the search pane.
2. Press *Go* or *Enter.*

   The search results are displayed in table format. The results show all resources whose name includes the search string.

Tip: Selecting any of the metrics listed in the table displays a chart showing a live view of the metric.

To display results including Min, Max and Count for each result:

- Select Show Min, Max and Count.

Tip: You can select Show Min, Max and Count after searching and the results will refresh with the new columns.

Using regular expressions

The Search pane accepts any regular expression using Perl 5 regular expressions language.

Tip: Perl 5 regular expression language is also used to define metric groupings. For more on how metric groupings are defined, see Creating a new metric grouping (see page 274).

To use regular expressions in search:

1. Select *Use Regular Expressions.*
2. Enter a regular expression in the Search pane.
3. Press *Go* or *Enter.*

   The Search tab displays the results.

Using Transaction Trace

The Transaction Tracer is a powerful tool enabling you to trace the activity of transactions as they flow through a Java Virtual Machine (or a Common Language Runtime [CLR] in a .NET environment) inside a production application.

See Using the Introscope Transaction Tracer (see page 203).

CA Technologies Prerelease Documentation. Authorized Use Only.

Using thread dumps

Viewing a thread dump can help you identify the source of JVM performance problems.

In these example situations, you can collect and analyze a thread dump. This agent JVM snapshot can help you understand the source of a slow down, hanging server, or unusually high CPU usage.

- A Workstation displays stall metrics but no transactions show up when running a Transaction Trace. This situation can happen because transactions are not getting completed, and the Enterprise Manager is getting incomplete information about the agent server hang.
- CPU usage for an Enterprise Manager is low, yet there are long response times. This situation can indicate that all of the threads in an operation are deadlocked, blocked, or waiting.
  - If a method is taking a long time to load, one thread can be using a large amount of CPU. Meanwhile, all the other threads wait for the single thread to complete its task before starting their next tasks.

Note: You must have thread_dump permission to collect a thread dump or load a previous thread dump. For more information, see the *CA APM Security Guide*.

Follow these steps

1. Select an agent node in the metric browser tree, then click the Thread Dumps tab.
2. Click the Collect New button.

   The header displays the thread dump time. The thread dump summary bar displays the total number of number threads and number of threads that are waiting, blocked, or running.

   Note: Successive continuous thread dumps can degrade CA APM performance. For best performance, pause between thread dumps.
3. To examine information about one thread, select it in the thread information table.

   Each thread is associated with a stack trace, which lists all methods in the order called and displays in the thread stack trace table.

   - Click the Hide redundancies check box in the thread stack trace table to hide the redundant methods that are displayed. If more than one method of the same class is called consecutively in the stack trace, then only the first method is displayed. If the check box is not selected, then all the methods in the stack trace are listed in the thread stack trace table

CA Technologies Prerelease Documentation
Authorized Use Only

When the redundant calls are hidden, the number of additional hidden calls displays in angle brackets to the right of the method name.

For example, if you have selected a thread and these methods display in the stack trace table:

java.net.PlainSockettlmpl.SocketAccept(Native Method)

java.net.PlainSockettlmpl.accept(PlainSockettlmpl.java:457)

java.net.ServerSockettimpl.Accept(ServerSockettimpl.java:473)

java.net.ServerSockettimpl.accept(ServerSocket.java:444)

com.ibm.rmi.transport.ListenerThread.run(ListenerThread.java:166)

When you select the Hide redundancies checkbox, these methods display:

java.net.PlainSockettlmpl.SocketAccept(Native Method) <1>

java.net.ServerSockettimpl.Accept(ServerSockettimpl.java:473) <1>

com.ibm.rmi.transport.ListenerThread.run(ListenerThread.java:166)

The later calls for the java.net.PlainSockettlmpl.SocketAccept and java.net.ServerSockettimpl.Accept methods are now hidden in the trace stack. These calls are included in hidden method count of <1>.

4. Use the search pane to search for a specific string within all the thread dump information.
5. List all threads, or threads in a deadlocked, blocked, running, or waiting state.
   - Click the Threads State drop-down list (to the right of the search pane) then select the state.
   - Use the Threads State drop-down to further filter a search.
6. To save current thread dump to a text file, click the Save as Text button. Save the thread dump text file to the location you select.

   Introscope saves all the thread dump details in the thread information table to a text file. You can send the file to another person or look at it using a text editor.

   Note: You cannot import the thread dump text file data into the Thread Dumps tab for viewing.
7. View thread dump details for a previous thread dump.
   a. Click the Load Previous button.
   b. Select one row in the Load a Previous Thread Dump dialog, then click OK.

   The data for the selected thread dump displays in the Thread Dumps tab.

   Note: You can select from all previous thread dumps, which Introscope saves to a different location than the saved thread dump text files. You can view one previous thread dump at a time.

CA Technologies Proprietary - Documentation. Authorized Use Only.

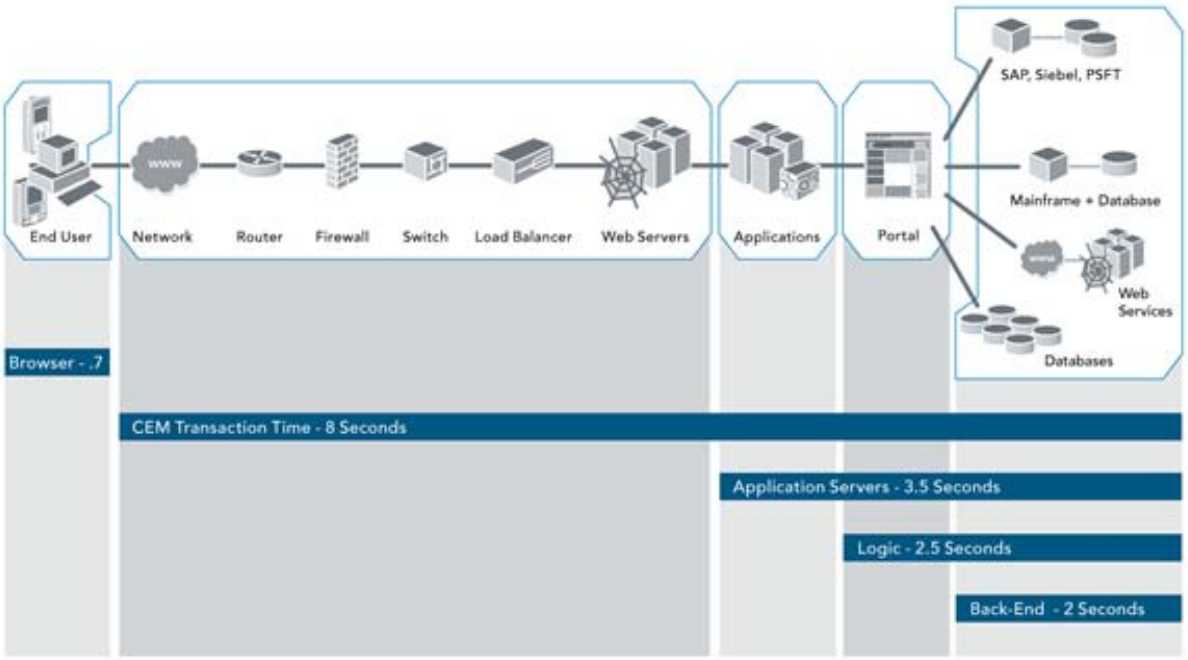
8. Compare threads from a healthy agent with threads from an agent experiencing problems.

- Open an Investigator for each agent, then collect a thread dump for each agent.

Diagnosing Problems with CA APM Customer Experience Transactions

You can find the root cause of slow time CA APM Customer Experience transaction defects by analyzing transaction times from various perspectives.

This diagram shows an example of how CA APM Customer Experience transaction time is measured.



Note: Times depicted in this graphic are for illustration purposes only.

Transaction-related times are reported both by TIMs and by Introscope agents. As such, it is possible for time comparisons to be confusing. For example, it is possible for the CA APM Customer Experience transaction time to be less than the application server time.

See Incident-related terminology for detailed descriptions of the time-related measurements.

Incident-related terminology

As you begin to investigate CA APM Customer Experience defects and incidents and the associated Introscope transaction traces, you should be familiar with the related terminology.

application server time

The measure of time that it takes the application server to process the transaction, as reported by Introscope. Introscope reports the response time from the first Blamed component that was involved in serving the response.

backend time

The measure of time that the backend system's suspected Blame component (for example a database) takes to complete, as reported by Introscope.

CA APM Customer Experience transaction time

The total elapsed time of a transaction, from the first request packet to the last response packet, as monitored by the TIM.

logic time

The measure of time that the suspected Blame component's program code takes to process the transaction, as reported by Introscope.

suspected Blame backend component

The most specific portion of the backend time that is identified as being the suspected cause of delay in a slow transaction. In Introscope, the suspected Blame backend component appears as the widest, but not necessarily the lowest, backend component in the graph.

Note: It is the *slowest lowest* backend component, not the *lowest slowest* backend component in the Introscope graphs.

The suspected Blame backend component is identified by taking the lowest backend component that takes longer than 25% of the overall backend time to complete.

suspected Blame component

The most specific portion of logic (or program code) that is identified as being the suspected cause of delay in a slow transaction. In Introscope, the suspected Blame component appears as the widest, but not necessarily the lowest, component in the graph.

Note: It is the *slowest lowest* component, not the *lowest slowest* component in the Introscope graphs.

The suspected Blame component is identified by taking the lowest (non-backend) component that takes longer than 25% of the overall transaction time to complete.

CA Technologies Confidential Document
Authorized Customer Use Only.

time to first response

The elapsed time from the *last* (see Note) packet of the request to the first packet of the response for the component.

The time to first response varies, based on the type of defect being tracked:

- component defect—the time to first response for that component.
- transaction defect—the time to first response for that transaction's identifying component.
- business transaction defect—the time to first response for the identifying transaction's identifying component.

Note: This setting can be changed to the *first* packet of the request to the first packet of the response. If you need this setting to determine inbound network latency, contact CA Support.
Changing this setting affects only *new* data (that is, existing data values are based on the time to first response setting that was in place at the time the data was collected).

See the *CA APM Glossary* for all CA APM terms and definitions.

Wily CEM problem resolution triage metrics on Introscope

Introscope agents produce metrics in terms of the CA APM Customer Experience business service and business transaction hierarchy, as defined in the CA APM Customer Experience domain configuration file. Metrics are visible in the Metric Browser tree.

For each business service, a Stall Count metric appears under the Business Service node for business transactions defined with post parameters.

Note: Stall count metrics do not correlate to any business transaction definition as a limitation of Introscope.

For each business transaction, the following metrics appear:

- Average Response Time (ms)
- Errors Per Interval
- Responses per Interval
- Concurrent Invocations

Note: For business transactions defined with post parameters, Stall Count and Concurrent Invocations metrics are not used as a limitation of Introscope.

CA Technologies Pre-release Documentation. Authorized Use Only.

Incident overview

The incident overview page gives you basic incident data and also provides links to critical defect information.

To view incident overview information:

1. Open the CA APM Customer Experience Console.
2. Select CEM > Incident Management.

   The Incidents page appears.
3. Click on the link for a specific slow time incident ID number.

   The Incident Overview page appears for the incident you selected. The page includes:

- Overview — of defect, numbers of defects and users, status data.

  You can click on the links in this section to learn more, for example:

  - Analysis graph of the defect condition
  - List of specific defects
  - Specific affected users
- Evidence collection — types collected (or not) for this incident.

  You can see if evidence collection has started (or perhaps if it should be collected in the future).
- Incident transaction trace — data and links for Introscope transaction trace; appears only when Introscope integration has been enabled.
- Problem resolution cycle — times related to this incident.

  You can see how long the defect condition has been a problem.
- Defect time distribution graph — shows where transaction time is being spent.
- Close incident — date and closed by; cause and resolution.

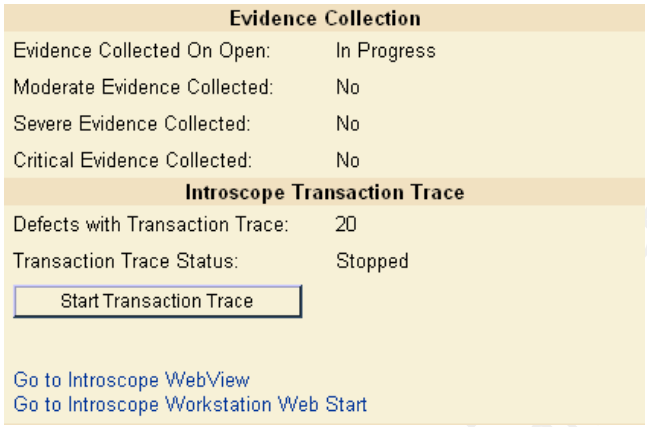
For detailed information about troubleshooting the incident, see Incident troubleshooting to find root cause (see page 196).

To see more information about affected users and groups, click the Affected User Groups or Affected Users links.

CA Technologies Prerelease Documentation. Authorized Use Only.

You can collect transaction-related incident information using Transaction Tracer. Manually start or stop the transaction trace session by using the Transaction Trace button.

You can also click the Go to Introscope WebView or Go to Introscope Workstation Web Start links to investigate the business service behind the transaction.

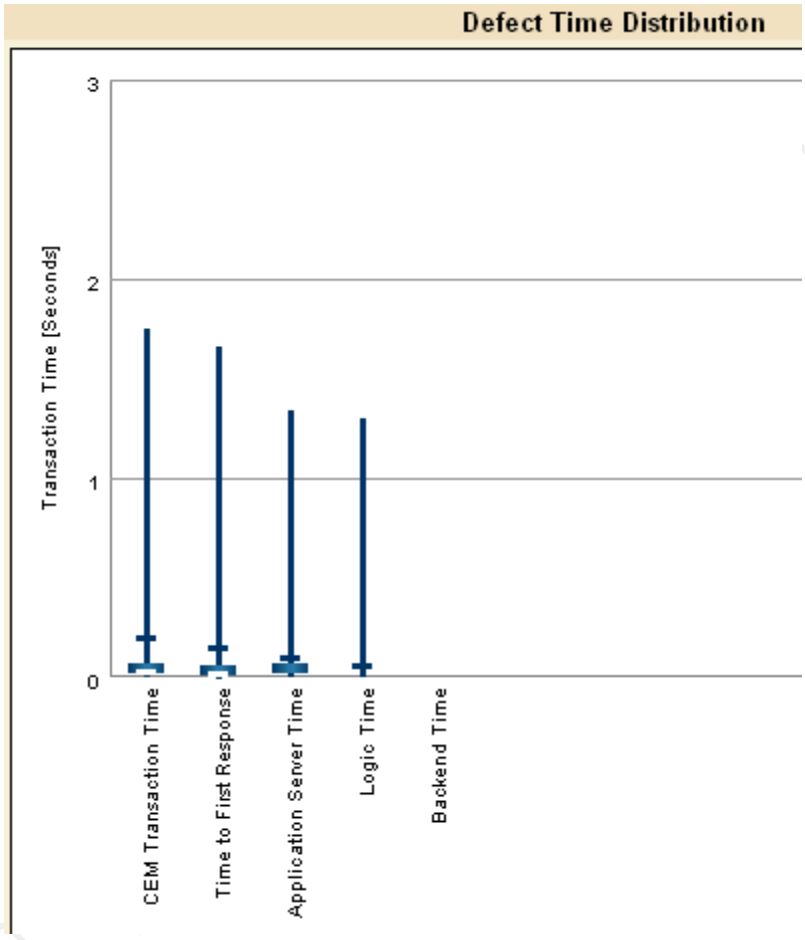


Important! Starting a transaction trace can cause performance degradation on the instrumented application. (If the related slow time defect specification is too low, it is possible to bring down the instrumented application.) See Configuring transaction traces for more information.

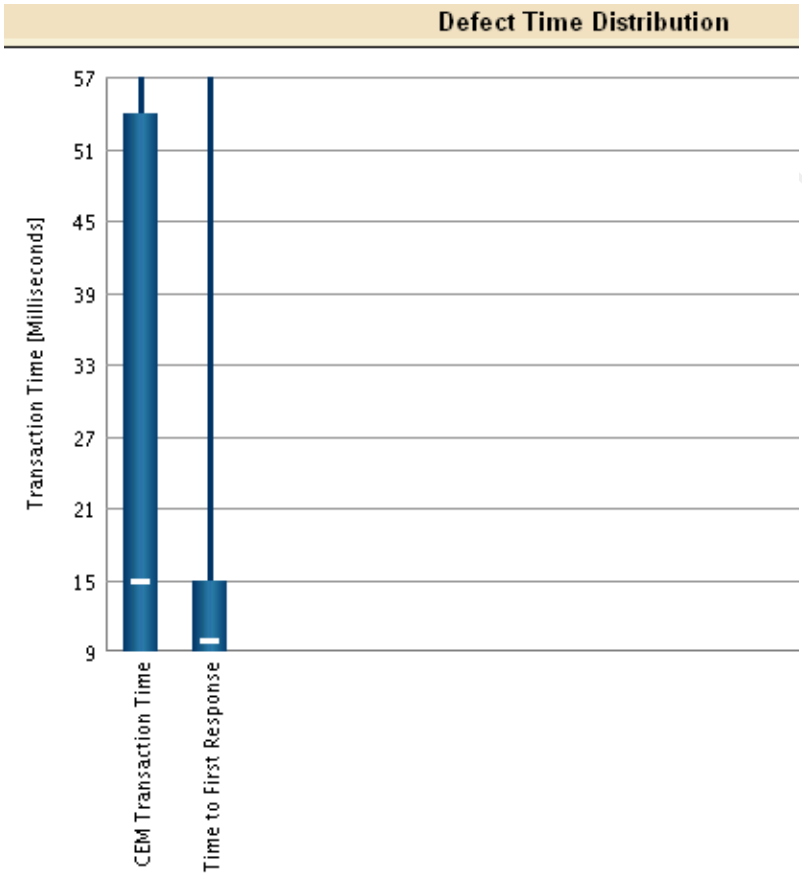
CA Technologies Preferred Authorized Use Only Documentation.

The defect time distribution graph shows you, at a glance, where transaction time is being spent in each tier.

Note: A tier will not display a bar in the graph if it has a zero value.



Note: If Introscope is enabled and the incident is a slow time incident, then the graph displays all tiers with data. Otherwise, only CEM transaction time and time to first response appear.



Using Web Start to Drill Down from CA APM Customer Experience to the Investigator

To drill down from the CA APM Customer Experience console to the Workstation from an incident:

1. Select CEM > Incident Management.

   The Incidents page appears.

2. Click on the link for a specific slow time incident ID number.

   The Incident Overview page appears, for the incident you selected.

3. Click the Go to Introscope Workstation Web Start link.

   The link opens a Workstation page, displaying the Metric Browser tab of the Investigator with the related business services for the CA APM Customer Experience agent.

   Tip: If you link to the Introscope Console instead of the Investigator, check your Introscope configuration. See the section on troubleshooting in the *CA APM Integration Guide.*

4. In the Metric Browser tree, locate the Business Segment node.

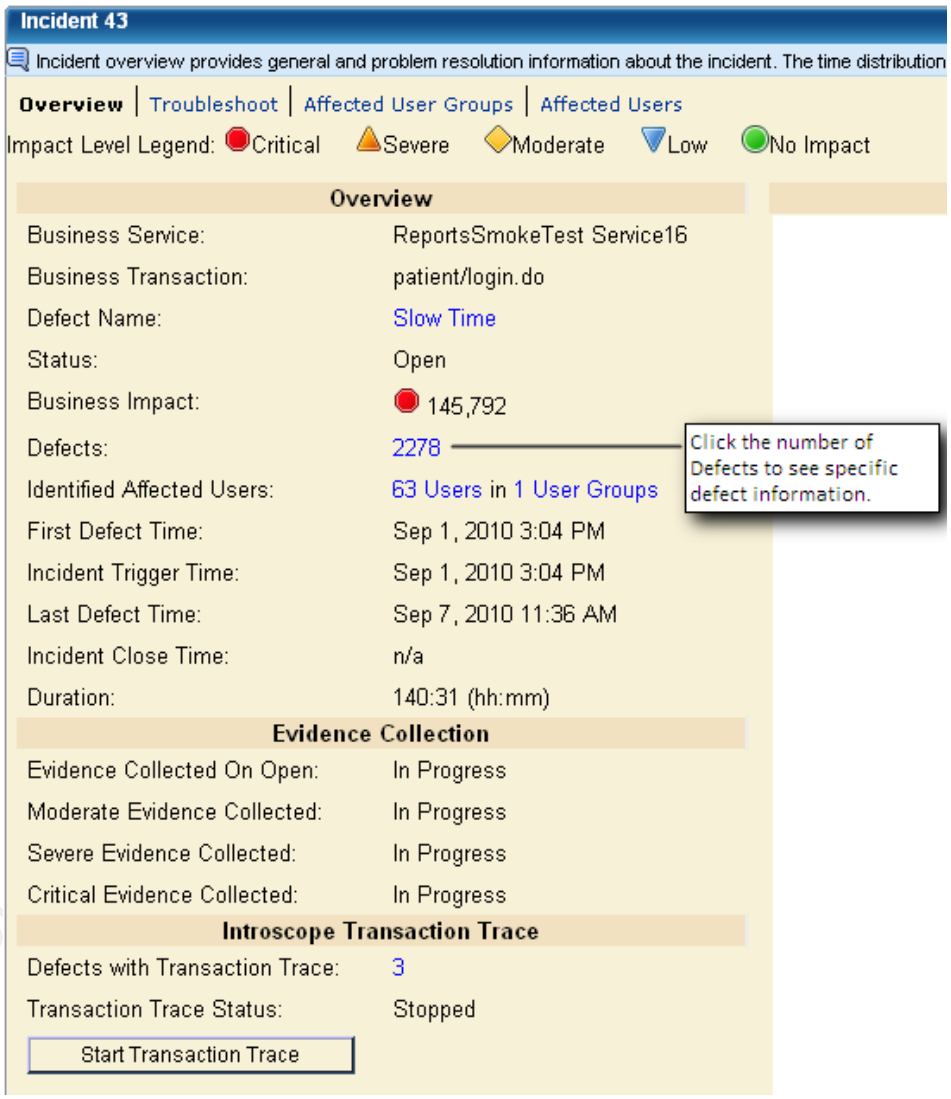
From here you can:

- Navigate and click through each node to analyze transaction trace metrics.
- Change the Investigator Time range to view specific metrics.
- Further analyze the incident, see Incident troubleshooting to find root cause (see page 196).
- If you want to view a specific defect for the CA APM Customer Experience incident, click the Defects link. Continue to the next section.

CA Technologies Prerelease Documentation. Authorized Use Only

To drill down from the CA APM Customer Experience console to the Workstation from a defect:

1. From the Incident Overview page, you can click the link for the number of Defects.



   The defects list appears, with all the defects related to the incident.

2. Select from the View list to see Introscope, Customer Experience, or All defect-related information.
3. Search to find the defect you want to review.
4. Click on the link for a specific slow time defect date and time. (You can also click the Yes link in the Transaction Trace column.)

   The defect details page appears.

5. Scroll down to view transaction trace and HTTP information.

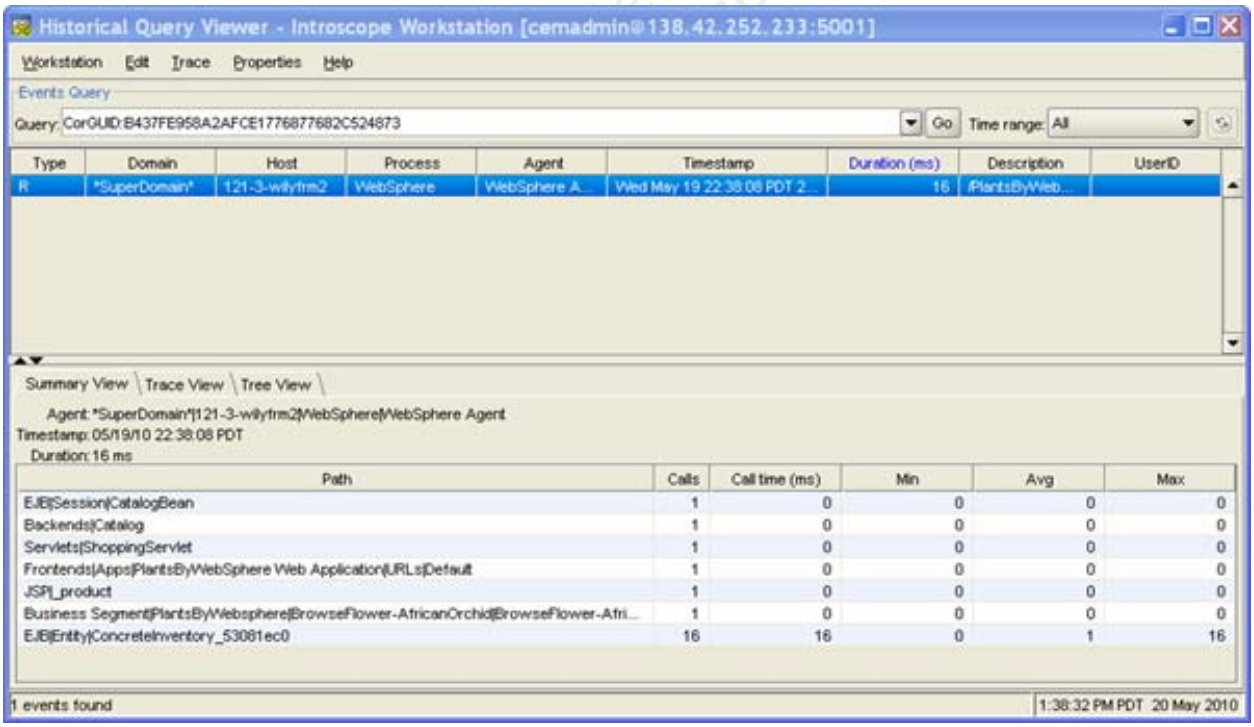


The Transaction Trace shows Introscope Blame data pointing to the possible root cause of the slow time defect from the servlet component to the backend component.

The HTML information allows you to see the HTML that the customer viewed on his or her computer at the time of the defect.

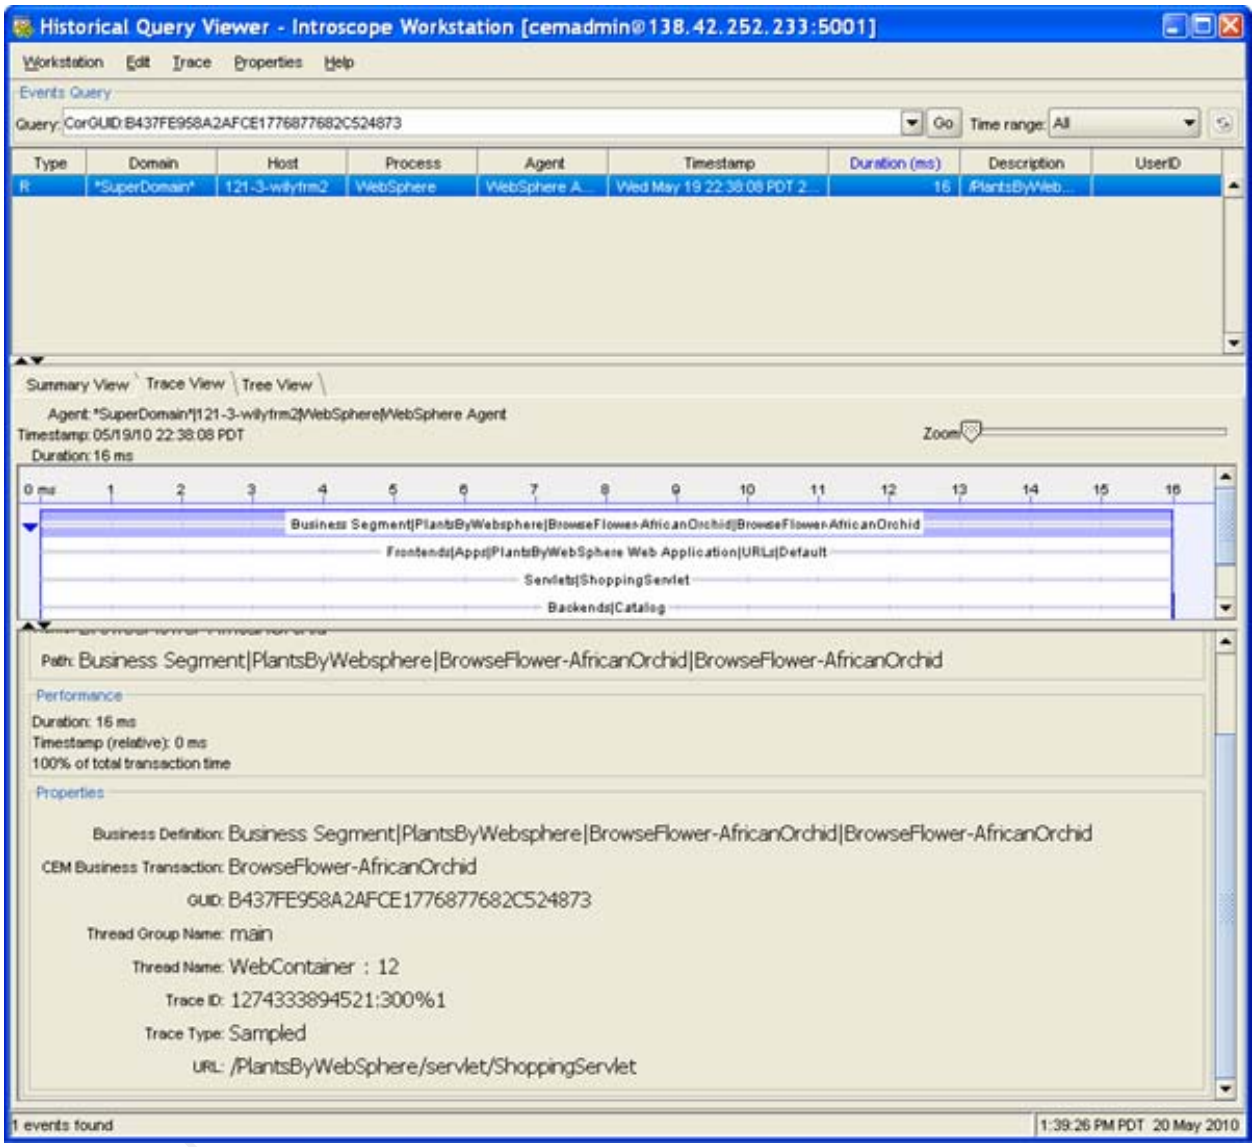
Note: Images and other elements of the page might not appear.

6. If you need more information (for example, errors or stalls), click the Go to Introscope Workstation Web Start link.



The link opens the Workstation Historical Query Viewer, displaying the CorGUID query for the GUID, which correlates to the CA APM Customer Experience defect. The standard trace event information displays.

7. Click on a trace event to view transaction trace details.



Note: If CA APM Customer Experience and Introscope system times are not in sync, the defect information might not appear immediately.

If you are viewing a trace from a 7.x .NET agent, the GUID might appear in the ASP.NET layer of the Transaction Trace stack.

Using WebView to Find More Information About an Incident

WebView provides an alternative browser-based view when you do not have access to the Workstation.

For more information about WebView, see the *CA APM WebView User Guide.*

To drill down from the CA APM Customer Experience console to Introscope WebView from an incident:

1. Select CEM > Incident Management.

   The Incidents page appears.

2. Click on the link for a specific slow time incident ID number.

   The Incident Overview page appears for the incident you selected.

3. Click the Go to Introscope WebView link.

   The link opens WebView.

   Note: The landing page reflects the same context as the last time you were in WebView.

To drill down from the CA APM Customer Experience console to Introscope WebView from a defect:

1. Select CEM > Incident Management.

   The Incidents page appears.

2. Select Defects.

   The defects list appears.

3. Search to find the defect you want to review.

4. Click on the link for a specific slow time defect date and time.

   The defect details page appears.

5. Click the Go to Introscope WebView link.

   The link opens WebView.

   Note: The landing page reflects the same context as the last time you were in WebView.

CA Technologies Prerelease Documentation. Authorized User Only.

Incident troubleshooting to find root cause

Incident troubleshooting charts highlight problem areas to help you find the root cause of defects. Defect distribution data include a root cause probability overview and pie charts by tier:

- Overview (see page 198)
- Client tier
- Network tier (see page 199)
- Web server tier (see page 199)
- Application server tier
- Logic tier (see page 200)
- Backend tier

From the incident overview information, you can navigate to the troubleshooting pie charts to view root cause probability and defect counts per tier for each incident.

To troubleshoot incident information:

1. Select CEM > Incident Management.

   The Incidents page appears.

2. Click on the link for a specific slow time incident ID number.

   The Incident Overview page appears for the incident you selected.

3. Click the Troubleshoot link.

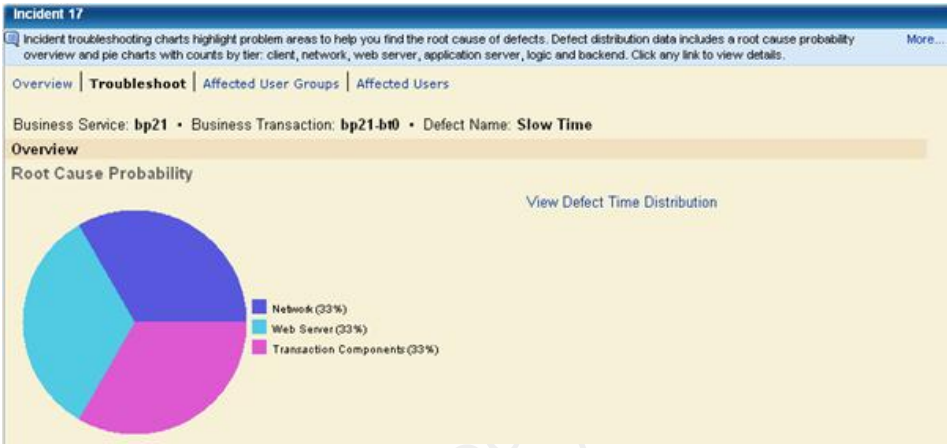
   Tip: Remember that the incident troubleshooting described here is available only for slow time incidents.

The pie chart in each tier represents the distribution of how many defects passed through the slice in that tier. For example, the Client pie chart shows how many defects in the incident were using each browser type. And if all the defects came from the same type of browser, then the problem might be browser-related.

CA Technologies Pre-release Documentation. Authorized use only.

Overview

The overview section provides clues about the potential root cause of the defects. It also displays the transaction components most involved in the slow time incident.



If one slice is much bigger than the others, then the probability is that the defect condition's root cause is in that tier.

Root cause probablity

Each slice in the overview chart represents the degree of variability in the pie for the corresponding data.

- A low level of variability (a bigger pie slice) suggests that a particular element is appearing consistently across many defects and therefore might be the root cause of the incident.
- A high level of variability (a smaller pie slice) can suggest normal operations, and therefore might not be the root cause. The percentage provides a clue as to the likelihood that this tier has an issue to focus on.

By summarizing this variation-evaluation process and computing it automatically, CA APM Customer Experience automates the analysis process you otherwise would need to perform yourself.

Transaction components

The transaction components chart displays, by default, the top 5 slowest transaction components associated with this incident. If there are more than 5 components, all the rest are combined into "Others (n)," where (n) shows the number of other transaction components involved.

CA Technologies Prerelease Documentation
Authorized Use Only

- Five evenly distributed (similar sized pie slices) components, plus one larger component (all the other components, besides the top five) suggests that no one transaction component is the root cause of the defects.
- An uneven distribution (one or more large pie slices) suggests that one transaction component might be involved in many of the defects. Look for the named component with the higher average transaction time.

Client tier

The client tier shows defect count by client and browser type.

- An even distribution (many small pie slices) suggests problems affect customers with all kinds of browsers.
- An uneven distribution (one or more large pie slices) suggests problems affect customers with the same kind of browsers or the total population uses a small set of browsers.

Network tier

The network tier shows the actual IP address ranges impacted by this incident, based on the client IP addresses observed in the monitored transaction.

- An even distribution (many small pie slices) suggests problems for customers in many locations.
- An uneven distribution (one or more large pie slices) suggests a problem affecting customers only in certain IP subnets. When it is just one IP range, it is clearly a geographic issue.

Web server tier

The web server tier displays defect count, showing which web servers served the defective requests (IP Address | MAC Address).

- An even distribution (many small pie slices) suggests defective transactions pass through a variety of web servers.
- An uneven distribution (one or more large pie slices) suggests a potential problem with a particular web server or set of web servers, or that the load balancer is not evenly distributing requests.

Application server tier

The application server tier displays defect count, showing which application server instances served defective transactions (Agent Hostname | Agent Process | Agent Name).

CA Technologies Internal Documentation. Authorized Use Only.

- An even distribution (many small pie slices) suggests defective transactions are served by a variety of application server processes.
- An uneven distribution (one or more large pie slices) suggests a problem with a particular application server process, or that the load balancer or web server is not evenly distributing requests.

For more information about viewing details of transactions using the Transaction Tracer, read Using the Transaction Tracer (see page 203).

Logic tier

The logic tier shows the logic component responsible for the most time in serving the request. For more information, see Incident-related terminology.

- An uneven distribution (one or more large pie slices) suggests one or more suspected Blame (logic) components that are the slowest method (program code) in a transaction. This might be a common scenario since all the defects in an incident are the same type of transaction and are often processed by the *same* code path.
- An even distribution (many small pie slices) suggests many different logic components contributed long processing times to the defective transaction. Since only one logic component is chosen (the suspected Blame component) for each defective transaction, this means that defective transactions use many different code paths; and none of them is consistently the slowest in defective transactions.

For more information about viewing details of transactions using the Transaction Tracer, read Using the Transaction Tracer (see page 203).

Backend tier

The backend tier shows the backend component responsible for the most backend time spent processing defective transactions. For more information, see Incident-related terminology.

- An uneven distribution (one or more large pie slices) suggests one or more suspected Blame backend components that are the slowest in a transaction. This might be a common scenario since all the defects in an incident are the *same* type of transaction and are often processed using the same backend.
- An even distribution (many small pie slices) suggests many different backend components contributed long processing times to the defective transaction. Since only one backend is chosen (the suspected Blame backend component) for each defective transaction, this means that defective transactions use many different backend components and none of them is consistently the slowest in defective transactions.

For more information about viewing details of transactions using the Transaction Tracer, read Using the Transaction Tracer (see page 203).

CA Technologies Internal Use Only. Unauthorized Distribution.

View Defect Time Distribution

You can view the distribution of time to determine where time is being spent in each defective transaction.

To view defect time information:

1. Select CEM > Incident Management.

   The Incidents page appears.

2. Click on the link for a specific slow time incident ID number.

   The Incident Overview page appears for the incident you selected.

3. Click the Troubleshoot link.

4. Click on View Defect Time Distribution links to view graphs specific to that tier. You can:

   - Change the graphs view by selecting a different Type (of time) from the list.
   - Scroll down to see the various available graphs for that defect condition.

Incedent troubleshooting hints

The following tips will help you navigate the incident troubleshooting pages to find the root cause of incidents.

- The overview pie charts show *variation*, where tier pie charts show defect *counts* for this incident.
- If you see no link when data appear in a pie chart, it is because you are not viewing a slow time incident. (Links show slow-time-related data.)
- If you see no data, there are two possible reasons: either all values are zero, or there are no data points.
- Tiers show *by category* (where): client, network (user group), web, application, logic, backend.
- Defect time distribution (links to data) shows *slow time* distribution.
- Type drop-down list shows the *type of time*: client, CEM transaction, application server, logic, backend.
- If graphs / links do not appear, there are no data points for the *tier category* (where) or *type* (of time).

Viewing tier category data

From the Incident > Troubleshoot page, you can click on any of the Show Data links to view tier category (where) data in table format.

CA Technologies Prerelease Documentation. Authorized Use Only

Chapter 5: Using the Introscope Transaction Tracer

Workstation users with appropriate permissions use Introscope Transaction Tracer to trace the activity of transactions as they flow through a Java Virtual Machine, or a Common Language Runtime (CLR) in a .NET environment, inside a production application.

This section contains the following topics:

About the Transaction Tracer (see page 203)
Starting, Stopping, and Restarting a Transaction Trace (see page 205)
Transaction Trace session options (see page 208)
Using the Transaction Trace Viewer (see page 209)
Using Dynamic Instrumentation (see page 217)
Printing a Transaction Trace window (see page 229)
Querying Stored Events (see page 229)
Saving and exporting Transaction Trace information (see page 234)

CA Technologies Prerelease Documentation
Authorized Use Only.

About the Transaction Tracer

Introscope defines a transaction as the invocation and processing of a service. In the context of a web application, it is the invocation and processing of a URL sent from a web browser. In the context of a web service, it is the invocation and processing of a SOAP message.

Introscope Transaction Tracer reduces the time required to identify a problem component in a transaction, enabling authorized users to trace the transaction activity at the component level. Transaction Tracer can trace synchronous transactions that cross boundaries in the homogeneous application server environments that support this capability:

- WebLogic Server 8.0 and later
- WebSphere 6.x.

In other environments, transactions can be traced within the boundaries of a single Virtual Machine (VM) or Common Language Runtime (CLR).

You view the results of a cross-process Transaction Trace query in the Trace View tab of the Transaction Trace Viewer.

Introscope saves Transaction Trace session data in the Transaction Events Database for a specified amount of time, and periodically aged out to reduce overhead.

You can configure the Introscope agent to capture Transaction Trace data based on the values of servlet or ASP.NET variables such as HTTP request headers, request parameters, session attributes, session ID, username, URLs and URL Query strings. In addition, Introscope agents automatically sample transactions; see Automatic Transaction Trace sampling (see page 204), below.

Note: Metric Shutoff state does not affect Transaction Trace data. If a managed agent is shut off, that agent does not report Transaction Trace data. If the agent is shut off while a Transaction Trace session is in progress, the agent does report the data collected before the shutoff request.

Automatic Transaction Trace Sampling

By default, Introscope agents sample transaction behavior by tracing each normalized unique URL in an application once per hour. You can view and analyze sampled traces from a selected historical time range:

- in the Introscope Workstation and WebView
- in the Traces tab in the Investigator

You can also configure automatic trace sampling even if no URL groups are configured by specifying the number of transactions to sample during a time interval; the default value is one transaction every two minutes. For more information, see the *CA APM Java Agent Implementation Guide*.

Transaction trace sampling is enabled by default. You can disable the behavior, change the sampling period, or de-randomize the timing of sampling as appropriate. For more information, see the discussion of Controlling Automatic Transaction Tracing Behavior in the *CA APM Java Agent Implementation Guide* and *CA APM .NET Agent Implementation Guide* as appropriate.

Transaction Trace overhead

A Transaction Trace session affects overhead from the time it starts until all transactions in process at the end of the session complete. You can specify the execution threshold at the millisecond level, but doing so increases the load on the system.

CA Technologies Pre-release Documentation. Authorized Use Only.

These Transaction Tracer features reduce the likelihood of trace sessions imposing unacceptable overhead:

- **Transaction Trace Session Timeout—**A Transaction Trace session times out after a user-defined period so that the Admin user cannot accidentally leave the Transaction Tracer on and negatively affect performance for a sustained period. At the end of the timeout period, the agent stops tracing new transactions and completes tracing for transactions in progress.
- **Anti-Flooding Logic—**To prevent excessive overhead, agent anti-flooding logic limits the number of transactions traced per 15 second interval to 200. After this limit is exceeded, the agent logs that the anti-flood threshold was exceeded, and does not report Transaction Trace data to the Enterprise Manager until that 15-second period has expired. After the 15-second period expires, the anti-flooding logic resumes reporting.

The *APM Sizing and Performance Guide* has more information about controlling Transaction Trace overhead.

Transaction Tracer compatibility with agents from previous releases

Introscope version 9.0 with Transaction Tracer enabled is compatible with agents from versions before 9.0, with these caveats:

- When you use Transaction Tracer with agents from version 5.3.2 and later, you can filter on parameters and threshold execution time.
- With 6.0 and later agents, Transaction Tracer can filter by errors, in addition to parameters and threshold execution time.

Starting, Stopping, and Restarting a Transaction Trace

To run a Transaction Trace session, you specify the agents whose transactions you want to trace, and how long to capture the data. You can specify filter options to limit tracing to transactions that:

- exceed a threshold execution time you define
- match parameter values such as User ID, request headers information, etc.
- have errors, if ErrorDetector is enabled

When the Transaction Trace Session starts, Introscope captures Transaction Trace data that is specified in the agent profile, for each transaction. The transactions that match the filter criteria appear in the Transaction Trace Viewer window, and are saved in the Transaction Events database.

CA Technologies Proprietary Documentation
Authorized Use Only.

Note: You can start Transaction Trace using a CLW (Command Line Workstation) command. For information about the command and its syntax, see the *APM Configuration and Administration Guide*.

Starting a Transaction Trace Session

To start a Transaction Trace session:

1. Select Workstation > New Transaction Trace Session.

   The New Transaction Trace Session window opens.

2. In the Trace transactions section of the window, specify the minimum duration for transactions to be traced. Select milliseconds or seconds from the drop-down list.

   Note: Sub-second durations can have a negative impact on performance.

3. To specify a transaction filter, click the checkbox to the left of the dimmed drop-down menu reading "User ID" in the *Trace transactions* section, and select a type from the list:

   - User ID—enter an operator and a parameter value.
   - Session ID—enter an operator and a parameter value.
   - URL, or URL Query—enter an operator and a parameter value.
   - Request Header—enter a data type name, a condition, and a value.
   - Request Parameter—enter a data type name, an operator, and a parameter value.
   - Session Attribute—enter a data type name, an operator, and a parameter value.

   Note: Data is only available for use in filters if the Introscope agent is configured to capture it. See the discussion of Controlling Automatic Transaction Tracing Behavior in the *CA APM Java Agent Implementation Guide* and *CA APM .NET Agent Implementation Guide* as appropriate for your environment.

   The table below lists the filter conditions:

Filter Condition	Condition Effect
equals	Transactions in which the parameter value matches the string specified are traced.
does not equal	Transactions in which the parameter value does not match the specified string are traced. Note: Transactions that do not include the parameter to which the filter applies are also traced.

CA Technologies Preliminary Documentation. Authorized Use Only.

contains	Transactions in which the parameter value contains the specified string are traced.
does not contain	Transactions in which the parameter value does not contain the specified string are traced. Note: Transactions that do not include the parameter to which the filter applies are also traced.
starts with	Transactions in which the parameter value starts with the specified string are traced.
ends with	Transactions in which the parameter value ends with the specified string are traced.
exists	Transactions that include the parameter to which the filter applies are traced, regardless of the parameter value.
does not exist	Transactions that do not include the parameter to which the filter applies are traced.

4. Enter the trace session duration in minutes.
5. In the Trace Agents section, select one or more agents for which to trace transactions:
   - To trace all agents that support Transaction Tracing, click Trace all supported Agents. This option traces supported agents that are currently connected, and any that connect during the Trace session.
   - To trace selected agents, click Trace selected Agent(s) and select agents from the list (CTRL + click to select multiple agents).
6. Click OK to start the Transaction Trace session.

   Transaction Trace results appear in the Transaction Trace Viewer window. For more information see Using the Transaction Trace Viewer (see page 209).

Stopping a Transaction Trace session

To stop a Transaction Trace session:

- Click Stop, or
- Select Trace > Stop Tracing Session.

Restarting a Transaction Trace session

Restarting the Transaction Trace session resets the timeout to the user-defined time period and continues to trace Transactions in the targeted agents using the same threshold criteria.

CA Technologies Pre-release Documentation. Authorized use only.

You can restart a Transaction Trace session:

- after a session has timed out.
- to restart a session you have stopped.
- to restart an in-progress session.

To restart a Transaction Trace session:

- Click Restart, or
- Select Trace > Restart Tracing Session.

Transaction Trace session options

Options for your transaction trace session include:

- The ability to turn off low-threshold execution time warnings (see page 208)
- The ability to review agents targeted for tracing (see page 208)

Turn Off Low-Threshold Execution Time Warnings

If you are running the Transaction Tracer and set the threshold execution time to less than one second—to perform a deep analysis, for example—you might see continual warnings. The warnings indicate increased overhead because of increased traces, so you might want to turn them off in a production environment.

To turn off the warnings about low-threshold execution time:

1. Select Workstation > User Preferences.
2. Select the Transaction Tracer tab.
3. Select *Don't warn when threshold is less than 1 second.*
4. Click *Apply.*

Reviewing agents targeted for tracing

To review the agents targeted for tracing:

1. Select Trace > Show Traced Agents.

   The Tracing Agent(s) dialog box appears.
2. When you are finished viewing the Tracing Agent(s) information, click OK.

CA Technologies Confidential Documentation. Authorized Use Only

Using the Transaction Trace Viewer

The Transaction Trace Viewer shows trace information for transactions that meet the criteria you specified for the trace session.

The table in the top pane of the Transaction Trace viewer lists transactions that were traced during the session. You can sort the rows by column by clicking on the column header. New transactions are inserted into the table in sorted order.

This table lists the columns in the transaction table:

Field	Description
Type	The type of information in the trace row, one of: ■ Transaction Trace (T) ■ Error (E) ■ Sampled (R)—A transaction chosen by random sampling. ■ Stalled (S)—A stalled transaction Error data only appears if ErrorDetector is enabled. Asterisk—If an asterisk appears after the type symbol, it means that some of the components in the transaction were truncated, or clamped. See Clamped transactions (see page 214). Only transactions of types T and E can be clamped. The types above apply to transactions available in Live mode. When querying historical transactions, other transaction types are available. See Query options and syntax (see page 231).
Domain	Domain to which the traced agent is mapped
Host	Host on which the traced agent is running
Process	Agent Process name
Agent	Agent Name
Timestamp	Start time, in the agent computer's system clock, of the invocation of the root component
Duration	Wall clock execution time of the root component
Description	The URL that was invoked to initiate this transaction, or the Introscope path to the component that initiated the transaction.
UserID	The ID of the logged-in user that is running the transaction (if it is configured and available)

CA Technologies Proprietary Documentation. Authorized Use Only.

The Transaction Tracer window includes three tabs:

- Summary view (see page 210)
- Trace view (see page 211)
- Tree view (see page 216)

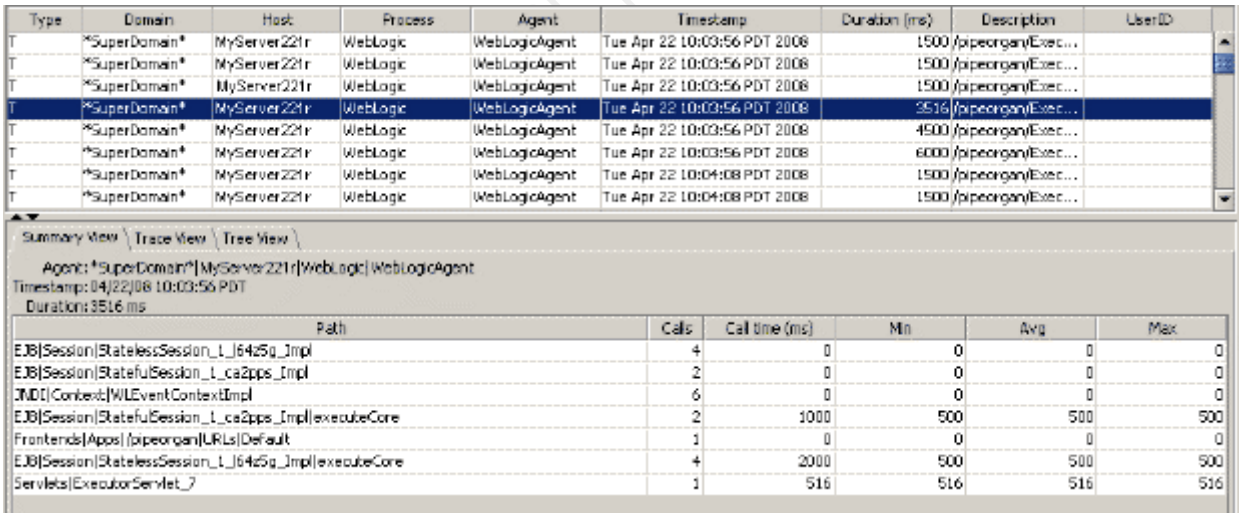
Summary view

The first time you select a transaction in the transaction table, the Summary View opens. When you select a transaction that has been opened before, it opens in the most recently selected view.

This information appears for the currently selected transaction in each tab:

- the fully qualified agent name
- start time, in the agent machine's system clock, of the invocation of the root component
- execution time of the root component in milliseconds

Summary View shows metrics for the components in the selected transaction. Metrics include the path, number of calls, the length of the call in milliseconds, and the minimum, average, and maximum call times.



Type	Domain	Host	Process	Agent	Timestamp	Duration (ms)	Description	UserID
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:03:56 PDT 2008	1500	/pipeorgan/Exec...	
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:03:56 PDT 2008	1500	/pipeorgan/Exec...	
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:03:56 PDT 2008	1500	/pipeorgan/Exec...	
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:03:56 PDT 2008	3516	/pipeorgan/Exec...	
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:03:56 PDT 2008	4500	/pipeorgan/Exec...	
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:03:56 PDT 2008	6000	/pipeorgan/Exec...	
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:04:08 PDT 2008	1500	/pipeorgan/Exec...	
T	*SuperDomain*	MyServer221r	WebLogic	WebLogicAgent	Tue Apr 22 10:04:08 PDT 2008	1500	/pipeorgan/Exec...	

Summary View | Trace View | Tree View

Agent: *SuperDomain*|MyServer221r|WebLogic|WebLogicAgent
Timestamp: 04/22/08 10:03:56 PDT
Duration: 3516 ms

Path	Calls	Call time (ms)	Min	Avg	Max
EJB|Session|StatelessSession_1_j64z5g_Impl	4	0	0	0	0
EJB|Session|StatefulSession_1_ca2pps_Impl	2	0	0	0	0
JNDI|Context|WLEventContextImpl	6	0	0	0	0
EJB|Session|StatefulSession_1_ca2pps_Impl|executeCore	2	1000	500	500	500
Frontends|Apps|/pipeorgan|URLs|Default	1	0	0	0	0
EJB|Session|StatelessSession_1_j64z5g_Impl|executeCore	4	2000	500	500	500
Servlets|ExecutorServlet_7	1	516	516	516	516

Release Documentation. Use Only.

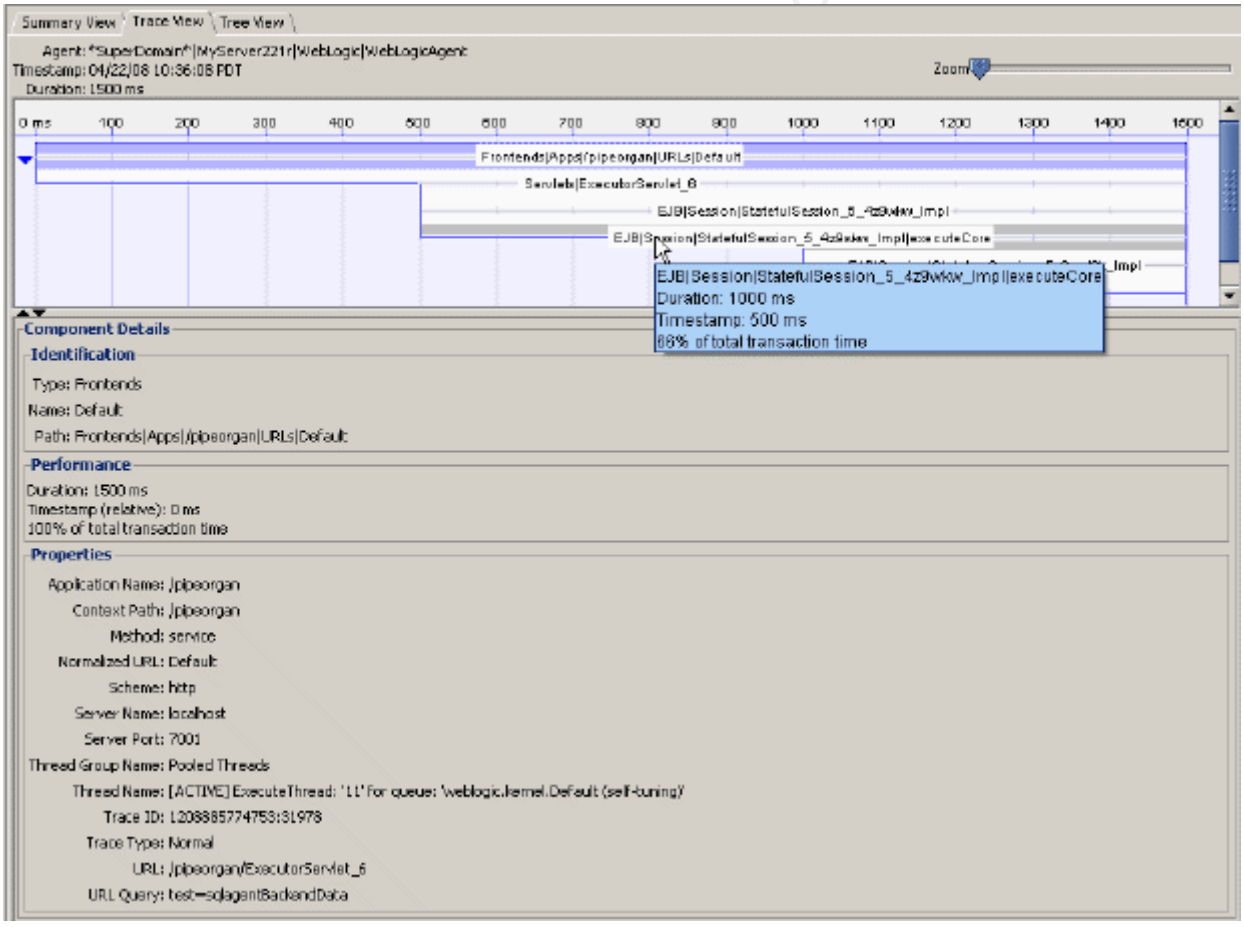
You can double-click one of the metrics listed in the table view to open the metric in the Browse tree.

At the bottom of the Trace window, the Transaction Trace status bar shows:

- the number of transactions that were collected in the session.
- the filter criteria for the Transaction Trace session.
- the remaining time before the current session times out.

Trace view

Trace View shows a selected transaction in a graphical stack display, sometimes referred to as an "upside down wedding cake" display, of the components which make up a transaction. When you select one of the components, you can see component details in the bottom pane of the viewer.



The Trace View shows:

- each component in the transaction as a bar
- the percentage of total transaction execution time for each component
- the calling relationships between components—the bars for components are displayed from top to bottom in calling order.
- transaction sequence over time—the placement of components from left to right indicates sequence. Relative wall clock time in milliseconds appears across the top of the Transaction Snapshot.
- errors within transactions (if ErrorDetector is enabled): red slices in the Transaction Snapshot represent errors within transactions. (See Reading and understanding error metrics (see page 407).)

Note: The default time range for traces in live mode is 20 minutes. Traces older than 20 minutes are not displayed in live mode; they will be aged out (not shown) after they are more than 20 minutes old.

In the Trace View you can:

- Hover your mouse pointer over a component to open a tooltip, as shown above. See Tooltips (see page 70).
- Right-click a component to open the Investigator and view component metrics.
- Right-click a component to instrument, in runtime, one, many or all of the methods the component calls. See Using dynamic instrumentation (see page 217) for more information about dynamically instrumenting methods.
- Select a component in the Trace View to open the Transaction Component Details pane.

Transaction Component Details

The component details of the Trace View shows this information:

- **Type—**High-level component (for example, EJB, Servlet, JSP in Java, and ASPX in .NET).
- **Name—**Name of the component.
- **Path—**Full resource name of component.
- **Duration—**Execution time of the selected component. Default unit is milliseconds; this can be set to other units (see Setting the Duration unit).
- **Timestamp (relative)—**Start time, in the agent host computer's system clock, of the invocation of the selected component.

CA Technologies Pre-release Documentation
Authorized Use Only.

- **% of total transaction time—**Percentage of total transaction time taken by selected component.
- **Properties—**Any optional properties reported by the component (for example, URL, URL Query, Dynamic SQL), or defined for collection in the Introscope agent profile (User ID, Request Header, Request Parameter or Session Attribute). You can select the text of any field in the Properties details and copy it using the keyboard commands CTRL+C.

Property	Description
User ID (Servlet, JSP, ASPX)	User ID of the user invoking the HTTP servlet request.
URL (Servlet, JSP, ASPX)	URL passed through to the servlet or JSP, not including the query string (text after the '?' delimiter in the URL
URL Query (Servlet, JSP, ASPX)	Portion of the URL that specifies query parameters in the HTTP request (text after the '?' delimiter in the URL)
Session ID (Servlet, JSP, ASPX)	The HTTP session ID associated with the servlet request, if any.
Dynamic SQL (Dynamic JDBC or ADO.NET Statements, when SQL Agent is installed)	Generalized dynamic SQL statement, as it would be seen in the aggregate form in the SQL Agent
Callable SQL (Callable JDBC or ADO.NET statements, when SQL Agent is installed)	Callable SQL (with the '?' still present)
Prepared SQL (Prepared JDBC or ADO.NET statements, when SQL Agent is installed)	Prepared SQL (with the '?' still present)
Method (Blamed Tracers; everything but servlets, JSPs and JDBC statements for Java, ASPX and ADO.NET for .NET)	Name of the traced method

CA Technologies Pre-release Unauthorized Use Documentation.

Tooltips in the Transaction Trace viewer

Hovering your cursor over any of the individual components, or layers, of the graphical depiction of a transaction provides details about the component in a tooltip.

The tooltip displays:

- Path
- Duration
- Timestamp (relative)
- % of total transaction time

See Transaction component details (see page 212) for definitions of this information.

WYISC--Sequence View for WUG

The Sequence View tab displays a transaction's components in the order in which they are called by a process.

This view is available if you have installed the SOA Performance Management extension components. For more information about this view, see the section on using the Sequence View in the *CA APM for SOA Implementation Guide.*

Correlation IDs in cross-process transactions

Introscope Workstation uses a unique identifier, the correlation ID, to link traced frontend and backend transactions. The sequencing of this ID is determined by the order in which frontends call backends in a transaction.

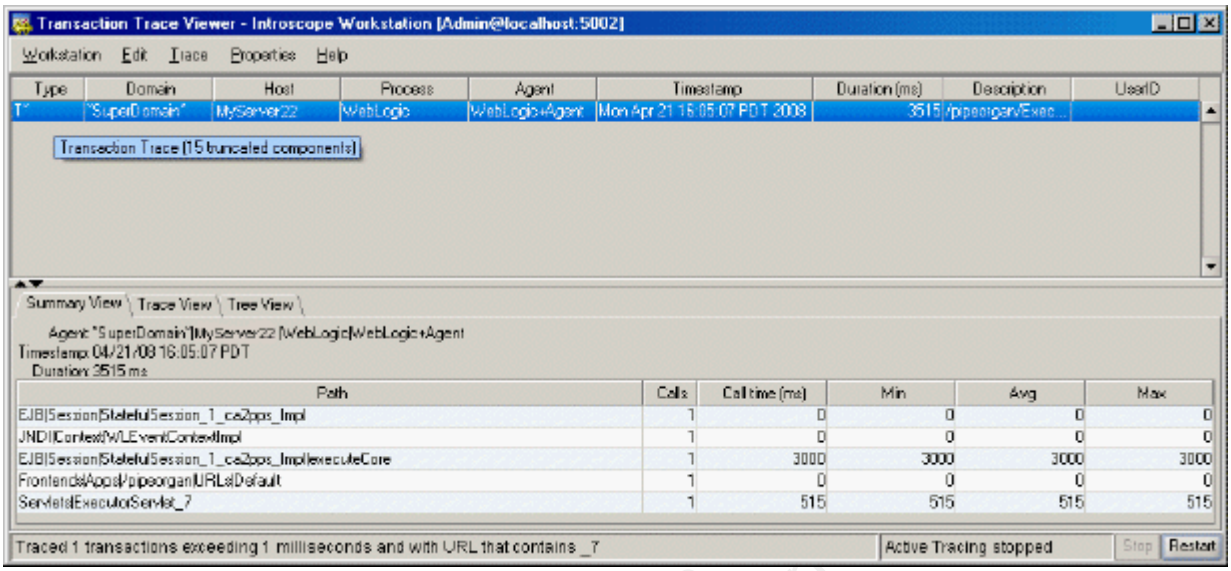
By using this correlation ID to recognize and trace the path of linked components in a transaction trace, you can get insight into which calls might be the source of a slow or stalled transaction.

Clamped Transactions

To prevent unusual Transaction Trace results from consuming too many cycles, a clamp on Transaction Trace components is set by default at 5000. (This setting, *introscope.agent.transactiontrace.componentCountClamp*, is specified in *IntroscopeAgent.profile*. For more information about working with the properties in this file, the discussion of Controlling Automatic Transaction Tracing Behavior in the *CA APM Java Agent Implementation Guide* or *CA APM .NET Agent Implementation Guide* as appropriate.)

CA Technologies Pre-release Documentation
Authorized Use Only

For traces producing clamped components—those exceeding the CountClamp—traces will be marked with an asterisk, as in the first row of the screenshot below:



Things to notice:

- The first row of traces is selected.
- The Type symbol is marked with an asterisk, signifying that some of the components in the transaction were truncated, or clamped.
- A tooltip indicates how many components were truncated. In the example above, 15 of the components of the selected trace exceeded the number specified in the *introscope.agent.transactiontrace.componentCountClamp* property.
- The components which were not truncated appear in the Summary View tab at the bottom of the viewer.

Note: Each agent has an IsClamped heuristic value, with 0=not clamped, and 1=clamped.

Appearance of exported XML file when transactions are clamped

When a trace component is clamped, the exported XML file will be well formed, and will include a parameter like:

```
<Parameter Value="15" Name="Components Not Shown"/>
```

To see a tooltip with more information about a trace:

1. Select one of the traces in the table.
2. Hover your cursor over the selected trace.

   The tooltip displays trace type and number of truncated, or clamped, components.

CA Technologies Proprietary Pre-release Access Not to be Used or Disclosed

To sort the traces by type:

- Click the heading of the Type column in the table.

Searching for clamped transactions

You can search for clamped transactions by issuing a historical event query. Following the instructions for querying historical transactions in Querying historical events (see page 230), use a string like this in your query:

```
componentsNotShown:[1 TO 9999]
```

This will help ensure that traces that had clamped transactions will be returned by the query.

Note: Because the historical event viewer search uses Lucene syntax, note:

- The word TO in the string is case sensitive.
- The search syntax is lexicographical, not numerical. For this reason, performing historical queries using "componentNotShown" as a query filter may return incorrect results.
- Strings beginning with * (asterisk) or ? (question mark) are not allowed.

Viewing errors with Transaction Tracer

You can use the Transaction Tracer to identify and view errors. This functionality is available if you have enabled ErrorDetector,which is discussed in the Extensions appendix in the section Viewing errors using the Transaction Tracer (see page 410).

About the Tree view in Transaction Tracer

Tree view is a hierarchical view of the transaction's components.



In the illustration, notice that three methods contribute to the selected transaction. Of the three, the third method, *runRequestCycle*, is decorated with a red indicator, and took 100% of the 1453ms it took the transaction to run. With that method selected, the tab displays additional information about the method in the Component Details pane.

Technologies, Inc. Prerelease Documentation. Authorized Use Only

Trace components that do not contribute a significant amount of time to the transaction are color-coded with a green icon.

To remove these components and view only the key transaction components:

- Choose Trace > Transaction Filter.

Aggregated Data for Multiple Transactions

In Transaction Tracer, you can select multiple transactions to see a representation of all components in the traces.

To view aggregated data:

1. Open a list of transactions by running a Transaction Trace and viewing them (see Using the Transaction Trace Viewer (see page 209)), or querying for them (see Querying stored events (see page 229)).
2. Select multiple transactions using CTRL-click or SHIFT-click.
3. Open the Summary or Tree view to see the transaction data aggregated.
   - Transaction Tracer shows the aggregated data in the table (it might be necessary for you to scroll down to see all the data).
   - The Tree View shows the aggregated data.

     In the Tree view, Transaction Tracer adds a node named Root if the selected transactions don't share a common root node.

Using Dynamic Instrumentation

Instrumenting a method means attaching byte code to the method, thereby enabling Introscope to monitor several aspects of the method's performance. (For background information about what instrumentation means and what it allows you to do, see the *CA APM Java Agent Implementation Guide* or *CA APM .NET Agent Implementation Guide* as appropriate for your environment.)

Instrumenting a method *dynamically* means inserting the instrumentation during runtime, without the need to restart the application server.

You can dynamically instrument one, more or all of the methods during a transaction trace session, and subsequently view metrics returned by the newly instrumented methods. This allows you to do dynamic application performing tuning.

Note: Only users whose administrators have granted them certain (usually administrative) permissions can use this functionality. Permissions are controlled in the *domains.xml* file. For more information, see the *CA APM Configuration Administration Guide.*

CA Technologies Product Reference Documentation. Authorized use only

When you instrument one or more methods through the Transaction Trace View:

- The instrumentation is temporary, lasting only for the duration of the transaction trace session.
- The instrumentation can be made permanent through a dialog, without the need to manually create a *.pbd* file.
- The instrumentation provides the five standard Introscope metrics (see page 349), viewable in the Investigator tree under the metric browser tab.

Note to users on the .NET operating environment

Dynamic instrumentation is supported on the.NET operating environment with limited functionality. Each of the topics in this section contains guidance on the extent to which the functionality outlined in the topic is supported on .NET.

In this section

This section tells you how to do the following tasks:

- View and instrument one, more, or all the called methods belonging to a transaction component. See Temporarily instrumenting one, more or all called methods (see page 219).
- View traces on temporarily instrumented methods. See Viewing and understanding traces on instrumented methods (see page 220).
- View the metrics collected on temporarily instrumented methods. See Viewing metrics collected on a temporarily instrumented method (see page 221).
- Make temporary instrumentation permanent on one of the dynamically instrumented methods. See Converting temporary instrumentation to permanent (see page 222).
- Remove temporary or permanent instrumentation from dynamically instrumented methods. See Removing temporary or permanent instrumentation (see page 224).
- Save instrumentation to a file and import the instrumentation to other agents. See Exporting instrumentation (see page 226).
- Adjust instrumentation on the tracer group level. See Modifying instrumentation level (see page 227).

Note: After executing one of the instrumentation changes described in this section, the agent may take several seconds to process the change. During this interim, no additional dynamic instrumentation changes can be made during this interim until the agent is finished making the changes. You may see an error message if you try to execute additional instrumentation changes.

CA Technologies Pre-release Documentation. Authorized Use Only.

Temporarily Instrumenting One, More or All Called Methods

Using the Transaction Trace Viewer, you can see the methods called by a selected trace component and temporarily instrument one or more of them.

To temporarily instrument one or more called methods:

1. Start a transaction trace (see Starting, stopping, and restarting a Transaction Trace (see page 205)).
2. When transactions begin to appear, click the Trace View tab.
3. Select one of the transactions displayed in the transaction table.

   When you select one of the transactions, the Transaction Trace Viewer displays the transaction's components in the viewer pane as a series of stacked bars sometimes referred to as the "upside-down wedding cake."
4. Right-click one of the components.
5. From the menu, select *View All Called Methods...*

   A dialog appears with a list of all the methods called by the selected transaction component.

   The dialog shows which methods are already instrumented, and which are Instrumentable, or available for instrumentation.
6. Select method to instrument.
7. Select *Add Instrumentation.*

You instrument one method at a time using the *View All Called Methods* dialog, and repeat the steps to instrument other methods.

Methods you selected to instrument will appear in the existing Transaction Trace as green segments, as long as this Transaction Trace is running. See Viewing and understanding traces on instrumented methods (see page 220).

The Stop button in the Transaction Trace viewer allows you to stop a trace before its time runs out. When you do this, any temporary instrumentation will disappear. This happens because, by definition, temporary instrumentation is not saved anywhere and lasts only for the duration of the trace.

To get information about instrumented methods:

1. Start a transaction trace.
2. When transactions begin to appear, select the Trace View tab.
3. Right-click one of the transactions displayed in the transaction table and select *View All Called Methods...*

CA Technologies Preliminary Documentation. Unauthorized Use Only.

4. In the *View All Called Methods* dialog, right-click one of the methods marked Instrumented and select *Get Instrumentation Info...*

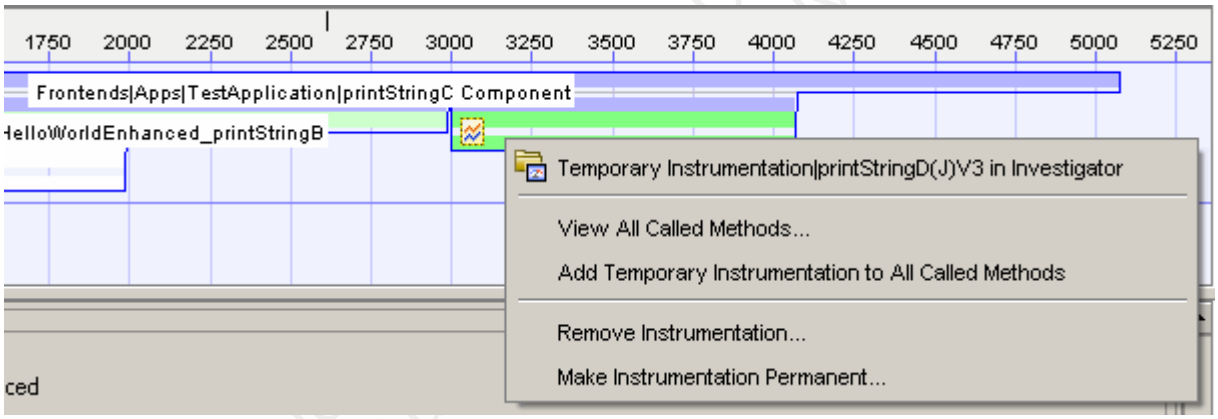
   A new dialog appears with information about the instrumented method.

5. Select the resource listed to get more information about the method.

Note: You can remove the method from instrumentation by right-clicking on the method and selecting Remove. For more information about removing instrumentation, see Removing temporary or permanent instrumentation (see page 224).

Viewing and understanding traces on instrumented methods

Temporarily instrumented methods appear in the Transaction Trace Viewer as green segments, as shown in the illustration below. Note there are two green segments; the one on the right is brighter green because it has been selected. A user has also right-clicked the segment to reveal the context menu, providing several actions that can be taken.



To understand the information they represent, note:

- Each instrumented method will appear only during the duration of the Transaction Trace session. Their instrumentation expires at the end of the session.
- Each instrumented method is identifiable by its name.
- You can identify problematic methods by noticing the duration (size) of the segment, as displayed in the Transaction Trace Viewer. The size of the segment is analogous to the time it took for the method to execute. Unexpectedly large segments are likely causes of slow transactions.
- You can hover your cursor over any of the methods that appear in the Trace View, and a tooltip will show metrics information about that method.

Note to .NET users

The option "Add Temporary Instrumentation to All Called Methods" is not available to .NET users.

Once you have identified problematic methods, you can:

- Convert temporary instrumentation on a method to permanent instrumentation. See Converting temporary instrumentation to permanent (see page 222).
- View metrics on the method. See Viewing metrics collected on a temporarily instrumented method (see page 221).
- Remove instrumentation from methods. See Removing temporary or permanent instrumentation (see page 224).

Viewing Metrics Collected on a Temporarily Instrumented Method

There are two ways to view metrics on temporarily instrumented methods:

- Hover your cursor over the method in the Trace View or in the View Detailed Instrumentation Info dialog. A tooltip shows metrics on the method.
- Jump from a segment in the Transaction Trace viewer directly to the Agent tree displayed in the metric browser tab. The node on the Agent tree will bear the name of the segment.

  Note: To name a segment/node, you configure its name in the Make Instrumentation Permanent dialog. See Converting temporary instrumentation to permanent (see page 222).

Note: Metrics are not collected on temporarily instrumented methods on .NET applications.

To view metrics in the metric browser tab:

1. Right-click the segment.

   A context menu appears. (You can see an illustration with this menu in the section entitled Viewing and understanding traces on instrumented methods (see page 220).

2. Select *<Method_Name> in Investigator*.

   A new Investigator window opens to the metric browser tab, with the node corresponding to the tracer segment highlighted.

CA Technologies Pre-release Documentation. Authorized Use Only

Convert Temporary Instrumentation to Permanent

After viewing the metrics from temporarily instrumented methods, you can make instrumentation permanent.

To convert temporary instrumentation to permanent:

1. In a Transaction Trace Viewer showing one or more temporarily instrumented methods (distinguishable by their green color and by the icon which signifies temporary instrumentation), right-click one of the segments.

   The temporary instrumentation icon is shown below:



2. Click *Make Instrumentation Permanent...*
3. In the Make Instrumentation Permanent dialog, enter the following information:

   Properties:

Property	Description
Node Name	Metrics will appear in the metric browser tab of the Investigator tree under this name.
Path (optional)	Metric path for the metrics. To create a new path, type the new path, like: Node Name|Subnode Name
Tracer Type	Tracer type to be used. Only the DynamicBlamePointTracer type is supported.
Calling Method	The name of the method that calls this class. Non-editable.
Instrumentation applies to all calling methods (check box)	Select to make permanence apply to all calling methods.

CA Technologies Prerelease Documentation
Authorized Use Only.

Narrow instrumentation to this class only (check box)	You can narrow this action to apply only to the selected class (i.e., the class represented by the selected Transaction Trace component), whose name is listed next to this option. The temporary instrumentation will be removed, and only this class will be assigned permanent instrumentation. By default, the "Make Permanent" action will apply to all classes derived from the original interface or abstract class. Note: When you narrow instrumentation using this option, you apply permanent instrumentation to only one class at a time. To apply to more than one class, repeat the steps in this section.

Groupings:

Optionally, you can assign the newly permanent instrumentation to an existing tracer set, or to a new one.

Property	Description
New Tracer Group	Select this to create a new tracer group, enter the name of the group, and assign the instrumentation to this group.
Existing Tracer Group	With this option selected, you can select one of the existing tracer groups from the drop-down selector.
New Label	Type the name of a new label to apply to the saved instrumentation. This string will be the basis for a the name of a new *.pbd* file.
Existing Label	With this option selected, you can select one of the existing labels from the drop-down selector. The labels correspond to existing *.pbd* files.

4. Click OK.
5. In the confirmation dialog, click OK or Cancel.

The newly created permanent instrumentation appears as follows:

- As a standard segment (still green colored) in the Transaction Trace Viewer without the "temporarily instrumented" icon.
- As metrics in the Investigator tree, in the location you specified.

CA Technologies Prerelease Documentation Authorized Use Only

Note to .NET users

Users on the .NET operating environment must restart their .NET application to see newly created permanent instrumentation as described.

Important: You should be careful about the level of instrumentation you export and subsequently use. The "TraceAllMethods" option should never be used in a production environment because of the performance implications; it is intended to be used in preproduction only, for the purpose of initially creating a custom PBD which is then pared down. Good practice is to use the search function to filter and pare down instrumentation before you export it from your sandbox environment to the test environment and certainly before using it in a production environment.

Note: The "Trace All Methods" option mentioned in the paragraph above is not available to .NET users.

Notes about instrumentation that has been made permanent

When you make instrumentation permanent, the instrumentation is saved in a PBD in the Dynamic directory. This directory will be automatically created if it doesn't exist yet. Existing PBDs are not overwritten.

Only one temporarily instrumented method at a time can be converted to permanent instrumentation.

Removing Temporary or Permanent Instrumentation

After viewing the metrics returned by instrumented methods, you can remove instrumentation from the methods.

When you remove permanent instrumentation through options in the context menu, the newly created PDB in the Dynamic directory is deleted if the PBD has no other instrumentation.

Note: The method of removing instrumentation mentioned in the paragraph above is not available to .NET users.

CA Technologies Prerelease Documentation
Authorized Use Only

You can remove instrumentation in either of three ways:

- By selecting a component from the graphical "wedding cake" view in the Transaction Trace Viewer.

  Note: Only this method is available to .NET users. The two other methods below are not.

- By selecting a row in the View Detailed Instrumentation Info dialog and clicking *Remove.*
- By removing labels from instrumented classes and methods, starting from the Investigator tree.

In addition, you can remove temporary instrumentation simply by stopping the Transaction Trace.

Remove instrumentation by selecting a component from the graphical view

1. Identify a segment from which you want to remove instrumentation.

   You can select either a permanently instrumented segment (colored green), or a temporarily instrumented segment (green with an icon which signifies temporary instrumentation).

2. Right-click the segment.

   The illustration above shows a green segment which the user has right-clicked.

3. Select *Remove Instrumentation...*
4. Click OK.

   The removed component or method does not appear in any subsequent traces.

Remove instrumentation by selecting a row from the View All Called Methods dialog

Note: This method of removing instrumentation is not available to .NET users.

1. Right-click a component, i.e. one of the rows of the table.
2. Select *View All Called Methods...*

   A dialog displays all called methods, indicating those which are already instrumented.

3. Right-click an already instrumented method.
4. Select *View instrumentation Info...*

   A dialog displays the instrumentation for the selected method, including:

   - whether the method is dynamic
   - whether the instrumentation is currently enabled

CA Technologies Prerelease Documentation. Authorized Use Only.

5. Right-click one row.
6. Select *Remove.*

   The removed row disappears from the dialog.
7. Select *Close* to close the dialog.

   The removed component or method does not appear in any subsequent traces.

Remove instrumentation by removing labels from classes and methods

Note: This method of removing instrumentation is not available to .NET users.

1. In the Investigator tree, right-click an agent node.
2. Select *Remove Dynamic Instrumentation.*

   The Remove Dynamic Instrumentation dialog opens. It lists the labels which have been assigned to classes and methods which the selected agent monitors. Each label corresponds to a *.pbd* file which resides in the *Dynamic* directory.
3. Select one or more labels. Use CTRL-click to multi-select.
4. Click OK to permanently remove the instrumentation represented by the selected labels.
5. In the confirmation dialog, click OK.

   The Browse tree will automatically be refreshed to display only metrics which are still instrumented. Subsequent transaction traces will not display any of the classes or methods whose instrumentation you removed.

Note: This method of removing instrumentation works on permanently instrumented classes only.

Exporting Instrumentation

Once you have used the dynamic instrumentation features in Workstation to instrument classes or methods, you can save the instrumentation to a file, then import it to other agents. The result is a *.pbd* file with the same functionality as one created through other means. (See the *CA APM Java Agent Implementation Guide* or *CA APM .NET Agent Implementation Guide* as appropriate for your environment for information about other ways *.pbd* files are created.)

You can only export permanent instrumentation.

Note to .NET users: This functionality is not supported for .NET.

CA Technologies Pre-release Documentation. Authorized Use Only

From the Browse tree:

1. Right-click the agent icon.
2. Select *Export Dynamic Instrumentation.*
3. Select either *All Instrumentation* or *Labeled Items*.
4. If you select *All Instrumentation,* skip to step 7.

   Note: If you choose Tagged Changes, you can search using regex.
5. Select a label.
6. Click OK.
7. In the Save As... dialog, enter a name for the *.pbd* file.
8. Click OK.

You can use the saved *.pbd* file to apply the same instrumentation to other agents. For information about how to use custom *.pbd* files, see the *CA APM Java Agent Implementation Guide* or *CA APM .NET Agent Implementation Guide* as appropriate for your environment.

Important: When exporting dynamic instrumentation, you should be careful about the level of instrumentation you export and subsequently use. The "TraceAllMethods" option should *never* be used in a production environment because of the performance implications; it is intended to be used in preproduction only, for the purpose of initially creating a custom PBD which is then pared down. Good practice is to use the search function to filter and pare down instrumentation before you export it from your sandbox environment to the test environment and certainly before using it in a production environment.

Modifying Instrumentation Level

Tracer groups are sets of instrumented classes. They are defined in *.pbd* files, and their main function is to allow you to turn instrumentation of a tracer group on or off to facilitate performance monitoring and triage. For complete information about how tracer groups are defined and used, see the *Introscope CA APM Java Agent Implementation Guide.*

Starting from the agent tree in the metric browser tab, you can:

- Dynamically enable or disable tracer group instrumentation.
- Reset tracer group instrumentation to original settings.
- Keep your changes permanent.

Note to .NET users: This functionality is not supported for .NET.

CA Technologies Internal Use Only
Author Review Documentation.

To dynamically enable or disable tracer groups:

1. In the Metric Browser tree, right-click an agent node.
2. Click *Change Instrumentation Level.*

   The Change Instrumentation Level dialog displays all tracer groups configured on the selected agent, with their current state.
3. Select one or more (using CRTL-click) tracer groups and:
   - To enable a group which is not currently enabled, click Enable.
   - To disable a group which is currently enabled, click Disable.

   Enabled means instrumentation exists for an individual tracer group.

   The dialog will display an asterisk in the rows you have enabled. However, the tracer group's state is not changed until you click OK.
4. Click OK to activate the changes.

   The agent tree in the metric browser tab updates to reflect the changes.

When you click OK on your changes, it modifies instrumentation on the agent, but your changes are not saved in the *.pbd* file in the *Dynamic* folder.

You can:

- Reset tracer groups to original settings.
- Make your changes permanent.

To reset tracer groups to original settings:

1. In the agent tree in the metric browser tab, right-click an agent node.
2. Click *Change Instrumentation Level* to open the Change Instrumentation Level dialog.
3. Click *Reset All.*

   When you select *Reset All,* you return the state of the Tracer Groups to the current permanent state.
4. Click OK.

To make tracer group instrumentation changes permanent:

1. In the Investigator tree, right-click an agent node.
2. Click *Make Instrumentation Level Permanent...*

   The Confirm Instrumentation Change dialog opens. This dialog summarizes the changes you are making permanent.
3. Click OK.

CA Technologies Prerelease Documentation. Authorized Use Only.

Printing a Transaction Trace window

To print the Transaction Trace window:

1. Select Workstation > Print Window.

   The Page Setup window opens. Defaults are letter size, portrait orientation.

2. Click OK to proceed, or change options then click OK.

   The Print window appears.

3. Select printing options, then click OK.

Note: Printing a page range is not supported (everything prints on one page).

The contents of the entire Transaction Trace window prints, scaled to fit on one page.

Querying Stored Events

Transaction Trace session results are automatically stored in the Transaction Event Database. Transaction events include Transaction Traces and errors, including stalls (if you have installed Introscope Error Detector.) The Transaction Event Database contains Transaction Traces that were automatically sampled by Introscope, as described in Automatic Transaction Trace sampling (see page 204). It also contains the results of Transaction Traces sessions you run yourself.

The Transaction Event database supports these types of queries:

- historical events (basic)—see Querying historical events (see page 230)
- similar events (to selection)
- correlated events (to selection)

Note: Be sure that you run some Transaction Trace sessions before you use the historical query, so that there is data to query.

Query Syntax

The sections below describe how to use the Historical Query facility to query stored errors. The query facility:

- **Is case-insensitive**—for query strings or values for query options.
- **Supports the asterisk (*) wildcard character**—Enter a fragment of a search term followed by the asterisk. (You may not start a search term with the asterisk character). For instance, to look for errors associated with a component whose name includes the string Shopping, use the query string Shopping*.

CA Technologies Product Documentation. Authorized use only.

- **Supports Boolean operators—**Search terms can use Boolean logic, such as "AND", "OR", "NOT". and "()" groupings.
- **Supports exclusion conditions—**Use "+JDBC -CICS" to look for transactions with JDBC but not CICS.
- **Supports query options—**Use the options described in Query options and syntax (see page 231) to limit your query error events that occurred in a particular timeframe, or are associated with particular users, or elements of the hosting environment (as identified by domain, agent, host, or process).

Querying Historical Events

To query historical transaction events:

1. Select Workstation > Query Historical Events.

   The Historical Query Viewer opens.

   The Query field will display, in a drop-down, up to twelve previous searches from this session, or previous sessions by the same Workstation user. This enables you to select one of your saved searches instead of retyping it.

   Tip: By default, the field remembers up to twelve searches; you can designate a different number of searches for the field to remember by editing the introscope.workstation.historical.query.history.limit property in *IntroscopeWorkstation.properties*.

2. In the Query field, enter a combination of:

   - the query option *type:* to include all transaction events that match the specified type.
   - a query string—to search for errors that contain or match a string. If you don't enter a query string, all errors events are returned.
   - query options—to limit your search based on event parameters, as defined in Query options and syntax (see page 231).

   Tip: As you begin typing in the Query field, the searches displayed in the drop-down will be limited to those that match what you have typed.

3. Use the Time Range option to filter your query based on a time range, if appropriate—see Viewing historical data (see page 64) for an explanation of how to use the Time Range option.

   If you don't select a time range, the query uses the default of All and does not apply a filter.

CA Technologies Pre-Release Documentation. Authorized Use Only

4. Click Go.

   Transactions that match the query are displayed in the Historical Query window—the format is similar to the Transaction Trace Viewer. For more information see Using the Transaction Trace Viewer (see page 209).

Note: Only 500 events can be viewed. If more than 500 events match the query, the oldest 500 are shown.

Query options and syntax

Queries use Lucene regular expression syntax to locate and substitute text strings. For information about Lucene syntax, see http://lucene.apache.org/java/docs/queryparsersyntax.html (http://lucene.apache.org/java/docs/queryparsersyntax.html).

Field	Description	Example
agent	Limits the search to events reported by a particular agent.	*agent:ControlledRangeAgent*
domain	Limits search to events related to component(s) in a given domain.	*domain:AcmeWest*
fullAgent	Limits search to events reported by specific agent(s), as specified by its full path: *domain|process|host|agent.*	*fullAgent:AcmeWest| Custom Metric Host| ControlledRange Agent*
host	Limits search to events that occurred on a particular host.	*host:Wmiddle01*
process	Limits search to errors related to component(s) in a given application.	*process:Custom Metric Host*
root	Limits search to events associated with specific component(s), as specified by metric path.	*root:servlets|accountServlet*

CA Technologies Prerelease Documentation. Authorized Use Only.

type	Specifies the type of event to include in query results. **errorsnapshot**—Limits search to error events. **normal**—Returns transaction events captured in user-initiated Transaction Traces. **sampled**—Returns transaction events that were captured as a result of Introscope's default transaction sampling. **whatsinteresting**—Returns "What's Interesting Events", which are generated when Application Overview heuristic values change. For more information see What's Interesting events (see page 141). Results for these types will have the following codes in the Type column: E, T, R and WI, respectively. This set of codes is slightly different from the codes available in the Transaction Trace viewer in Live mode (see Using the Transaction Trace Viewer (see page 209)).	*type:errorsnapshot* *type:normal* *type:sampled* *type:whatsinteresting*
url	Limits search to events associated with the specified transaction URL path prefix. The path prefix is the portion of the URL that follows the hostname. In the following URL: *http://burger1.com/bWar/burgerServlet?ViewItem&category=11776&item=55562630&rd=1* ... the path prefix is: */bWar/burgerServlet*	*url:/bWar/burgerServlet*
urlParams	Limits search to events associated with the specified transaction URL parameters. URL parameters follow a question mark (?) in the URL. In this URL: http://ubuy.com/ws/shoppingServlet?category=734&item=3772&tc=photo the URL parameter portion is: ?category=734&item=3772&tc=photo Note: urlParams cannot start with a wildcard character.	*urlParams:category=734**
user	Limits search to events for transactions associated with the specified Username.	*user:jdoe*
message	Limits search to events associated with the specified message.	...

CA Technologies Pre-release Documentation. Authorized Use Only.

traceDataCreationType	Limits search according to whether data was of one of the following types: 0 = millisecond data (default) 1 = microsecond data 2 = nanosecond data The duration and startTime query keywords should be used only with this keyword.	...
duration	Limits search by event duration. Use with traceDataCreationType keyword. When traceDataCreationType is set to 1 or 2, substitute "durationinnanos" for "duration"	See example for traceDataCreationType
startTime	Limits search by event start time. Use with traceDataCreationType keyword. When traceDataCreationType is set to 1 or 2, substitute "starttimeinnanos" for "startTime"	See example for traceDataCreationType
componentsNotShown	Limits search to events where a given component is not shown	...
durationencoded	...	...
time	Limits search to events before or after a specified time.	...
agent	Limits search to events recorded by a specified agent.	...
traceID	Limits search to events with a specified trace ID.	...
error	...	...
isolatedTo	...	...
app	Limits search to events generated by a specified application.	...
shallow	...	...

CA Technologies Prerelease Documentation. Authorized Use Only.

Using special characters

If the following special characters are part of your query, Lucene syntax allows you to escape them with a backslash (\) character:

```
+ - && || ! ( ) { } [ ] ^ " ~ * ? : \
```

For example, to search for (1+1):2, use the query:

\(1\+1\)\:2

Note: The * (asterisk) and ? (question mark) characters are not supported at the beginning of a query.

Querying for similar events

In Introscope you can query for events that are similar to a selected event. For example, similar events might be events that all contain the same components (Servlet > EJB > SQL) with varying response times. Introscope considers events similar if 60% of the strings within them (component names, SQL tables names, and so forth) overlap.

Note: Even if a transaction type event is selected, both transactions and errors might be returned in the results (errors are only be returned if ErrorDetector is installed).

To query for similar events:

- With a window of query results open, select a table row, then select Trace > Similar Events.

Introscope lists similar events in the Historical Query window.

Querying for correlated events

In Introscope you can query for events that are correlated—those that are part of the same larger transaction. For example, a browser response time event is correlated with a servlet transaction event.

Note: Even if a transaction type event is selected, both transactions and errors might be returned in the results.

To query for correlated events:

- With a window of query results open, select a table row, then select Trace > Correlated Events.

Introscope lists correlated events in the Historical Query window.

Saving and exporting Transaction Trace information

In Introscope:

- You can save Transaction Trace data as an XML file that can be opened later in a Transaction Trace window.
- You can export Transaction Trace data as a text file for review in a text editing program.

CA Technologies Prerelease Documentation. Authorized Use Only.

Saving Transaction Trace data

To save Transaction Trace data to an XML file:

1. In the Transaction Trace Viewer, select the Transaction Traces to save:
   - CTRL + click to select multiple Transaction Traces.
   - Edit > Select All to select all Transaction Traces in the window.
2. Click Save As...
3. You can open the file now, or select a location to save the file into, enter a filename, and click Save.

Opening Saved Transaction Tracer XML Data

You can open and view saved Transaction Trace data in a new Transaction Trace window. These files can be shared through e-mail or stored on a shared network drive to enable users to collaborate on problem analysis.

When opening saved Transaction Trace data:

- you cannot restart the Transaction Trace session being viewed.
- links from Transaction Trace components to their metric paths are unavailable if the metric paths aren't live in the Enterprise Manager to which the Workstation is connected.

To open saved Transaction Trace data in an XML file:

1. Select Workstation > Query Historical Events
2. Select Trace > Open Saved Events (XML).
3. Select the XML file from the browser window, and click Open.

   The data in the XML file appears in a new Historical Query window.

Note: When you view saved historical events in an XML file, correlated events will be displayed, but will not be shown as correlated. To see correlation for historical events in a Transaction Trace, view an active trace (see Querying for correlated events (see page 234)).

Now you can:

- export a Transaction Trace as a text file (see page 236)
- select Transaction Traces within the data and save them as a new XML file.

CA Technologies Pre-release Documentation. Authorized Use Only.

Exporting selected Transaction Trace to a text file

To export selected Transaction Traces to a text file:

1. In the Transaction Trace Viewer, select the Transaction Traces to export:
   - CTRL + click to select multiple Transaction Traces
   - Edit > Select All to select all Transaction Traces in the window.
2. Select Trace > Export.
3. Select a location to save the file, and name the file (default name is *<root component type>_<root component name>.txt*.), and click OK.

CA Technologies Prerelease Documentation. Authorized Use Only.

Sample Transaction Trace XML File

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<TransactionTracerSession EndDate="2005-03-15T17:28:13.953-08:00" Version="0.1"
Duration="32" StartDate="2005-03-15T17:28:13.921-08:00" User="Admin">
    <TransactionTrace Duration="32" Domain="SuperDomain"
EndDate="2005-03-15T17:28:13.953-08:00" AgentName="WebLogic Agent"
Host="rnadimpalli-dt3" StartDate="2005-03-15T17:28:13.921-08:00"
Process="WebLogic">
        <CalledComponent MetricPath="Servlets|ActionServlet"
ComponentName="ActionServlet" Duration="32" ComponentType="Servlets"
RelativeTimestamp="0">
            <CalledComponents>
                <CalledComponent MetricPath="JSP|__register"
ComponentName="__register" Duration="16" ComponentType="JSP"
RelativeTimestamp="16">
                    <CalledComponents>
                        <CalledComponent MetricPath="JSP
TagLib|HtmlTag|doStartTag" ComponentName="doStartTag" Duration="0"
ComponentType="JSP TagLib" RelativeTimestamp="16">
                            <Parameters>
                                <Parameter Value="doStartTag" Name="Method"/>
                            </Parameters>
                        </CalledComponent>
                        <CalledComponent MetricPath="JSP
TagLib|BaseTag|doStartTag" ComponentName="doStartTag" Duration="0"
ComponentType="JSP TagLib" RelativeTimestamp="16">
                            <Parameters>
                                <Parameter Value="doStartTag" Name="Method"/>
                            </Parameters>
                        </CalledComponent>
                        <CalledComponent MetricPath="JSP
TagLib|MessageTag|doStartTag" ComponentName="doStartTag" Duration="0"
ComponentType="JSP TagLib" RelativeTimestamp="16">
                            <Parameters>
                                <Parameter Value="doStartTag" Name="Method"/>
                            </Parameters>
                        </CalledComponent>
                        <CalledComponent MetricPath="JSP
TagLib|MessageTag|doStartTag" ComponentName="doStartTag" Duration="0"
ComponentType="JSP TagLib" RelativeTimestamp="16">
                            <Parameters>
                                <Parameter Value="doStartTag" Name="Method"/>
                            </Parameters>
                        </CalledComponent>
</TransactionTrace>
</TransactionTracerSession>
```

CA Technologies Proprietary/Confidential Information. Authorized Use Only.

Chapter 6: Introscope Reporting

Reporting provides critical information for a variety of functions within an enterprise. For example, reports enable business managers to assess applications' impacts on the business; they enable capacity planners to determine resource consumption; and they give Service Level Agreement administrators an understanding of whether goals are being met.

Introscope includes report templates for creating reports quickly, and enables you to create your own templates with custom graphs and tables.

This section contains the following topics:

Creating Report Templates (see page 239)
Working with report templates (see page 259)
Introscope sample report templates (see page 261)

CA Technologies Pre-release Documentation. Authorized Use Only.

Creating Report Templates

A report template defines which metric data to track, the time range of the reported metric data, and how to present the data in graphical and tabular form. After you save a report template, any user can generate a report at any time.

To create a report template:

1. In the Management Module Editor, select Elements > New Report Template.

   Note: The New Report Template menu item is disabled if you do not have write permission.

   The New Report Template dialog opens.

2. Specify the initial elements for the report.

   a. Type the Name for the new report template.

   b. Select *Force Uniqueness* to verify that the report name is unique.

   If you select this option and you then enter a name that is not unique, Introscope adds a number to the name to make it unique.

   Note: The appended number appears after the report template is created, when you view it in the Management Module Editor. If you don't select Force Uniqueness and an identical report template name exists, Introscope displays an error message and does not create the report.

c. Select a Management Module from the drop-down list box to choose the Management Module that will contain the report.

d. Optional: Instead of selecting an existing Management Module to contain the report, click Choose, then click New Management Module and assign a name to the new Management Module.

e. Click OK.

For more information about creating Management Modules, see Creating and working with Management Modules *(see page 264)*.

The new report template is added to the Management Module Editor, and the settings pane opens.

3. In the settings pane, select the Active check box if you are ready to activate the report template.

   When you generate an Active report template it appears in the list of report templates in the Console, Investigator, and Management Module Editor. See Generating reports from report templates (see page 259).

   Tip: It's a good idea to leave a new report inactivated after you create it, so that you can test-generate the report without having it appear in the list. After you test the report and it is ready for use, click Active to make it available.

4. Click *Open Template Editor* to define report data.

   In the Report Editor you specify the purpose of the report, when and how long it runs, and how the results look.

5. Use the toolbar to add elements to your report.

Now you can:

- Add report elements, such as charts, to the report—see Adding report elements to reports (see page 241).
- Define report properties—see Defining properties in the Report Editor (see page 243).

CA Technologies Proprietary Documentation. Authorized Use Only.

Adding Report Elements to Reports

You can add graphical elements such as charts and graphs, based on metrics or metric groupings, to your report.

To add a graphical report element to a report:

1. If the report template editor is not already open, open it:
   a. With the Management Module Editor open, select the report in the pane on the left.
   b. Click *Open Template Editor.*
2. Right-click the Report listed in the upper left pane, and select *Add.*

   A list of available elements appears.
3. Select one of the element types.

   A new set of tabs appears.

   In the following steps, you configure settings for the report element. To save your work as you go, click Apply at the bottom of the edit window.
4. Configure text settings for your new report element using the Text tab.
   a. Specify the title to appear with the report element.

      By default, Use Metric Grouping Name as Title is selected. If you choose this, the element will take the name of the metric grouping whose data it displays. (You associate the element with a metric grouping in step *5d* below.)

      You can also click *Enter Title* and type a new title to appear with the report element.
   b. Optional: Enter a description for the report element. This will appear in a tooltip with the element.
5. Configure Data Properties for the report element using the Data Properties tab.
   a. Set time range.

      The time range is defined by a Start Time and End Time. The report element will display data bound by these times.

      The Template Default Time Range is set in the default report properties (see step 3 in Defining Properties in the Report Editor (see page 243) to set the default time range). You can choose to accept the default time range, or click Override Template Default Time Range.

CA Technologies Proprietary Documentation
Authorized Use Only

To set the time range, click the calendar icon by the Start Time field.

A calendar dialog appears, with the current date ("Today") circled.

- Use the calendar dialog to set the date, and edit the clock time in the text field after the dialog is closed.
- Repeat to set the End Time.

b. Set the report duration using the Duration field.

Note: If you have specified a Start Time and End Time, leave the Duration field blank.

c. Use the Unit dropdown to match the numbers you entered in the Duration field.

d. Select a metric grouping to associate with the report element.

- Click the drop-down next to the Metric Grouping label.
- A list of available metric groupings appears.
- Select one of the available metric groupings.

e. Optional: Filter the metrics associated with a metric grouping, or define a new metric grouping.

To filter the metrics associated with a metric grouping, click Choose and enter a regular expression.

To create a new metric grouping, click Choose, click New Metric Grouping, and use the dialog to create a new metric grouping based on a management module. For information about defining metric groupings, see the *CA APM Configuration and Administration Guide*.

f. Set values for element attributes in the table of element attributes:

6. Set the display properties for the report element in the Display Properties tab.

For information about display properties, see Defining properties in the Report Editor (see page 243).

7. When you have finished setting all properties for the report element, click Ok.

CA Technologies Prerelease Documentation. Authorized Use Only

Defining properties in the Report Editor

Each element in the report—graphs, tables, bar charts, and pie charts—has properties that you can edit by selecting a properties tab. When you select the Report Element (the top element in the list, which is labeled with the report title) you see tabs that enable you to specify default properties:

- **Cover Page—**these properties apply to the selected element only: a title for the report, a logo to include on the cover page if appropriate, and a description of the report.
- **Default Data Properties—**specify defaults for the whole report: time range of the data (start and end time), the reporting period (for example, 15 seconds or 1 minute), and a specification of the metric data to report.
- **Report Properties—**specify formatting properties that apply to this report only (whether to show the title page and table of contents), and properties that apply to the whole report (time zone and language).
- **Default Display Properties—**define the default appearance of graphs and tables for the whole report.

Note: Changes to the properties in the Default Data and Default Display tabs affect all elements in the report. Individual element customizations will not be affected by the changes in default properties.

To define properties in the Report Editor:

1. Click the Cover Page tab to specify the purpose of the report.
2. Enter the information that will appear on the report's cover page:

To add	Do this
Report Title	Type a title for the generated report; the title appears on the title page with the table of contents.
Logo	Click Choose to browse for your logo or other graphic file. Any graphic chosen here appears in the upper left corner of the title page. Supported formats are .jpg, .gif or .png.
Report Introduction	Type text that describes the contents of the generated report. The introduction appears on the title page above the table of contents.

CA Technologies Prerelease Documentation
Authorized Use Only

3. Click the Default Data Properties tab to specify the default time and data parameters for all elements.
4. You can accept the default data properties, or set new ones:

For	Do this
Start Time and End Time	When you specify a time range, you can specify a specific start date and end date, or specify a time period such as *24 hours*. You can specify a time range for the report in one of these ways: ■ Type a specific start and end date and time, or click the calendar icon to select start and end dates. ■ Leave the Start Time blank and use the Duration and Unit parameters to specify how long the report runs. ■ Leave the End Time blank and use the Duration and Unit parameters to specify how long the report runs. ■ Type *Now* for the End Time and use the Duration and Unit parameters to specify how far back in the immediate history to report on. **Note:** When you type a specific start or end date and time, use the format *mm/dd/yy hh:mm* (or *dd/mm/yy hh:mm*, depending on the machine's regional settings) and then specify AM or PM—for example, you would type 12/15/06 10:00 AM for English Regional.
Duration	Type a number to specify how long the report runs. This number works in conjunction with the Unit value—for example, you might type *24* for the duration if the Unit is *hours*. **Note:** See the explanation of Start Time and End Time for a description of how the Duration and Unit parameters work in conjunction with Start Time and End Time.
Unit	Select a time unit from the drop-down list. Options are minutes, hours, days, or weeks.
Default Period	Click the field to activate the drop-down menu, then select a default reporting interval for the report. You can choose to aggregate all data over the interval, or choose a specific reporting interval—for example, 15 seconds, 15 minutes, a day, or a week. If you choose a specific interval, the data is averaged over the specified interval. The default period value is Auto; this chooses the period automatically, based on the selected Start and End Time range.

CA Technologies Product Documentation
Authorized Use Only

Default Agent Override Expression	Type the default expression to use if you want to override other agent expressions: ■ If you are entering data properties for the report element, and therefore for the *overall* report, all elements in the template use this expression. The value you enter here overrides the metric grouping or Management Module settings. ■ If you are entering data properties for an individual element, the value you enter here overrides values entered for the top-level element, as well as the metric grouping or Management Module settings. This field is optional. If you leave it blank, Introscope reports on the agents based on the metric grouping setting. If the metric grouping is set to inherit the agent expression from the Management Module, Introscope reports on the agents based in the Management Module. **Note:** When you generate a report you can specify an agent expression that overrides the template agent expression. See Generating reports from report templates (see page 259).
Start Time of Reference Data	Enter a date and time if you want to overlay a graph with metric data from the same metric grouping, but from a different time range. When you use an overlay, Introscope identifies the metric data that is plotted on the graph, and overlays it with data from the same metric grouping, but from you specified time range. The length of the period is the same as that of the base metric grouping. To specify a start time for the reference data, you can: ■ Type a date and time, using the format *mm/dd/yy hh:mm* (or *dd/mm/yy hh:mm*, depending on the machine's regional settings) and then specify AM or PM—for example, you would type 12/15/06 10:00 AM for English Regional. ■ Click the calendar icon to select a start date. When you use the calendar to select a start date, Introscope sets the time to the current time—to change the time, type over it.

1. Click the Report Properties tab to specify settings for the report's formatting, time zone, and language.
2. Enter the settings for the report:

To...	Do this
Show title page	Click On to generate a title page for the report.

CA Technologies Proprietary & Confidential Information
Authorized Use Only

Include table of contents	Click On to create a table of contents on the title page.
Add report signature	Type a signature to appear at the bottom of the title page.
Time zone	Click the row to open the list of time zones and choose a time zone. The default is Use Time Zone of Client. The report uses the selected time zone for the Report Date, and Start and End dates.
Language	Click the row to open the list of languages. Choose a language to format the report's date and time according to its standard. For example, the Italian date/time standard is 9-mar-2008 15.50; the Japanese standard is 2008/03/09 15:50. The language settings also determine the font used to display the report in PDF files. To display Asian Language text properly in PDF files, be sure to set the language appropriately. **Note:** In reports set to a non-English language, some English words will still appear where they represent labels, the internationalization of which is not supported. The default is Use Client Locale, which bases the date and time formatting on the language used on the client machine. **Note:** Producing reports in Asian languages requires that some additional components were installed on your Workstation during Introscope Installation. See the *CA APM Installation and Upgrade Guide* discussion "Configuring the Workstation for Asian-Language Reports" for information.

3. Click the Default Display Properties tab.

   You can accept the default properties, or set new ones to determine how the graphs and tables in the report look after the report is generated.

   This tab, like the other Default tabs, enable you to set default property values for all elements in the report. For example, by setting *Row Limit* to 10, you ensure that all tables in the report have a maximum of 10 rows. You can, however, override this value for a particular table element in the report by selecting the element, clicking the Display Properties tab, then entering a new *Row Limit property.*

   Note: Use the scroll bar on the right of the Default Display Properties tab to see all the properties.

CA Technologies Proprietary Document
Authorized Use Only

4. Click the Display Properties tab to set the default display properties.

   In reports, Average Min, Average Max, Mean, Absolute Min and Absolute Max are defined as follows:

   - Average Min—The unweighted average of the minimum values of all periods.
   - Average Max—The unweighted average of the maximum values of all periods.
   - Mean—A weighted average, calculated as follows:

     (tv1 + tv2 + tvn...) / dp

     where *tv* is the total of all values for a period, and *dp* is the total count of data points for all periods. This gives greater weight to periods with more data points.
   - Absolute Max—The actual largest or highest single value across all periods.
   - Absolute Min—The actual smallest or lowest single value across all periods.

The table below contains additional information on display properties and the steps necessary to configure them.

Note: In this step, it is possible to set display property attributes Sort Rows, Sort By, and Value Format only for the report element types Metric Data Table and Bar Chart. These attributes cannot be set for report element types Metric Data Bar Chart and Metric Data Graph.

For	Do this
Aggregate Data by Group	If on, combines data across metrics by summing or averaging all metrics in a group (based on the *Aggregate Using* property) When metrics are grouped, only the group's summary values appear in a report, instead of the individual metric-level values. The aggregated summary rows are presented like metric-level rows in a table or a plot in a chart, but their labels show the group name instead of the individual metric name. The group name becomes the label for the data item, replacing the Item Label regular expression. Use the Group Definition regular expression property to determine the group—see Setting custom group definitions (see page 252).
Aggregate Using	If *Aggregate Data by Group* is on, set this property to Sum or Average, to specify how grouped metrics appear in a report.
Fill Time Markers	If on, the time between the Marker Start and Marker End time is highlighted in the report.
Fill Y Axis Markers	If on, the area between the Y Axis Marker Start and End values is highlighted in the report

CA Technologies Prerelease Documentation. Authorized Use Only.

Group Definition	When either *Aggregate Data by Group* or *Subtotal by Group* is on, use this property to define the group. You can select a group from the drop-down list, or create a custom regular expression. The group options from the menu are: ■ Fully qualified metric name ■ Agent location ■ Agent location - Metric Name ■ Agent name ■ Host ■ Metric Category ■ Metric Category: Metric Name ■ Metric Name ■ Servlet Name Selecting one of these options inserts the appropriate regular expression. To create a group using a custom regular expression, see Time series bar charts (see page 256).
Item Label	Select a label for the item to appear in the legend: ■ Fully qualified metric name ■ Agent location ■ Agent location - Metric Name ■ Agent name ■ Host ■ Metric Category ■ Metric Category: Metric Name ■ Metric Name ■ Servlet Name Selecting an option inserts the appropriate regular expression. You can use variables or regular expressions to create labels. See Setting custom group definitions (see page 252).
List Agents	This setting allows you to choose whether to display a list of the agents whose metrics are being displayed. ■ On (default)—the list of the agents will be displayed ■ Off—the list of the agents will not be displayed

CA Technologies Pre-release Documentation. Authorized Use Only.

Min/Max Bars	Plots the minimum and maximum values in each period for any given metric. You specify how you want the minimum and maximum bars to appear: ■ Show None (shows only the mean value) ■ Show Max Only ■ Show Min Only ■ Show Min and Max
Red Line Value	Specify the Y axis value where a red line is drawn to represent an alert trigger value, with a Red Line Label if you specify one.
Red Line Label	Type a label for the red line.
Row Limit	Specify a value to filter to show only values above or below the limit, depending on whether *Sort Rows* is set to ascending or descending.
Show Average Lines	If On, shows the averages of the metrics in the graph.
Show Fractions of a Second	If On, shows the fractional parts of a second, up to six decimal places to the right. For example: ■ *03:22.5123456* for 3 minutes, 22 seconds and 123456 ms. ■ *00:00.25* for 250 ms. ■ *3.13s* for 3130 ms.
Show Legend	If On, a legend is included for the selected graph. The legend shows which metrics correspond to each plot in the graph according to the color of the plot and, if Show Shapes is on, according to the shape used to mark each data point.
Show Shapes	If On, Introscope draws shapes at each point, in addition to plotting the line between points. For graphs with many metrics or with a high density of data, showing the shapes might obscure the data, but if you omit the shapes the only way to correlate plots with the legend entries is by using color. If a plot consists of only one data point in the given time range, it does not appear in the graph unless shapes are On. In particular if you set the period to *Aggregate All* it plots a single value in the chart, but if shapes are off nothing appears. You need at least two data points for a line to be plotted.
Show Volume	If On, the number of metric data points within each period is plotted as a bar in the report. If more than one metric appears on the chart, the volume bars overlay each other.

CA Technologies Proprietary Document. Authorized Use Only.

Sort By	Select how to sort the columns: ■ Metric/Group Label ■ Mean ■ Average Min ■ Average Max ■ Absolute Min ■ Absolute Max ■ Count ■ Sum
Sort Rows	Select Ascending or Descending to sort the rows.
Subtotal Data by Group	In tables, you can set the *Subtotal Data by Group* to sort the items by group and then subtotal them—when Aggregate Data by Group is on, the Subtotal Data by Group attribute has no effect. Use *Group Definition* to define how metrics are divided into groups, to provide a label for the group. **Note:** Data in tables is always summarized across the entire time range. The Value column is labeled Sum or Mean, depending on the Aggregate Using setting. Choosing Sum adds up every metric value.
Summary Row Label	Type text to appear as the label for the summary row.
Table Columns	Select a value to specify which columns appear in the report: ■ Show All Columns includes Mean (or Sum, depending on how the Aggregate Using property is set), Average Min, Average Max, Absolute Min, Absolute Max, and Count ■ Show Mean, Min, Max, Count ■ Show Mean, Count ■ Show Text Value of Metric Only results in a single column labeled *Value*, which shows the metric unformatted. This is most often used for String metrics that would otherwise appear as zero. **Note:** For a text string value to be reported, the time range for data must be a Live Range of the last 8 minutes.

CA Technologies Internal/Restricted Use Only. Documentation. Authorized Use Only.

Value Format	Select a value to use for the table value display format, and for the Y axis format (except for pie charts): ■ General ■ Use M(illions) and B(illions) ■ Memory Value Format (MB, GB, KB) ■ Percent (%) ■ Percent x 100 (%) ■ Show two decimal places ■ Millisecond as HH:MM:SS (shows milliseconds in hours, minutes, and seconds) use for metrics whose values are milliseconds ■ Microsecond as HH:MM:SS (shows microseconds in hours, minutes, and seconds) use for metrics whose values are microseconds ■ Millisecond as d, h, m, s (shows milliseconds in days, hours, minutes, and seconds—for example, *3h 22m 36s*)
X Axis Label	Type a label to appear along the X axis of the graph.
X Axis Marker Start Time, X Axis Start Marker Label, X Axis Marker End Time, and X Axis End Marker Label	You can use these attributes to bracket a period within a report chart, and to label the start and end points for that occurrence. Start/end date/time formats are expressed, for example, as: `3/31/99 11:30 AM` You can also use the calendar widget which appears when you put your cursor in the Start Time or End Time field. Labels are text strings. The specified period will appear bounded by vertical lines in the report chart, with labels.
X Axis Marker Start Offset in Seconds and X Axis Marker End Offset in Seconds	These settings provide an alternative to setting absolute date values for the start and end markers. The values are an offset, in seconds, from the start of the graph to where the marker appears. Offsets are useful when a report's date range is relative to the report's start and end date and are not absolute time ranges—from Now - 1 hour to Now, for example. For an X Axis marker to appear, you must set either the date or the offset. If both are set, the date is used; if neither is set, no marker appears.
X Axis Time Format	Click the row to choose from a list of time formats

CA Technologies Proprietary Documentation. Authorized Use Only.

Y Axis Format	Click the row to choose from a list of formats—for example, Memory Value Format (MB, GB, KB), or Percentage (%).
Y Axis Label	Type a label to appear along the Y axis of the graph.
Y Axis Marker Start Value, Y Axis Marker Start Label, Y Axis Marker End Value, and Y Axis Marker End Label	Use the Marker Start and End Values to bracket values on the Y Axis, and label those values.
Y Axis Upperbound and Y Axis Lowerbound	Type values in these fields to specify values on which to report. You would use these properties if, for example, you have a metric that might fall far outside the range of values—say, 50 seconds as opposed to 1 second. If you specified the Upperbound property in this situation as 0.8 and the Lowerbound property as 0.2, the report would only report between those values.
Yellow Line Value	Specify the Y axis value where a yellow line is drawn to represent an alert trigger value, with a Yellow Line Label if you specify one.
Yellow Line Label	Type a label for the yellow line—for example, *Response time is slow*.

Setting custom group definitions

You can use group definition to define grouping for these elements:

Bar Charts	Bar Charts are a simple way to show summary data. The values in a bar chart are the same as you would see in a table, but you can additionally use *Group Definition* to group the bars. You use the *Group Definition* property to group bars in the chart and define the label that appears underneath each group of bars. By default it is the agent. To disable grouping, enter a literal value for the group definition and that will appear as a single label underneath the chart. Use the Item Label property to define what appears in the legend.
Pie Charts	Pie Charts are useful for showing relative values of summary or grouped data, defined by the Group Definition property to divide metrics into groups. Set the Aggregate Data By Group property to *on.* Use the Item Label property to define what appears in the legend.

CA Technologies Proprietary Document. Authorized Use Only.

Aggregating Data	When you use the *Aggregate Data into Groups* property, Introscope combines the metrics in a group by summing or averaging, depending on how the *Aggregate Using* property is set. The aggregated data becomes a new data item and appears as a single row in a table, or a plot in a chart. The group name becomes the label for the data item, and the *Item Label* property no longer applies.
Subtotalling	You use *Group Definition* to define how metrics are divided into groups, to provide a label for the group, and to subtotal rows. The *Subtotal Data by Group* property is similar to aggregation. In tables, both properties combine rows, but in subtotalling the individual metric rows appear; with *Aggregate Data by Group* turned on, only the subtotal rows appear. In tables, you can set the *Subtotal Data by Group* to sort the items by group and then subtotal them—when Aggregate Data by Group is on, the Subtotal Data by Group attribute has no effect. **Note:** Data in tables is always summarized across the entire time range. The Value column is labeled Sum or Mean, depending on the Aggregate Using setting. Choosing Sum adds up every metric value for every data point in the entire time range.

You can use variables and regular expressions to:

- extract a common part of a metric string and thus define a group.
- format data item labels.

Using variables

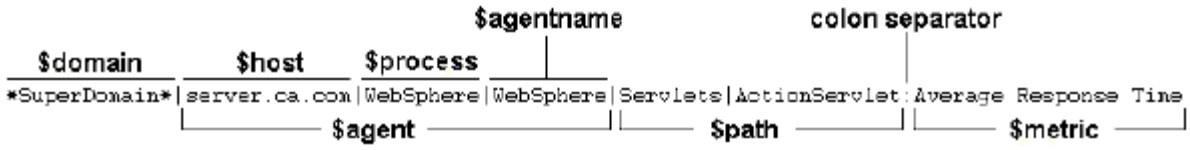
Use these variables to extract parts of the fully qualified metric string.

Variable	Substitution
$host	host part of an agent
$proc	process part of an agent
$agentname	agentname part of an agent. Compare with *$agent*.
$agent	full agent spec: host, process, agent
$metric	The part of the metric identifier to the right of the colon (:).
$path	The part of the metric identifier to the left of the colon (:).
$path[n]	The indexed segment of the path (base 1). If out of range, return empty string
$path[-m]	Path segment *m* counting from the end.

CA Technologies Prerelease Documentation
Authorized Use Only

$path[m:n]	The part of the path from segment *m* up to and including segment *n*. If either value is negative, then the segment is counted from the end.
$domain	Domain; for example **SuperDomain**
$regex	Defines the beginning of a regular expression string. See Using regular expressions (see page 254).

For example:



Using the example above:

This string using variables and plain text...	**is displayed:**
$host - $path[-1]	damien.ca.com - ActionServlet
$agentname servlet *$path[-1]*	WebSphere Servlet ActionServlet
Servlet *$metric*	Servlet Average Response Time

Using regular expressions

You can also use regular expressions to define grouping.

Regular expressions use these patterns:

Variable	Description
$regex['pattern']	The part of the full metric URL which matches the given regular expression. If the regex has a group, then only extract that group. Otherwise extract whatever is matched. If nothing matched, return the full metric. This will be needed to represent old settings.

CA Technologies Prerelease Documentation. Authorized Use Only.

$regex['pattern','replacement']	Replace the part of the full metric URL which matches the given regular expression pattern with the given substitution pattern. Any capturing groups in pattern can be inserted into replacement using $ variables, where *$1* is the first group, *$2* is the second, etc.

For the full qualified metric

```
*SuperDomain*|foo.company.com|WebSphere|WebSphere|Servlets|ActionServlet:Aver
age Response Time
```

This string using variables and plain text...	**will display as:**
$regex['(\w).company.com']* servlets	foo servlets

Using regular expressions to match a range of metrics

Consider an example where this regular expression is used as the item name:
`\|Servlets\|.*:Average Response Time.*`

Let's say that this matches five different servlets on each of two agents. If you show these metrics on a chart with default settings you will see 5 * 2 = 10 plots on the chart.

You can group the metrics by Servlet or by agent. The default is by agent, because the default group definition is:
(.?\|.*?\|.*?)\|*

If you set Aggregate Data by Group to *on*, you will see only two plots—one for each application server that is the aggregation of all servlets on that application server.

If you change the group definition to be a regular expression matching the servlet name, the metrics for a particular servlet on both application servers will be aggregated into a single plot, giving you 5 plots, one for each servlet.

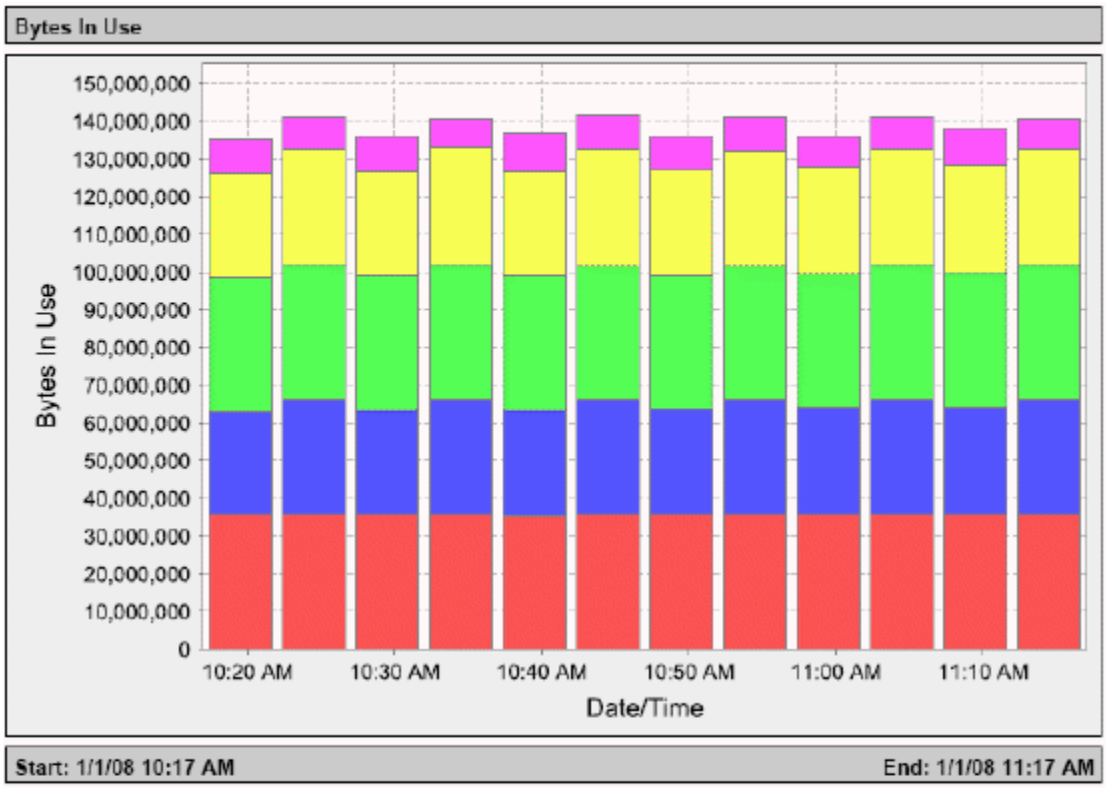
In this case the group definition might be: *Servlets\|(.*):* to match the exact Servlet name part of the metric.

A complete guide to the supported regular expression syntax is located at Sun's Java API Pattern class page—http://java.sun.com/j2se/1.5.0/docs/api/java/util/regex/Pattern.html.

CA Technologies Internal/Partner Documentation. Authorized Use Only.

Time Series Bar Charts

You can configure report bar charts to show time series data.



Period: 5 minutes

Agents: myServer1|App Servers|wilyFONE_01
myServer1|Web Services|websvc-1
myServer1|.NET|FoneBook
myServer1|Web Services|websvc-2
myServer1|App Servers|wilyFONE_02

Legend:

- myServer1|App Servers|wilyFONE_02|GC Heap:Bytes In Use
- myServer1|Web Services|websvc-2|GC Heap:Bytes In Use
- myServer1|App Servers|wilyFONE_01|GC Heap:Bytes In Use
- myServer1|Web Services|websvc-1|GC Heap:Bytes In Use
- myServer1|.NET|FoneBook|GC Heap:Bytes In Use

To specify a custom time range for the report:

1. Select the report template.
   a. If you want to create a new report template, see Creating report templates (see page 239).
   b. If the report template already exists, open the Management Module Editor, expand the tree structure, and select the report template to edit.
2. Verify that the *Active* check box is selected.
3. Click *Open Template Editor.*

   The Edit Report dialog opens. For more information about the tabs in this dialog and what you can do with it, see Defining properties in the Report Editor (see page 243).
4. Click the Default Data Properties tab.
5. With the Report selected in the left pane:
   a. Right-click the title of the report template in the left pane, and click *Add.*
   b. Select *Metric Data Bar Chart* from the available elements.

   The Metric Data Bar Chart element you added will appear in the list of report elements under the Report.
6. Click the Data Properties tab to define properties for the chart.
7. Set the time range:
   a. Select *Override default time range.*
   b. Enter start and end date and time values.
   c. Verify that the Duration and Unit settings agree with Start and End Time values. These do not automatically reset based on the Start and End Times.

   Note: Setting the time range to a relatively small period will cause graphic display elements in the chart to overlap and reduce readability.

CA Technologies Prerelease Documentation. Unauthorized Use Only.

8. Select a metric grouping to associate with the report element.
    a. Click the drop-down next to the Metric Grouping label.

        A list of available metric groupings appears.

    b. Select one of the available metric groupings.

    c. Optional: Filter the metrics associated with a metric grouping, or define a new metric grouping.

        To filter the metrics associated with a metric grouping, click Choose and enter a regular expression.

        To create a new metric grouping, click *Choose,* click *New Metric Grouping,* and use the dialog to create a new metric grouping based on a management module. For information about defining metric groupings, see the *APM Configuration and Administration Guide*.

9. Set the Period for the chart. This sets the interval on the Y axis. For example, setting the Period to 5 minutes will produce a series of bars each representing 5 minutes, as in the example on page 161.
10. Click the Display Properties tab to set display properties.
11. Set the Item Label:
    a. Click in the pane to the right of Item Label.

        The pane becomes a drop-down.

    b. From the drop-down, choose *Fully Qualified Metric Name.*

        After selection, Fully Qualified Metric Name is displayed as: *$agent|$path:$metric*.

12. Set *Aggregate Data by Group* to On.
13. Set Group Definition:
    a. Click in the pane to the right of Group Definition.

        The pane becomes a drop-down.

    b. Click *Fully Qualified Metric Name.* After selection, Fully Qualified Metric Name will be displayed as *$agent|$path:$metric*.

14. Apply your changes, then click Ok.

CA Technologies Prerelease Documentation
Authorized Use Only.

Applying and reverting changes

To apply changes to a report:

- Click the Apply button.

  The Apply button saves your changes to a report without closing the report, allowing you to continue working.

To save changes and close the report:

- Click the Ok button.

To revert changes to a report:

- Click the Revert button.

  The Revert button returns the report to:

- the state it was in after you last clicked Apply, or
- if you haven't clicked Apply, to the state it was in when you opened it.

Working with report templates

This section contains:

- Copying or deleting report templates (see page 259)
- Generating reports from report templates (see page 259)

Copying or deleting report templates

To copy or delete report templates:

1. Right-click the template.
2. Select Copy Report Template <*name*>, or Delete Report Template <*name*>.

Generating reports from report templates

To generate reports, the report template must be active and the Enterprise Manager must be running. Introscope produces reports in these formats:

- PDF
- HTML
- XML (no pages)
- XML (Embedded Pages)

CA Technologies Prerelease Documentation. Authorized Use Only

- Multi-Sheet Excel (*.xls)
- Single Sheet Excel (*.xls)
- Word (*.rtf)
- Comma Separated (*.csv)
- Text (*.txt)
- Jasper Report (*.jrprint)

Note: Any user with read permission can generate a report from a report template.

To generate a report from a report template:

1. Select an active report template in one of these ways:
   - In the Management Module Editor, right-click on a report template and select Generate Report from Report Template <*name*> from the menu.
   - In the Management Module, Investigator, or Console, select Workstation > Generate Report.

   The Choose Report Template dialog box opens.
2. Select a report template from the list and click Choose to open the Generate Report dialog box.
3. Specify the report's start and end dates.

   Note: Time ranges for the report are calculated according to the time zone of the Workstation generating the report. The day starts and ends at midnight.
4. If you want to override the template agent expression, specify a different agent expression or click Select to choose an expression.
5. Click Generate Preview.

   The Preview pane shows the report results.
6. Now you can use the Preview buttons to manipulate the report output:
   - Click Save to open the Save dialog box. Specify a location and file name, and choose a format in which to save the report.
   - Click Print to open the Print dialog box and specify a printer.
   - Click the navigation arrows to move forward and backward through the report, or type a page number in the page number field.
   - Click the page views to choose how the report appears.
   - Click zoom to choose the view magnification.

CA Technologies Prerelease Documentation. Authorized Use Only.

Introscope sample report templates

Introscope includes sample report templates that are based on the sample dashboards and Management Module that are included with Introscope. You can customize these sample report templates and edit them to match corresponding business needs.

Application Capacity Planning report

The Application Capacity Planning report includes the graphs listed in this table of contents. The report shows trends in J2EE Application server resource utilization over a period of time. The default is one day, for a three-month period.

Production Application Health

The Production Application Health report includes the graphs listed in this table of contents. The report shows overall application health. It reports on the performance of EJBs, JSPs, servlets, SQL statements, available JDBC connections, and idle threads over the last 7 days.

QA/Test Application Performance

The QA/Test Application Performance report includes the graphs listed in this table of contents. The report shows all the characteristics of the application from a performance point of view in a QA or test environment. These include a component performance view as well as resources view.

CA Technologies Prerelease Documentation.
Authorized Use Only

Chapter 7: Creating and Using Management Modules

Management Modules are collections of objects and settings which together enable the graphical presentation of metric information in Console dashboards, as well as various actions for Introscope to take under circumstances you configure. For these purposes, they contain various elements, such as alerts, actions and dashboards.

This chapter describes how to create or edit Management Modules using the Management Module Editor to configure presentation, detection, and notification logic for managing Introscope-enabled applications.

This section contains the following topics:

About Management Modules (see page 263)
Creating and working with Management Modules (see page 264)
Configure Metric Groupings (see page 273)
Create and Edit Dashboards (see page 278)
Creating dashboard text and graphics (see page 295)
Creating and managing custom hyperlinks (see page 302)
Monitoring performance with alerts (see page 306)
Using Calculators (see page 336)
Using JavaScript calculators (see page 338)
Deploying Management Modules (see page 343)

CA Technologies Proprietary Documentation
Authorized Use Only

About Management Modules

A Management Module contains a set of Introscope monitoring configuration information. You configure Introscope's monitoring logic by using Management Modules and elements, which organize metric data in the Workstation.

Management Modules for each domain contain elements. Elements are objects that contain and organize metric data with monitoring logic, for presentation in the Workstation. Introscope elements are:

- Metric groupings
- Alerts (which includes Simple and Summary Alerts)
- Actions (which includes Shell Command, E-mail, and Workstation actions)
- Dashboards
- Calculators

- Report templates
- SNMP collections

A Sample Management Module is included in the SuperDomain when you install Introscope. This Sample Management Module contains pre-configured dashboards that include commonly used performance-monitoring logic. You must create other Management Modules for other created domains.

The section Pre-configured Introscope Dashboards (see page 50) describes the contents of the sample dashboards.

Triage Map Configuration Management Module

Another useful Management Module is the Triage Map Configuration Management Module. When you add a new Management Module element, this Management Module is selected by default in the dropdown list of available Management Modules. This Management Module does not contain any preconfigured content, but it does save definitions of alerts and other objects you create directly from the application triage map.

Permissions, Domain Enforcement and Element Editing

Agents are partitioned into domains. Users are given access to certain domains, and can only create elements and Management Modules that reference data in domains to which the users belong.

To create or edit elements, you must have the appropriate permissions. To perform most changes to elements, you need write permission to the domain in which the element is contained. Some functions require a specific permission.

Keep in mind that when creating or modifying an element that elements in individual domains can only reference other elements in the same domain. Elements in the SuperDomain can reference elements in any domain.

For more information about domains and user permissions, see the *CA APM Configuration and Administration Guide.*

CA Technologies Prerelease Documentation. Authorized Use Only

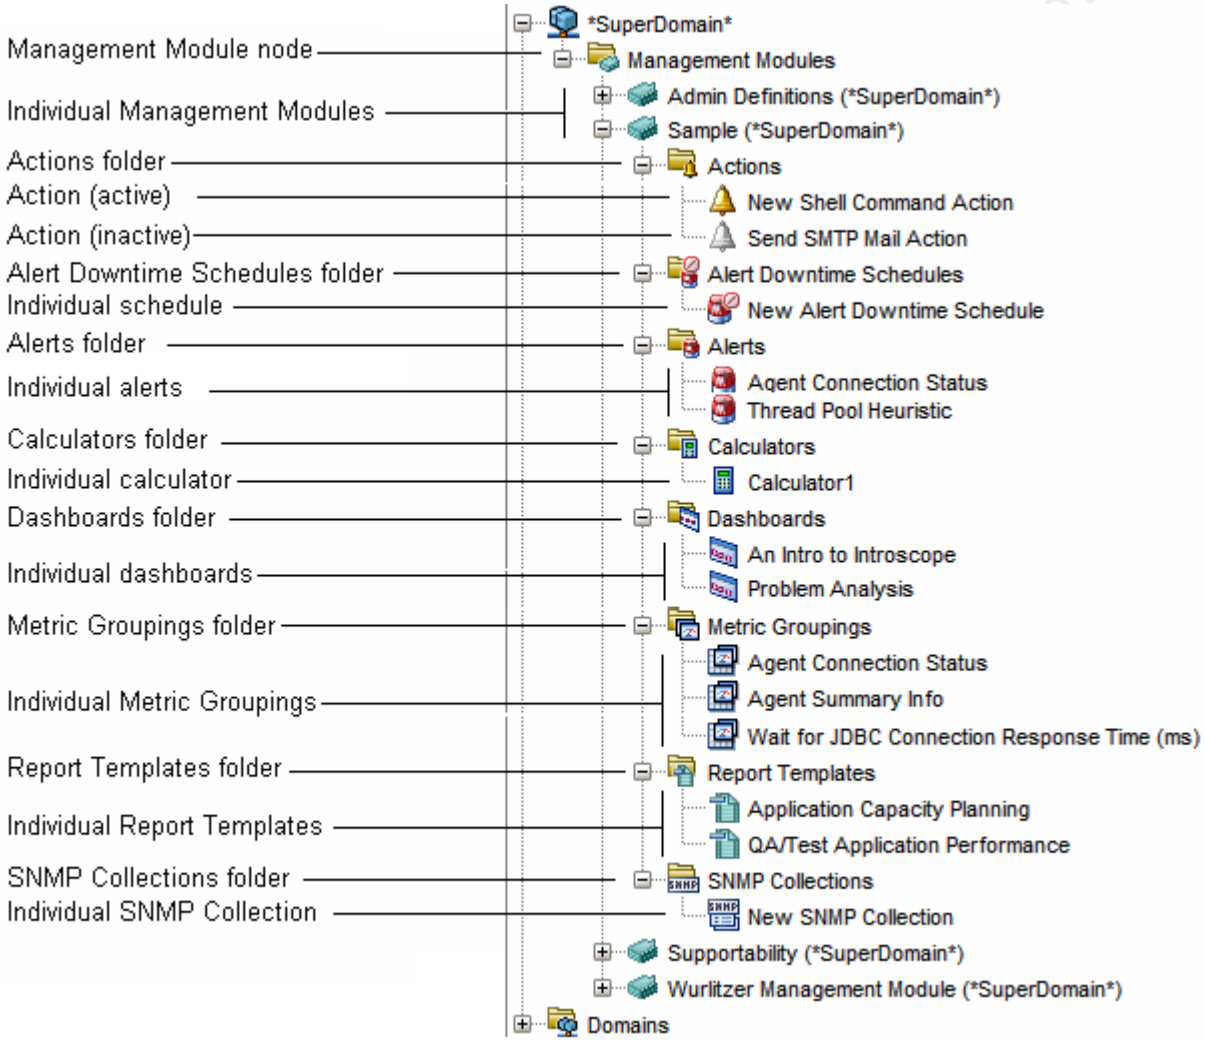
Creating and working with Management Modules

Management Modules organize dashboards and other elements so that you can conveniently find, copy, and edit them. Management Modules are stored as *.jar* files in your *<EM_home>/config/modules* directory. They can also exist in subdirectory domains beneath the *<EM_home>/config/modules* directory, if a user defines those domains and the Management Module .jar beneath them. Management Modules can be moved between Enterprise Managers—see the *APM Installation and Upgrade Guide*.

A Management Module can be defined as editable/non-editable, or active/inactive. If a Module is not editable, the elements within it are also not editable. If a Management Module is inactive, the elements within it are also inactive.

Elements in the Management Module Editor

The illustration below shows the parts of the Management Module Editor tree:



This table describes Management Module elements:

Element	Description
Management Module	A container that holds elements.
Actions	Actions specify actions to be triggered by alerts.

Alerts	Notifications of possible problems in your application, generated by comparing metric values against user-defined threshold values and producing a status.
Calculators	A calculator sums or averages metric data to produce custom metrics.
Dashboards	Objects which contain data viewers, shapes, text, and images.
Metric Groupings	Objects which specify which metrics to act upon; used as building blocks for elements such as alerts.
Report Templates	Customized report templates.
SNMP Collections	Objects which define which collected metrics are to be included in a published MIB.

Searching for Management Module elements

You can search for any Managment Module element using Lucene syntax regular expressions.

To search for Managment Module elements:

1. In the Management Module editor, select a domain or a Management Module node.
2. Click the Search tab.
3. Enter a regular expression, using Lucene syntax, in the Filter pane. See Query options and syntax (see page 231) for information on how regular expression search works.

Note:

- Special characters must be escaped. See Using special characters (see page 233).
- Beginning your search string with the asterisk (*) or question mark (?) character will evoke an error. These characters are not permitted at the beginning of a Lucene expression.

As you type in the Filter pane, matches appear with each keystroke. Matches appear in table format in the Search tab, displaying the following information for each element matching the search:

- Element name
- Management Module to which the element belongs
- Domain to which the Management Module belongs

CA Technologies Prerelease Documentation
Authorized Use Only.

To jump from search results to the element in the Management Module Editor tree:

- With results of your search visible in the viewer pane, double-click one of the results.

The Management Module Editor tree opens, with the selected element highlighted. The Search tab is replaced by the configuration tab for the selected element.

To return to the search results:

- Click the Back arrow to return to the Search tab.

The viewer pane displays the most recent search query string and results.

Copying search results

The search results table supports copy-paste, so you can use Management Module editor names and other parts of the results table in other applications.

To copy results from the Management Module search:

1. Use the mouse to highlight all or part of a row in the search results table.
2. Use a keyboard command to copy, or right-click the highlighted text and choose Copy.

The highlighted text will go on the clipboard, and you can paste it in a text editor or other application.

Using hyperlinks in the Management Module Editor

The Workstation Console and Management Module Editor provide hyperlinks between related items. For example, a metric in the Management Module Editor tree contains a link to any dashboard it is used in.

Links for an item in a dashboard or the Management Module Editor appear on the Links menu, with link types separated by a horizontal line. In the Management Module Editor, links on top of the menu are for Management Module Editor tree items; links on the bottom of the menu are for dashboards. In the Console, links on bottom of the menu are for Management Module Editor tree items; links on the top are for dashboards.

Hyperlinks can be viewed in two ways:

- right-click on the item in the tree, select the Links submenu
- select an item in the tree, select Properties > Links

If no links are available for a selected object, the Links menu is disabled.

CA Technologies Preselease Documentation. Authorized Use Only.

Agent expressions for Management Modules

In Introscope, metric groupings (and their agent and metric expressions) filter data that matches the agent and metric criteria. All metric groupings in a Management Module can share a single set of Agent Expressions. You can then specify at the metric grouping level whether to use the shared Agent Expression, or the metric grouping's Agent Expressions.

Using Management Module Agent Expressions simplifies monitoring logic configuration. You can make a change to the Management Module agent and metric expressions, and have it apply to all metric groupings contained within. If something in your deployment changes (the machine name, for example), it is easy to make the change at the Management Module level, and it applies to everything within. Or, you can copy a precisely configured Management Module and change the Agent Expression to monitor a different agent.

For simplicity, CA Technologies suggests that you either use Management Module Agent Expressions, or metric grouping Agent Expressions, but not a mix of both within a single Management Module. You might also use only metric grouping Agent Expressions if you want to monitor a specific set of metrics from a specific set of agents.

For steps to specify Agent Expressions in metric groupings, see Configuring metric groupings (see page 273).

Naming Management Modules and elements

These are the rules for naming Management Modules and elements:

- Management Modules within the same domain must have unique names.

  Non-unique Management Module names are allowed in separate Domains.

- Same types of Management Module elements within a single Management Module must have unique names.

  For example, you could have one alert and one calculator, both named "Bytes In Use," but you couldn't have two alerts both called "Bytes In Use."

- Non-unique Management Module element names can exist, if they are in separate Management Modules.

  For example, you could have two Alerts, both named Servlet Alert A, with one in the Sample Management Module and one in a Module you created called Test Module.

CA Technologies Prerelease Documentation
Authorized Use Only

To make naming easier, Introscope provides a Force Uniqueness option for creating and naming a Management Module or element:

- If the Force Uniqueness option is on and you enter a name that is not unique, Introscope adds a number to the name to make it unique.

  The appended number appears after the report template is created, when you view it in the Management Module Editor.

- If the Force Uniqueness option is off and an identical report template name exists, Introscope displays an error message and does not create the report.

Administering Management Modules

This section has instructions for creating, copying and deleting Management Modules, as well as making them active/inactive or editable/non-editable.

Creating a new Management Module

To create a new Management Module:

1. From any Workstation window, select Workstation > New Management Module Editor.
2. In the Management Module Editor window, select Elements > New Management Module.
3. Enter a name for the Management Module in the Name field (this name appears in the Management Module Editor tree.)
4. Enter a .jar file name for the Management Module, using alphanumeric characters without spaces (to comply with all operating systems).
5. In the Domain Name field, use the pull-down menu to select which domain contains the Management Module.
6. Click OK.

   The Management Module appears in the Management Module Editor tree. Modules are active and editable when they are created.

Copying a Management Module

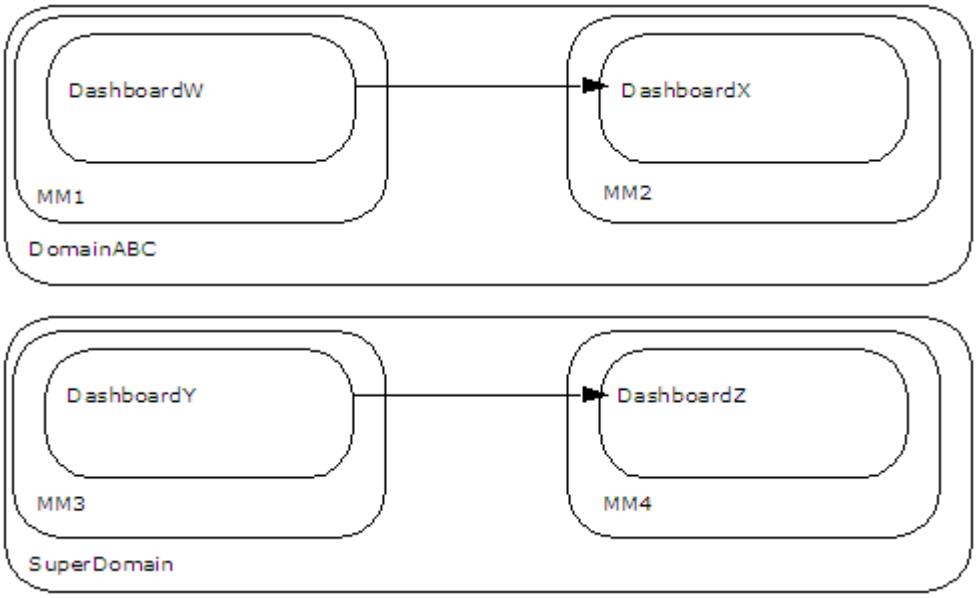
You can copy a Management Module within the same domain, or, within limits, to other domains. Copying a Management Module also copies all the elements in it.

Copying a management module observes these rules:

- When a Management Module contains no dependencies on other Management Modules, you can copy it from any domain to any domain.
- When a Managment Module contains a dependency, that Management Module can be copied only within its domain or to the SuperDomain.

CA Technologies Prerelease Documentation. Authorized Use Only.

For example, in the following arrangement, DashboardW in Management Module MM1 has a dependency on an element in DashboardX in MM2; both are members of DomainABC. Furthermore, DashboardY in MM3 has a dependency on an element in DashboardZ in MM4; both are members of SuperDomain.



In the example, the following are true:

- Management Module MM1, which has a dependency, may be copied to SuperDomain, because SuperDomain always includes the scope of other domains.
- Management Modules MM2 and MM4, neither of which has a dependency, may be copied within their domains, or to any other domain.
- Management Module MM3 may be copied within SuperDomain but may not be copied to DomainABC because of its dependency on MM4, because other domains do not include the scope of SuperDomain.

If you copy a Management Module from one Enterprise Manager to another, the copy is independent of the original Management Module—subsequent edits you make to the original Management Module are not replicated in the copy.

CA Technologies Internal / Authorized Use

To copy a Management Module:

1. In the Management Module Editor, right-click on a Management Module and select Copy Management Module *<Name_of_Management_Module>.*
2. Enter a name for the copied Management Module in the Name field.
3. Enter a .jar file name for the Management Module, using alphanumeric characters without spaces, to comply with all operating systems.
4. The Domain Name field shows the Management Module will be copied to the *Super Domain*.
5. Click OK.

   The new Management Module appears in the Management Module Editor tree. It is active and editable.

Deleting a Management Module

Deleting a Management Module deletes all the elements in it.

To delete a Management Module:

1. In the Management Module Editor, right-click on a Management Module and select Delete Management Module <Name_of_Management_Module>.
2. Click Yes.

Note: Before deleting a Management Module, deactivate it. See Making a Management Module active or inactive (see page 271).

Making a Management Module active or inactive

If the Management Module is made inactive, everything it is made inactive.

To make a Management Module active or inactive:

1. Select the Management Module in the Management Module Editor tree.
2. In the settings pane for the Management Module, check or uncheck the Active box.
3. Click Apply.

Making a Management Module editable or non-editable

If a Module is not editable, the elements it contains are also not editable. A non-editable Management Module is identified by a small padlock on its icon in the Management Module Editor tree.

Custom Management Modules you create are editable, but you can make them non-editable to prevent others from changing them.

CA Technologies Prerelease Documentation Authorized Use Only.

WARNING: Once a Management Module is made non-editable, you cannot return it to its editable state. For instructions on how to make a custom Management Module non-editable, see Making a Management Module editable or non-editable (see page 271).

Defining agent expressions for a Management Module

To define agent expressions for a Management Module:

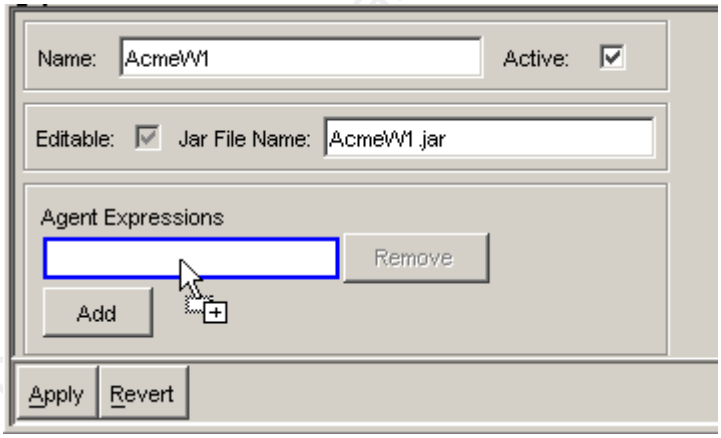
1. Select the Management Module in the Management Module Editor tree.

   The Management Module's settings appear in the settings pane.

2. Click Add.

   A blank Agent Expressions field appears. You can supply Agent Expressions information, in one of two ways:

   - Type in the information in a regular expression.
   - Open another Investigator window, select an agent or metric and drag the information to the Agent Expressions field, so that a the line appears around the Agent Expressions field. For example:



3. Click Apply.

Note: The Agent Expressions defined here are not automatically applied to metric groupings within. You must specifically choose to use the Management Module's Agent Expressions instead of the metric grouping's Agent Expressions. For information on this process, see Configuring metric groupings (see page 273).

Configure Metric Groupings

Metric groupings are Management Module objects that save the following information:

- The *agent expression* -- a regular expression in Perl 5 that filters input to the metric by specifying the data up to and including the agent name.
- The *metric expression* -- a regular expression in Perl 5 that specifies the Resource (the chain of folders leading to the metric) and the metric.
- The Management Module to which the metric grouping belongs.

Metric grouping example

Look at this example from the Supportability Management Module. The **disk usage (mb)** metric grouping uses these:

Metric Grouping Agent Expression:

```
(.*)\|Custom Metric Process \(Virtual\)\|(.*)
```

Metric Expression:

```
Enterprise Manager\|Data Store\|(.*)Disk Usage \(mb\)
```

Specifying expressions for metric groupings

To populate these fields, you can either type in the information using Perl 5 regular expressions language, or you can select and drag metrics and agents from the Investigator into the fields.

Agent Expressions can be defined per Management Module. These Agent Expressions can then be applied to metric groupings within a Management Module.

By default, every metric grouping uses its own Agent Expressions to match agents. If you want instead to use the Agent Expressions from the Management Module, choose this option in the metric grouping's settings panel. If you select this option, the matching agents automatically change if the Management Module's Agent Expressions change.

Metric name structure

A fully qualified metric name looks like this:

Domain|*Hostname*|*Process*|*AgentName*|*Resource:Metric*

For example, a fully qualified metric name of a metric in a Resource looks like this:

Acme|*c1737019-a*|*AcmeUSA*|*AcmeWest*|*GC Heap:Bytes In Use*

CA Technologies Pre-release Documentation. Authorized Use Only.

If a metric is located within two Resources, the name looks like this:

Acme|c1737019-a|AcmeUSA|AcmeWest|Servlets|FileServlet:Responses Per Second

If there are deeper resource layers, resources are separated by the pipe character (|).

See Using variables (see page 253) for more information on how metric names are constructed.

Note: Users in Domains other than the SuperDomain see the metric name without domain information in the following syntax:
Hostname|Process|AgentName|Resource:Metric. For example:
`c1737019|AcmeUSA|AcmeWest|GC Heap:Bytes In Use`

Creating a new metric grouping

You can:

- Create a new metric grouping from an existing metric (see page 274).
- Create a new metric grouping from the Elements menu (see page 275).
- Adding another metric to a metric grouping (see page 276).
- Customizing metric groupings (see page 276).

Note: Newly-created Management Module elements may not correctly display data if you try to view them using a historical time range immediately upon creating and saving them. You may need to wait several minutes to view the correct historical data.

A note about using Japanese characters in metric names

Because data that is captured by the agent is not localized, users on a Japanese locale machine will not see metric names displayed in the browse tree using Japanese characters, if the metrics have been created using Japanese characters. To avoid this problem, create metrics using Roman characters.

Create a new metric grouping from an existing metric

To create a new metric grouping from an existing metric:

1. Right-click the metric and select New Metric Grouping from Metric <Name> from the menu.
2. Accept the default name for the metric grouping.
3. Choose a Management Module to contain the metric grouping in one of these ways:
   - Select a Management Module from the drop-down list box.
   - Click Choose, select a Management Module from the list, then click Choose again.

CA Technologies Internal / Partner Use Only

4. Click OK.

   The new metric grouping you created is highlighted in the Management Module Editor tree, under the Management Module in which you saved it.

   Note: The metric grouping is active when it is created, and cannot be de-activated.

   In the Settings pane for the metric grouping, metric grouping Agent Expressions and Metric Expressions fields already contain the metric information.

   Note: Filtering based on agent name uses the property introscope.agent.perfmon.agentExpression. A valid agent expression would be:
   `ProcessName|AgentName`

   or
   `*|MyAgent`

5. Select the Description tab to enter descriptive text and any important information about the metric groupings in the Description Text field.

   This field should contain no more than 64 KB of data. After it is applied, it will be persisted in the management module jar.

6. Select which Agent Expressions to use:

   - Select Use Management Module Agent Expressions to use the Agent Expressions defined for the Management Module
   - Select Use Metric Grouping Agent Expressions to use Agent Expressions defined for this metric grouping

7. Click Apply.

Create a new metric grouping from the Elements menu

To create a new metric grouping from the Elements menu:

1. From any Workstation window, select Workstation > Management Module Editor.
2. In the Management Module Editor window, select Elements > New Metric Grouping.
3. In the Name field, enter a name for the metric grouping.
4. Choose a Management Module to contain the metric grouping in one of these ways:

   - Select a Management Module from the drop-down list box.
   - Click Choose, select a Management Module from the list, then click Choose again.

5. Click OK.

   The metric grouping you just created is highlighted in the Investigator tree, and its settings appear in the settings pane.

   Note: The metric grouping is active when it is created, and cannot be de-activated.

CA Technologies Preferences Documentation. Authorized Use Only.

6. Enter specific agent and metric information in the Metric Grouping Agent Expressions and Metric Expressions fields.

   You can enter information in one of these ways:

   - type in the information in a regular expression
   - open another Investigator window, select a metric, and drag the information to the metric groupings window, so that the blue line appears around the Metric Grouping Agent Expressions and Metric Expressions fields.

   When you release the mouse button, the Metric Grouping Agent Expressions and Metric Expressions fields automatically fill with metric information.

7. Select which Agent Expressions to use:

   - Select Use Management Module Agent Expressions to use the Agent Expressions defined for the Management Module.
   - Select Use Metric Grouping Agent Expressions to use Agent Expressions defined for this metric grouping.

8. Click Apply.

Adding another metric to a metric grouping

To add another metric to an existing metric grouping:

1. In the metric grouping settings pane, click Add.

   This adds another Agent/Metric Expressions field pair to accept the second metric's information.

2. Launch another Investigator window.
3. In the second Investigator window, select the metric you want the new metric grouping to display, drag it to the settings pane, and this time drop it on the Metric Grouping Agent Expressions field.
4. Reselect the metric from the second Investigator window, and drag it onto the Metric Expressions field.

Customizing metric groupings

If you require more customized information than that achieved from dragging and dropping data, you can customize the regular expressions in the metric grouping by editing the Metric Grouping Agent Expressions and Metric Expressions fields to specify the metrics to match.

CA Technologies Prerelease Documentation. Authorized Use Only.

To customize metric groupings:

1. To edit the definition, follow these rules:
   - Separate the successive levels of the Investigator tree with backslash-pipe symbols. (The backslash acts as an escape character.)
     - In Metric Grouping Agent Expressions: *Host*\|*Process*\|*AgentName*
   - Use ([^\|:]*) to represent one Resource segment.
     - In Servlets\|([^\|:]*):*Average Response Time \(ms\)*
   - An escape character (backslash) is required for separators and parentheses \| and \(and \).
     - In Servlets\|Servlet1:*Average Response Time \(ms\)*
   - To match several things with one expression, you can include lists of things between parentheses using pipe characters.
     - In Servlets\|Servlet(1|14|18):*Average Response Time \(ms\)*.
   - If there are no Resource folders between the AgentName and the metric, enter only the metric name. Otherwise, separate the Resource folders with backslash-pipe symbols and precede the metric name with a colon (:).
     - In Metric Expressions: *resource*\|*subresource:Metric*
     - In Metric Expressions: *resource:Metric*
     - In Metric Expressions: *Metric*

     For example, in Metric Expressions you specify the average JDBC query time for a servlet called OptionReport as Servlets\|OptionReport\|JDBC:Average Query Time.
   - Use (.*) to represent "any."
     - For example, Cherubim\|PhoneHome\|(.*) followed by Sockets:Output Bandwidth would specify the output bandwidth for all sockets for any instance of the process PhoneHome running on the host Cherubim.
     - An entry of File System:(.*) in the Metric Expressions field means that the data to be displayed is the file input and output metrics found in the Investigator under File System. In contrast, File System:File Input Rate displays the file input rate only.
     - Use (.*)\|(.*)\|(.*) in the Agent field to make the metric grouping display data from any server, any process, and any agent. Or, you can specify any or all of the segments to match agents with a given host, process and/or agent name.
2. If necessary, click Add to specify additional metrics for the metric grouping.
3. Click Apply.

CA Technologies Proprietary. Documentation. Authorized Use Only.

Create and Edit Dashboards

Users with write permission to a domain can create and edit data viewers and other dashboard objects such as imported images, shapes, lines, and text.

Introscope dashboards allow total layout control of objects on a dashboard.

You create data viewers in the dashboard editor window by doing one or both of the following:

- Creating data viewers automatically by dragging and dropping data from the Investigator onto a dashboard
- Creating an empty data viewer in the dashboard editor, then adding the data to the viewer

Both of these options are described in detail in this chapter.

About dashboard objects

There are two types of dashboard objects: data viewer objects, and shapes, images and text objects.

Data viewer objects

Depending on the type of metric or element selected, the Workstation can display the data in a data viewer as these objects:

- Graph
- Application Triage Map
- String Viewer
- Bar Chart
- Text Viewer
- Dial Meter
- Alert Status Indicator
- Equalizer

You can import these objects to a new or existing dashboard through either of these methods:

- Dragging and dropping (see page 285) from the Investigator Browse Tree (or, in the case of the Application Triage Map element, dragging and dropping from the Map Tree).
- Using the tools palette

CA Technologies Prerelease Documentation. Authorized Use Only.

Formatting text in string or text objects

You can format the text that appears in a dashboard as a string viewer or text viewer.
See Formatting text in string viewers and text viewers (see page 283).

Shapes, images and text objects

You can add text blocks, images, shapes, and lines to the dashboard, to help explain and illuminate your data.

For example, you could:

- add a conceptual diagram of your application environment
- add your company logo to your dashboards
- insert images of your products
- add text blocks to the dashboard to explain dashboard elements in your company's language
- draw a rectangle without a fill to visually group items on a dashboard
- draw arrows to point to an object and add emphasis
- draw connectors between objects to create a simple flow chart

Graphic and text objects

Graphic and text objects correspond to simple tools found in image editing programs, with the addition of the imported graphic.

- Rectangle
- Basic Line
- Rounded Rectangle
- Straight Connector
- Oval
- Elbow Connector
- Polygon
- Imported Graphics
- Scribble (freehand drawing tool)
- Text box

CA Technologies Prerelease Documentation. Authorized Use Only.

Creating dashboards

By creating new dashboards, you can create collections of different data viewers for different uses. For example, you might have one dashboard containing database information, and one dashboard for system Alerts.

Creating a new dashboard in the Console

You can create a new dashboard from a command in the Console.

To create a new dashboard in the Console:

1. In the Console, select Dashboard > New dashboard.
2. Enter a name for the new dashboard, and choose a Management Module to contain the dashboard.
   - Select a Management Module from the drop-down list.
   - Click Choose, select a Management Module from the list, and click Choose again.
3. Click OK.

   The new dashboard opens in the dashboard Editor.
4. Edit the dashboard to suit your needs, as described in Editing a dashboard (see page 281).
5. Select File > Save.

   The new dashboard appears in the Management Module Editor tree, under the domain and Management Module in which you saved it.
6. If you are finished editing the dashboard, select Workstation > Close Window.

Creating a new dashboard in the Management Module Editor

To create a new dashboard in the Management Module Editor:

1. Select Elements > New dashboard.
2. Enter a name for the new dashboard, and choose a Management Module to contain the dashboard.
   - Select a Management Module from the drop-down list box.
   - Click Choose, then select a Management Module from the list and click Choose again.

CA Technologies Prerelease Documentation. Authorized Use Only.

3. Select the Description tab to enter descriptive text and any important information about the dashboard in the Description Text field.

   This field should contain no more than 64 KB of data. After it is applied, it will be persisted in the management module jar.

4. Click OK.

   The new dashboard appears, highlighted, in the Management Module Editor tree.

Editing a dashboard

To add or manipulate dashboard contents, you can:

- open the dashboard directly for editing in the Console.
- open the dashboard for editing from the Management Module Editor tree.

Only a user with write access to a domain or SuperDomain can edit a dashboard.

To open a dashboard in the dashboard Editor:

1. In the Console, select the dashboard tab to activate it.
2. Select Dashboard > Edit Dashboard.

   The dashboard editor opens.

To open a dashboard for editing from the Management Module Editor:

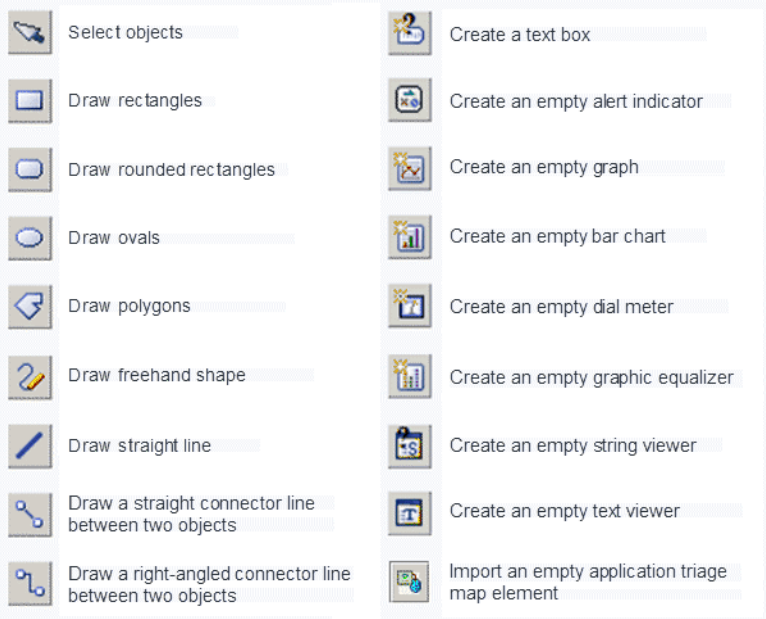
1. Select the dashboard in the Management Module tree.
2. With the dashboard selected, you can edit any of its elements in the editor pane.

About the tools palette

The tools palette contains all the tools for creating and editing dashboard objects. It contains standard drawing tools, tools for connecting objects, and tools for adding text. It also contains tools for drawing empty data viewers, onto which you can place data.

CA Technologies Prerelease Documentation. Authorized Use Only.

The illustration below describes the tools in the palette.



Resizing a dashboard

To resize a dashboard workspace area:

1. In the Console, select Dashboard > Edit dashboard.
2. Select Edit > Change dashboard Properties.
3. Enter new width and height values (in pixels) in the fields.
4. Enable Snap to Grid and set a grid size, if you want dashboards to snap to a specific location as you drag them.
5. Enable Clear Previous Lens Settings if you want the lens cleared each time a user selects this dashboard. See Dashboard links support agent lens (see page 302) for more information.
6. Click OK.

   The dashboard workspace area resizes to the defined size.
7. Select File > Save to save dashboard changes.

Saving a copy of a dashboard

To save a copy of a dashboard:

1. With a dashboard open in the dashboard Editor, select File > Save As.
2. Enter a name for the copy of the dashboard.

CA Technologies Prerelease Authorized Use Only.

3. Choose a Management Module to contain the dashboard.
   - Select a Management Module from the drop-down list box.
   - Click Choose, then select a Management Module from the list and click Choose again.

   Note: You can only choose a Management Module that belongs to a domain to which you have access, and that has access to all elements and metrics viewed in the dashboard.
4. Click OK.

   The new dashboard appears in the Management Module Editor tree, under the domain and Management Module in which you saved it.

Renaming a dashboard

1. With a dashboard open in the dashboard Editor, select File > Save and Rename.
2. Enter a new name for the dashboard.
3. Click OK.

Note: You can also rename a dashboard by selecting it in the Management Module Editor, and editing the name in the Preview pane.

Deleting a dashboard

To delete a dashboard from the Management Module Editor:

1. In the Management Module Editor tree, select the dashboard to delete.
   - Right-click on it, and select Delete *<Dashboard_Name>*.
   - Select Elements > Delete *<Dashboard_Name>*.

   The Delete Confirmation dialog box opens.
2. Click Yes to delete the dashboard.

Formatting text in string viewers and text viewers

To format the text belonging to string or text objects:

1. Open a dashboard for editing.
2. Place a string viewer or text viewer object on a dashboard, or select an existing one. See Creating data viewers in a dashboard (see page 284).

CA Technologies Prerelease Documentation. Authorized Use Only.

3. With the object selected, choose Properties > Text.

   In addition to the usual options for choosing the font, size, style, color and alignment of text, the dashboard text editing menu supports an optional "wrap text" property. Using this property allows the text to be constrained to a specific size, and to wrap text within that size, as opposed to the automatic re-sizing behavior that is the default.

4. In the dialog, apply formatting as desired and click OK.

   The string viewer or text viewer will display text formatted in the format you chose.

Domain enforcement in dashboard editing

The domains and Users features partition agents into specific Domains. Users have access only for certain Domains. Introscope enforces domain access as you create and edit dashboards. Any time you create or modify a data viewer in the dashboard editor, the operation is validity-checked against the domain visibility rules.

- Elements in dashboard objects in the SuperDomain can reference elements and data in any domain
- Elements in dashboard objects in a user-defined domain can only reference elements and data in the same domain

For more information on Domains and domain enforcement, see the *APM Installation and Upgrade Guide*.

Data type	Default data viewer type
Metric	graph
Metric grouping	graph
Alert	status indicator
Calculator	graph

CA Technologies Prerelease Documentation Authorized Use Only

Create Data Viewers in a Dashboard

There are two ways to create data viewers in Introscope:

- Select one of the following:
  - a metric in the Investigator tree
  - a metric grouping
  - an element from the Management Module Editor tree

  ... and drag it onto a dashboard, automatically creating the default data viewer for that type of information.

- Create an empty data viewer, then add information to it. (See Creating an empty data viewer and adding data (see page 288)).

Note: When you create or make changes to a dashboard, always save the dashboard as the last step. Although it isn't necessary to save changes to a dashboard after every individual edit, it is important that you save changes frequently to make changes available to other Workstation users with access to that dashboard.

Creating a data viewer automatically

The easiest way to create a data viewer is to select a metric (or element, etc.) from the Management Module Editor tree, and drag-and-drop it onto the dashboard Editor window. You can drag-and-drop in two ways:

- Drag the data onto an existing data viewer in a dashboard (the new data replaces the existing data in the viewer).
- Drag the data to an empty area on a dashboard, which automatically creates a new viewer to contain it.

These are the default data viewer types for each of the data types:

Data type	Default data viewer type
Metric	graph
Metric grouping	graph
Alert	status indicator
Calculator	graph

Create a Data Viewer by Dragging and Dropping Data

The easiest way to create a data viewer automatically is to select something (metric, metric grouping, alert, etc.) from the Investigator tree and drag it onto the dashboard.

Note: Dashboards with objects imported by dragging and dropping from the Investigator are not supported in WebView. WebView users who open the dashboards will see a box where the object should be, with an error message.

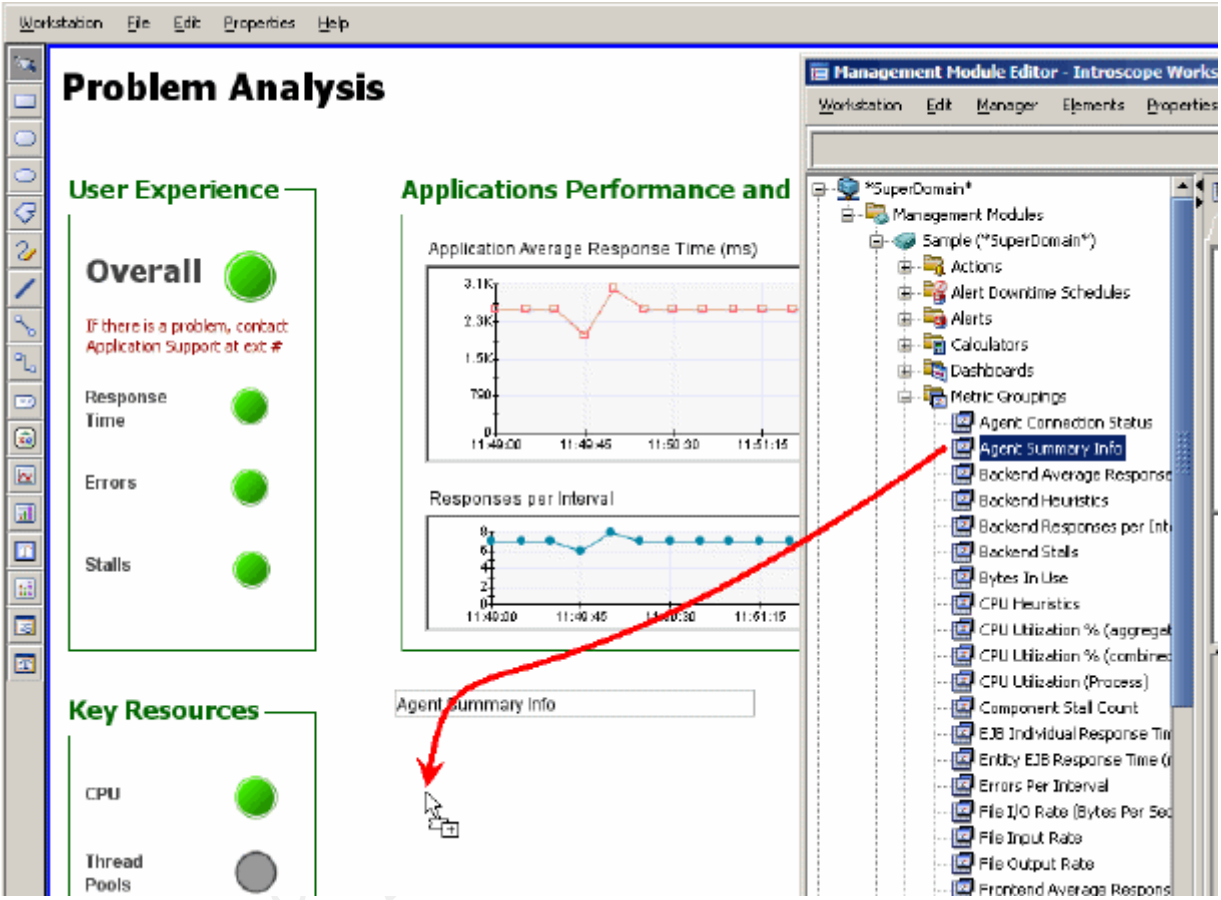
To create a data viewer by dragging and dropping data:

1. In the Console, select Dashboard > Edit dashboard.
2. Open an Investigator window and position it so both the Investigator window and dashboard are visible.

CA Technologies Prerelease Documentation. Authorized Use Only.

3. In the Investigator, click and hold a metric in the tree. Drag it to the dashboard and drop it when you see the blue highlighted line all around the dashboard.

   The illustration below shows what it looks like when you drag a metric from the metric browser tree into a dashboard.



Note: If you attempt to drag something to a dashboard in violation of domain enforcement rules, you do not see the blue highlight on the dashboard and nothing appears on the dashboard when you release the mouse button.

If you drop the selection when only an existing data viewer is highlighted, and the data viewer type is compatible with the selection, the information you are dragging replaces what is in that viewer.

4. Select File > Save to save dashboard changes.

Dragging contents of a tab view to a dashboard

You can also add the contents of a tab view from the metric browser tab to a dashboard by dragging and dropping. See Embed an Investigator Tab View in a Dashboard (see page 298).

Add an Application Triage Map to Dashboards

You can import any Application Triage Map as a dashboard object. Once it is placed in the dashboard, its elements have the same interactive functions as they have when the map is displayed in the Investigator. You can view current alert states, hover to see metrics, double-click to open the Locations pane, select map elements, and so on. In addition, the scrollbars, toolbar, and Locations pane all function normally, and the context menus appear as usual on right-click.

Note that actions that cause a different map or data overview tab to be displayed in place of the current map will instead bring up an Investigator window with the requested map or overview.

For more information on how the Application Triage Map works, see Using the Application Triage Map (see page 87).

Before you begin:

- Open a new or existing dashboard for editing, and identify the spot to place the map element.

To import the Application Triage Map to a dashboard:

1. Click the map import button:

   The cursor changes to a rectangle creator.

2. Click and drag the cursor diagonally across the open area on the dashboard.

   When you release the mouse button, the application triage map element fills the space you designated.

3. Set the map element to be populated with data from a Business Service, Business Transaction, or Frontend.

   a. Right-click the map element placeholder.

   b. Under Data Type, select from among Business Service, Business Transaction, or Frontend.

   c. Under Name, select the individual data source.

   d. Click OK.

   The map will display data from the selected data source. See more information on using the Data Options dialog (see page 288).

NOTE: Adding multiple maps on a dashboard will lead to degraded performance. CA Technologies recommends no more than five maps on a single dashboard.

CA Technologies Prerelease Documentation. Authorized Use Only

Creating an empty data viewer and adding data

In this process, you create a data viewer first, then specify the data to appear in the viewer. You create a data viewer using the Tools Palette.

There are two ways to add data to an empty data viewer.

- Drag and drop data from the Investigator tree (see page 288).
- Add data through the Data Options dialog (see page 289).

Creating an empty data viewer

To create an empty data viewer:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select a dashboard object creation tool from the Tools Palette.
3. Click and drag the location and size of the empty data viewer on the dashboard area.
4. Release the mouse button.

   The data viewer appears on the dashboard as drawn. You can now manipulate the data viewer.
5. Select File > Save to save dashboard changes.

Adding data to a data viewer by dragging and dropping

You can add data to a data viewer by dragging and dropping from the Investigator or the Management Module Editor.

You can use this function to:

- Add data to an empty data viewer.
- Replace the data displayed in existing data viewer.

To add data to a data viewer by dragging and dropping:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the data to add to the dashboard:
   - To add a metric to the dashboard, open an Investigator window and position it so both the Investigator window and dashboard are visible, then click and hold a metric in the tree.
   - To add an element to the dashboard, open the Management Module Editor, and click and hold an element in the tree.

CA Technologies Prerelease Documentation. Authorized Use Only

3. Drag the metric or element to the dashboard and drop it when you see the blue highlighted line around the data viewer.

   Note: If you try to drag a metric or element to a dashboard that violates domain enforcement rules, you do not see the blue highlight on the dashboard, and nothing populates the data viewer when you release the mouse.

   If the data viewer already contains data, it is replaced by your new data when you drop it.

4. Select File > Save to save dashboard changes.

Adding data to data viewer using the data options dialog

To add data to data viewer using the data options dialog:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Right-click on the empty data viewer and select Data Options.

   The Data Options window opens.

3. From the left side of the window (Data Type), select a data type for this data viewer. Data Type options differ depending on the data viewer type. Possible options are Metric grouping, alert, calculator, or metric.
4. From the Data Selection list, select the data selection to appear in the data viewer.
5. Click OK.

   The data viewer is populated with the data selection.

6. Select File > Save to save dashboard changes.

How to tell which data appears in a data viewer

Some data viewers (such as alerts) do not display the name of the metrics that supply its data.

To discover which data is displayed in a data viewer:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Right-click on the data viewer and select Data Options...

   The Data Type and exact Data Selection that are used are highlighted in the Data tab.

3. Click Cancel or OK to close the Data Options window.

CA Technologies Proprietary & Documentation. Authorized Use Only.

Setting data-viewing properties of a data viewer

In the dashboard Editor window, you can change these properties of a data viewer:

- Viewer display type
- Scaling options
- Sorting/filtering options
- Whether labels are on or off
- View period

Data viewer display options

In the Workstation, you can define views for almost all types of metrics. These views can appear as several view types, depending on the data defined in the metric.

Because metric data can consist of different information (text, dates, counters, numbers, and so forth), not all data can appear in every data viewer type. For example, the data from the metric Java Version cannot appear as a Graph, because its data is text.

To find out what view display types are available for the selected view, right-click on the data viewer and look at the View As... submenu.

To change the data viewer display type:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the data viewer whose display type you want to change in one of these ways:
   - Right-click and select View As and the new data viewer display type from the submenu
   - Select Properties > View As and select the new data viewer display type from the submenu
3. Select File > Save to save dashboard changes.

Changing scaling options for a data viewer

You can change the scale of graph charts while viewing live data in Introscope Workstation, to provide a more readable view. You change the scale of a chart by setting a minimum and maximum value for the chart's data axis. See Changing the scale of graph charts (see page 148).

CA Technologies Prerelease Documentation. Authorized use only.

Changing Sorting and Filtering Options

You can define sorting and filtering options for Graph and Bar Chart viewers.

In the Sorting/Filtering Options tab, you can:

- Enable/disable filtering.
- Switch between top/bottom metric filtering.
- Specify how many metrics to filter (default is 10).
- Add, remove, or clear included or excluded metrics.

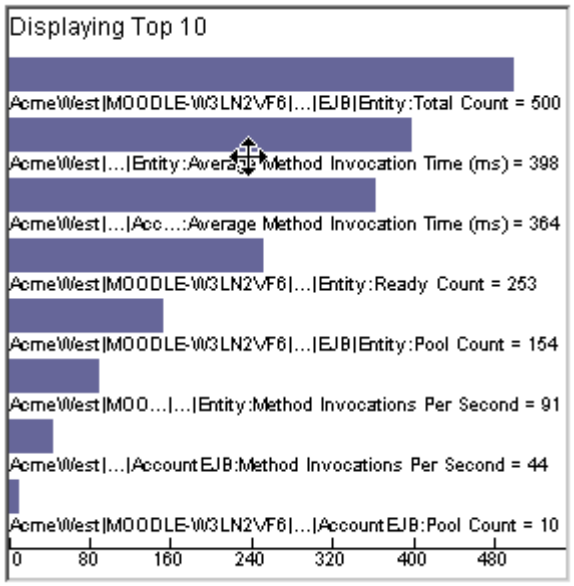
Note: While sorting/filtering is in use, these metric viewing options are not available: bring to front, send to back, and hide/show metric.

Defining type and number of metrics shown in Filtered View

To define type and number of metrics shown in Filtered View

1. Check the Enable Filter checkbox.
2. In the Show the Top/Bottom pull down menu, choose whether to show the Top N or Bottom N metrics.
3. In the X Metrics field, enter the number of metrics to view in the Filtered Views list. The default number is 10.
4. Click OK.

   The data viewer shows the number of top or bottom metrics you defined, and identifies the number of metrics it is displaying.



CA Technologies Pre-release Documentation. Limited Use Only.

5. Select File > Save to save dashboard changes.

Including selected metrics

To include selected metrics:

1. On the Sort/Filter tab, in the Included Metrics area, click Add.
2. The window lists metrics that currently match the metric grouping, sorted alphabetically. Metrics that are already defined in the Included or Excluded metrics lists do not appear.
3. Click to select the metric to add to the Included Metrics list. Select multiple metrics using Shift or Ctrl key.
4. Click OK, then OK again.

   The included metrics appear in the Included Metrics list in the Sorting/Filtering Options window.

   A message notifies you if no additional metrics are available.
5. Select File > Save to save dashboard changes.

Excluding selected metrics

You can exclude any metrics, regardless of whether they are ranked in your defined top or bottom viewer.

To exclude selected metrics:

1. In the Excluded Metrics area of the Sort/Filter tab, click Add.

   The window lists metrics that currently match the metric grouping, sorted alphabetically. Metrics already defined in the Included or Excluded metrics lists are not shown.
2. Click to select the metric to add to the excluded metrics list. Select multiple metrics using Shift or Ctrl key.
3. Click OK, then OK again.

   The excluded metrics appear in the excluded metrics list in the Sorting/Filtering Options window.

   A message notifies you if no additional metrics are available.
4. Select File > Save to save dashboard changes.

CA Technologies Prerelease Documentation
Authorized Use Only

Removing selected metrics from included or excluded metric lists

To remove individual metrics from the included or excluded metrics lists:

1. In either the Included or Excluded metrics list, click on the metric you want to remove from the list. Select multiple metrics using Shift or Ctrl key.
2. Click Remove.
3. Click OK.

   The metric is removed from the Included or Excluded metric list.
4. Select File > Save to save dashboard changes.

Clearing all included or excluded metrics

To clear all metrics from the included or excluded metrics list:

1. In either the Included or Excluded metrics list, click Clear All.
2. Click OK.

   All metrics in the list are removed.
3. Select File > Save to save dashboard changes.

Changing the view resolution

You can change the view resolution on every viewer type except for an alert status indicator. Note that this only changes the view resolution in the data viewer, not the collection resolution for any referenced metric groupings or other elements.

To change the view resolution for a data viewer:

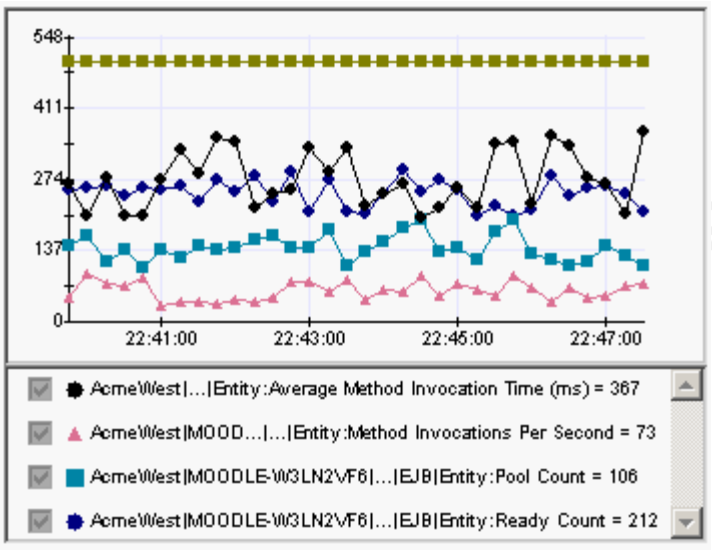
1. Open the dashboard in the dashboard Editor.
2. Select the data viewer with the Selection tool and right-click and select Data Options and then click the Miscellaneous tab.
3. Select a new view resolution and click OK.

   The viewer displays historical data using the new time-resolution value. Changes do not affect live mode.
4. Select File > Save to save dashboard changes.

CA Technologies Prerelease Documentation. Authorized Use Only.

Turning labels on and off

Labels that you can turn on and off refer to the metric information shown in a data viewer. The illustration below shows labels under the graphic chart.



You can turn labels on and off in:

- Graphs
- Bar charts
- Dial meters
- Graphic equalizers
- String viewers
- Text viewers

To turn labels on and off in a data viewer:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the data viewer with the Selection tool.
3. Select Properties > Data Options, then click the Miscellaneous tab.
4. Turn Labels on or off in one of these ways:
   - Check the Show Labels or Legend checkbox to turn Labels on.
   - Uncheck the Show Labels or Legend checkbox to turn Labels off.
5. Click OK.
6. Select File > Save to save dashboard changes.

CA Technologies Prerelease - Authorized Use Only

Changing alert status indicator options

Alert status indicator viewers can display:

- An alert status indicator, which contains three symbols; only one of the three symbols is colored to show the current status. For example:



- A single alert status indicator, which contains only one indicator; the indicator changes color and shape according to current status.

To change the options for an alert status indicator:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select an alert status indicator with the Selection tool and right-click and select Data Options, then click the Alert tab, or, select Properties > Data Options, then click the Alert tab.
3. In the Alert Type field, use the pull-down menu to choose either a Single Indicator or Multiple Indicator option.
4. Click OK.

   The alert status indicator changes to the selected view.
5. Select File > Save to save dashboard changes.

Creating dashboard text and graphics

This section has instructions for using the drawing palette to create text and graphics on dashboards:

- Adding shapes and lines to a dashboard (see page 296)
- Drawing connector lines and adding arrowheads (see page 296)
- Coloring shapes, lines and connectors (see page 296)
- Creating and editing text (see page 296)
- Inserting an image on a dashboard (see page 297)
- Manipulating dashboard objects (see page 299)

CA Technologies Prerelease Documentation. Authorized Use Only.

Adding shapes and lines to a dashboard

Use the dashboard Editor to add shapes, lines, text, and images to your dashboards. You can add the following shapes to a dashboard:

- Rectangle
- Rounded rectangle
- Oval/circle
- Polygon
- Scribble (freeform drawing tool)
- Line

After you draw a shape or line with any tool, the Selection tool is automatically selected, so you can then move or resize the shape or line.

Drawing connector lines and adding arrowheads

You can show a relationship between two or more dashboard objects by connecting them with either a straight connector or an elbow connector. Using a connector line enables you to move the two dashboard objects while maintaining the connection.

Coloring shapes, lines and connectors

You can add a fill color to shapes and a pen color to shapes and lines, selecting from standard or custom colors.

Creating and editing text

The Label tool is useful for adding descriptive text boxes to add context to your data viewers.

Note that you can only change the font attributes for user-created text blocks (created with the Label tool). You cannot change font attributes for Legend text in a data viewer.

Creating text on a dashboard

To create a text block on a dashboard:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Click the Label tool.

CA Technologies Preferred Use Documentation. Authorized Use Only.

3. Click anywhere on the dashboard to create a placement point. A text box opens.
4. Type in your text. To create text boxes with multiple lines, press Enter at the end of a line of text—the cursor moves to the next line in the box for you to resume typing.
5. To exit the text entry field, click anywhere on the dashboard outside the text field, or choose another tool.
6. Select File > Save to save dashboard changes.

Editing text

To edit existing text:

1. Click the Label Tool and click on the text block. The text is highlighted and you can edit it.
2. Click anywhere outside the text block to deselect it.
3. Select File > Save to save dashboard changes.

Changing text attributes

You can change the font, font size, color, background color, and style of text in a text block.

To change text attributes:

1. Click the Selection tool to select the text block to modify.
2. Select File > Save to save dashboard changes.

Changing text size freehand

You change the size of text freehand by resizing the text block.

To change text size freehand:

1. Use the Selection tool to select the text block to resize.
2. Click on the yellow dot and drag in or out to resize text.
3. Click anywhere outside the text block to deselect it.
4. Select File > Save to save dashboard changes.

Inserting an image on a dashboard

An easy way to give your dashboards more context is to import an image—for example, representations of network components, or a company logo. Introscope provides some basic network images (located in the *<EM_Home>/images* directory), or you can insert any graphic file of *.JPEG* or *.GIF* format.

CA Technologies Prerelease Documentation. Authorized Use Only.

Note: Introscope does not support animated *.GIF* files.

To insert an image onto a dashboard:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select Edit > Insert Image to browse for an image.
3. Browse to select an image, then click Open.

   The image is inserted in the upper left corner of the dashboard.
4. Click the Selection tool, then click on the image to select it and move it.
5. Select File > Save to save dashboard changes.

Embed an Investigator Tab View in a Dashboard

You can embed the contents of a tab view in a dashboard by dragging and dropping.

To add the contents of an Investigator tab view to a dashboard:

1. Create a new dashboard, or open an existing dashboard in edit mode.
2. In a separate window:
   a. Open a new Investigator
   b. Browse in the Investigator tree to the metric you want to include in your dashboard.
   c. In the Investigator viewer pane, click on the tab whose contents you want in your dashboard.
3. Click on the tab in the Investigator and drag it to the editable dashboard.

The contents of the tab view are imported to the dashboard. Note that when you drag a metric or tab in live mode, the clock at the bottom of the object will continue to tick.

Note: The following actions are not supported:

- Dragging and dropping from the triage map tab
- Dragging and dropping from the location map tab

See Tab Views in the Metric Browser Tab (see page 131) for a description of Investigator tab views.

CA Technologies Prerelease Documentation. Authorized Use Only.

Manipulating dashboard objects

You can use the dashboard editor to perform these actions with dashboard objects:

- Select
- Move
- Resize
- Resize graph legend type
- Cut, copy, paste, delete
- Align
- Arrange dashboard objects front to back
- Grop and ungroup
- Connect with connectors

Dashboard objects, except for connectors, can be moved and placed anywhere on the dashboard area. Both line connectors and elbow connectors are repositioned automatically when the objects to which they are connected are moved.

Most dashboard objects can be resized, with some restrictions:

- If multiple dashboard objects are selected, only the dashboard object whose handle is used is resized.
- Grouped dashboard objects cannot be resized.
- Polygons are resized in a special way.
- Scribbles cannot be resized, although you can reposition the points in the scribble, thus changing the length of the segments that make up the scribble.
- When data viewers are resized, they simply get bigger or smaller—the representation of data does not change (scale does not change, no more or fewer points are visible).

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the object with the Selection tool

   The dashboard object's handles appear.

CA Technologies Prerelease Documentation. Authorized Use Only

3. To select multiple dashboard objects, hold down the Shift key, and click on each object to add it to the selection.

   You can also use the Shift key to deselect an object from a multiple-object selection.

4. Make changes as needed.

   To move a dashboard object, use the selection tool and drag the object. You can also move the object in small increments by selecting the object, then using arrow keys on your keyboard.

   Note: Dashboard objects, except for connectors, can be moved and placed anywhere on the dashboard area. Both line connectors and elbow connectors are repositioned automatically when the objects to which they are connected are moved.

   To resize a dashboard object, drag the object's handles. To resize a polygon, select it and drag the yellow dot.

   To cut, copy or paste a dashboard object, select it, then use commands under Dashboard > Edit Dashboard.

   Note: Introscope only allows objects to be pasted into dashboards within the same domain as the dashboard they were copied or cut from.

   To delete a dashboard object, select it and press your Delete key.

   Note: Dashboard objects are not displayed as elements in the Management Module Editor tree; you must delete dashboard objects in the dashboard Editor window. If multiple items are selected, they all are deleted.

5. Select File > Save to save dashboard changes.

Resizing graph legend size within a graph data viewer

Resizing a graph by using its handles resizes the overall graph size.

To resize the graph legend area:

1. Make sure the graph is showing Labels/Legend.
2. Click and drag the yellow diamond up or down to make the Graph legend area smaller or larger relative to the Graph area.

Note: If you replace the data in this Graph with data that contains additional metrics, you might need to resize the Graph to show all the metrics (although scroll bars appear if the metrics cannot all be viewed).

CA Technologies Proprietary Documentation. Authorized Use Only.

Aligning dashboard objects

To align two or more dashboard objects:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the dashboard objects to be aligned.

   If only one dashboard object is selected, the Align commands are disabled.
3. Select an Align command from the Edit > Align menu:

   The dashboard objects are aligned as specified.
4. Select File > Save to save dashboard changes.

Arranging dashboard objects front to back

You can layer dashboard objects, and move them in front of and in back of each other.

To move objects to front and back:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the object to be moved with Selection tool and move it in one of these ways:

- Right-click on object and select either Bring to Front or Send to Back from the menu.
- Select Edit menu and then select either Bring to Front or Send to Back.

1. Select File > Save to save dashboard changes.

Grouping and ungrouping objects

In dashboards with many objects, it might be helpful to group the objects to make placement easier. You can also group multiple grouped objects.

To group selected objects:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the objects to be grouped (shift+click to add objects to a selection), and select Edit > Group, or right-click on the group of objects and select Group from the menu.

   Selected objects are grouped into one item, represented by one set of handles.

   Note: You cannot resize grouped objects.
3. Select File > Save to save dashboard changes.

CA Technologies Prerelease Documentation. Authorized Use Only.

To ungroup a grouped object:

1. Select the object to be ungrouped, and select Edit > Ungroup, or right-click on the grouped object and select Ungroup from the menu.

   The grouped object is separated into individual objects, all with their handles highlighted.

   Note: If the grouped object contained another grouped object, this object remains grouped.

2. Select File > Save to save dashboard changes.

Creating and managing custom hyperlinks

You can use custom links to create hyperlinks between dashboard objects and other dashboards, or external Web pages. A dashboard object can have multiple links of multiple types. Custom hyperlinks are available to any Introscope user.

Dashboard links support agent lens

You can associate an agent lens with a dashboard link, so that it is applied each time the link is clicked. For example, if you have an Overview dashboard with alert status indicators for multiple agents, you can link each alert to the same dashboard, and set a dashboard lens for each link to specify the agent associated with the alert.

You can set Edit > Change dashboard Properties > Clear Previous Lens Setting for a dashboard, to clear the lens each time a user selects that dashboard. For example, if the Clear Lens option is set for the Overview dashboard, when a user returns to the Overview dashboard from another dashboard to which a different lens is applied, that lens is cleared so that the Overview dashboard shows data for multiple agents, as intended.

dashboard lens settings are part of the navigation history. When you navigate to a previously viewed dashboard using the Back button, a lens that was previously applied is reapplied.

Creating a custom link to a dashboard

To create a custom link to a dashboard:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the dashboard object to contain the custom link. These dashboard objects cannot contain links: Line, Scribble, Connector, and Elbow Connector.

CA Technologies Pre-release Documentation. Authorized Use Only.

3. With the dashboard object selected, right-click on the dashboard object and select Object Links from the menu. Or, select Properties > Object Links.

   The Object Links dialog box opens.

4. Click Add.

   The Add Object Link dialog box opens.

5. Select the dashboard Link radio button.

6. Select a dashboard from the dashboard drop-down list and click Choose.

   The Select Agent Lens dialog opens.

7. Select a single agent, or select multiple agents (click and drag, or CTRL/click) on which to filter.

   Note: You can begin typing an agent name, hostname, or process name in the Search field. As you type, the agent list filters to match what you type.

8. Click OK in the Select Agent Lens dialog, then click OK in the Add Object Links dialog.

   The new dashboard link appears in the Object Links dialog box.

9. Click OK to exit the Object Links dialog box.

10. Select File > Save to save changes to the dashboard.

    The link is now accessible to users in the Workstation.

Creating custom link to an external Web page

To create a custom link to an external Web page:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.

2. Right-click on the dashboard object and select Object Links from the menu, or select Properties > Object Links.

   The Object Links dialog box opens.

3. Click Add. The Add Object Link dialog box appears.

4. Select the Web link radio button.

5. In the Name field, enter a name for the Web link. Keep the name short and descriptive, because it appears under the Links menu.

CA Technologies Prerelease Documentation. Authorized Use Only

6. In the URL field, enter the address of the Web page link.

   Note: The URL must be fully specified or it will not work correctly on all platforms. For example, instead of entering *wilytech.ca.com*, you must enter *http://wilytech.ca.com*.

7. Click OK.

   The new web link appears in the Object Links dialog box.

8. Click OK again to exit the Object Links dialog box.

9. Select File > Save to save changes to the dashboard.

   The link is now accessible to users in the Workstation.

Defining default links

A default link can be accessed by double-clicking it. If a dashboard object contains only one custom link, that link is automatically treated as the default link.

Note: There can only be one default link per object.

To specify a link as a default link:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.

2. Select the dashboard object that contains the custom link and right-click on the dashboard object and select Object Links from the menu, or select Properties > Object Links.

   The Object Links dialog box appears.

3. Select the row that contains the link for which to define a default link for.

4. Click Set As Default.

   The defined default link appears in bold in the list in the Manage Links dialog box.

   - To change the default link, click on a different link, and click Set As Default. The default link changes.
   - To clear the default link, click Clear Default.

5. Click OK.

6. Select File > Save to save changes to the dashboard.

   This quick link can be accessed by a user by double-clicking on the object with the custom default link.

CA Technologies Pre-Release Documentation. Authorized Use Only.

Editing custom links

You can perform these edits to links:

- for a dashboard link, choose a different dashboard
- for a web link, edit the name or URL
- change the link type from dashboard to Web link
- apply a lens to a dashboard link.

To edit custom links:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the dashboard object to contain the custom link and select Properties > Object Links, or right-click on the dashboard object and select Object Links from the menu.

   The Object Links dialog box opens.
3. .Select the row that contains the link to edit.
4. Click Edit.

   The Edit Object Link dialog box opens.
5. Edit the link as appropriate. Click OK.
6. Click OK again to exit the Object Links dialog box.
7. Select File > Save to save changes to the dashboard.

Removing links

To remove a custom link:

1. In the Console, open the dashboard to edit by selecting Dashboard > Edit dashboard.
2. Select the dashboard object to remove the custom link from and right-click on the dashboard object and select Object Links from the menu, or select Properties > Object Links

   The Object Links dialog box opens.
3. Select the row that contains the link to delete.
4. Click Remove.
5. Click OK.
6. Select File > Save to save changes to the dashboard.

CA Technologies Pre-release Documentation. Authorized Use Only.

Monitoring performance with alerts

There are two types of alerts: Simple and Summary.

- A Simple Alert takes status information produced by a comparison as input, compares it to user-defined threshold values, and outputs a status.
- A Summary Alert bases its status on the status of multiple Simple Alerts and other Summary Alerts.

Both types of alerts appear together under the Alerts node in the Investigator tree.

About Simple Alerts

A Simple Alert takes status information produced by a comparison as input, compares it to user-defined threshold values, and outputs a status. A Simple Alert has one of four states:

- Not reporting data—this could happen if the Simple Alert is not matching any metrics, if the metrics it is matching are not reporting (perhaps because they are shut off), or if the Simple Alert itself is inactive
- Green (OK)
- Yellow (Caution)
- Red (Danger)

You can define actions to be triggered for Caution or Danger states. Simple Alerts can use Danger and Caution action delays (SmartTrigger functionality) to determine when to initiate specified actions.

The Simple Alert is the foundation alert in Introscope. Simple Alerts can trigger actions and notifications, or can provide input for a Summary Alert. To create a Simple Alert, see Creating Simple Alerts (see page 309).

How alerts are defined using heuristic metrics

Each alert indicator on the sample dashboards is based on Introscope's automated heuristic modeling of standard key performance indicators as described in Application Overview (see page 133). Every key performance indicator has a matching heuristic metric.

CA Technologies Prerelease Documentation. Authorized Use Only.

The values for heuristic metrics are 1, 2 or 3:

- A value of 1 indicates that the current state of the key performance indicator appears normal.

  For example, if the application's overall response time usually varies between 600ms and 1000ms and the current value is 835ms, the response-time heuristic metric reports a 1.

- A value of 2 this indicates that the current state of the heuristic's key performance indicator is outside of normal.

  For example, if the application's CPU is usually between 30% and 60% and the current value is 75%, the heuristic value might be two.

- A value of 3 indicates that the current state of the heuristic's key performance indicator is outside of normal to a large degree.

  For example, if an application typically has no stalls or occasionally has one stall but suddenly, the application's database stops responding to requests. The number of stalls might increase to a comparably high number such as ten. In that situation, the stall heuristic for the application would report a value of 3.

By defining alerts in terms of the heuristic metrics rather than fixed thresholds, the work of determining normal values for key performance indicators shifts from the APM administrator to APM itself.

Eliminating alerting on transient spikes

A technique that is useful for defining alerts is the *At least N of the last M periods* property, which defines the number of instances in which the status of Danger triggers an alert.

In production environments, key performance indicators might spike for a short period of time. For example, a CPU might spike over a 15-second period, then return to normal in the next 15-second period.

It is undesirable for Introscope to alert on this type of spike. By telling Introscope to alert only if a condition lasts for eight out of the last eight periods (each period is 15 seconds, so two minutes out of the last two minutes), alerts are only generated for conditions that are real problems, rather than random spikes.

Example: Configuring an alert for agent disconnection

Agent disconnection is a critical event, since if an agent disconnects from the Enterprise Manager, you will no longer be able to collect or monitor data from that agent. You can set up an alert to trigger a notification that will notify you if this happens.

CA Technologies Prerelease Documentation. Authorized Use Only

To set up an alert for agent disconnection:

1. In the Workstation, expand the tree to *SuperDomain* > Custom Metric Host (Virtual) > Custom Metric Processs (Virtual) > Custom Metric Agent(Virtual) > Agents.
2. Under the Agents node, expand to *<Host_Name>*|*<Process_Name>*|*<Agent_Name>*.
3. Right-click the ConnectionStatus metric.
4. Choose New Simple Alert from Metrics "ConnectionStatus."

   This metric has the following values:

   - 3 = disconnected

     Disconnected means the agent has been manually disconnected.
   - 2 = connected, slowly or no data
   - 1 = connected
   - 0 = unmounted

     Unmounted means the agent has disconnected after a certain amount of time (depending on what the administrator has configured) during which it has reported no data to the Enterprise Manager.
5. Type a name for the new alert and click OK.
6. Set Comparison Operator to Less Than.
7. Set Trigger Alert Notification to Whenever Severity Changes.
8. Set the Danger Threshold and the Caution Threshold.

   Your threshold settings depend on how sensitive you want to make the alert.

   - A very sensitive setting would be Danger=2, Caution=2, At least *1* of the last *10* periods.
   - A less sensitive setting might be Danger=3, Caution=2, At least *3* of the last *10* periods.
9. Click Active.
10. Click Apply.

CA Technologies Proprietary - Documentation. Authorized Use Only.

Creating Simple Alerts

There are three ways to create a Simple Alert in the Workstation:

- Select the data (a metric, metric grouping, or resource), and create a Simple Alert from the data.

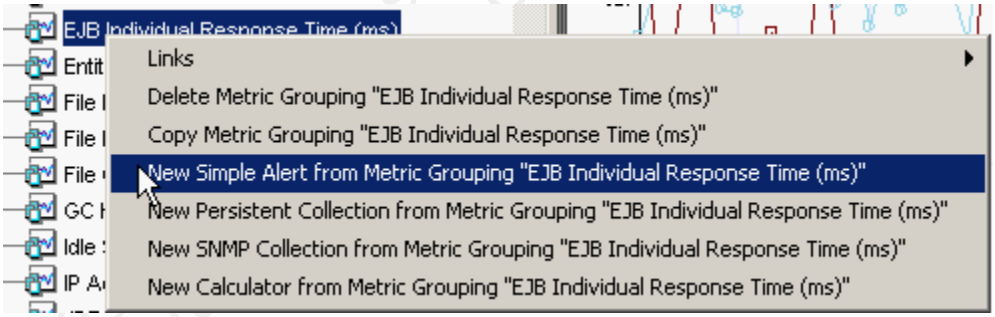
  Note: If you create a Simple Alert from a metric, a metric grouping with the same name as the Simple Alert is automatically created.

- Create a Simple Alert from the Elements menu, then add the metric information using regular expressions.
- Create a Simple Alert from a metric displayed in the Map tree.

Creating a Simple Alert from existing data

These instructions describe creating a Simple Alert by first selecting data (a metric grouping), then creating a Simple Alert using the right-click context menu.

1. In the Management Module Editor tree, right-click on the metric grouping from which to create the Simple Alert. From the menu, select New Simple Alert from Metric Grouping *<Name>*.



2. In the Name field, enter a name for the Simple Alert. The Management Module is the same one that the metric grouping belongs to.

   Note: Use informative names for alerts. When the recipient of an alert notification recieves a notification, in email or in some other way, often the only information they have is the alert name itself. So you should use a name which, as clearly as possible, helps identify the source of an alert.

3. Select the Description tab to enter descriptive text and any important information about the alert in the Description Text field.

   This field should contain no more than 64 KB of data. After it is applied, it will be persisted in the management module jar.

4. Click OK.

   Proceed to the section, Configuring Simple Alert settings (see page 311), to configure specific Simple Alert settings.

CA Technologies Product Documentation. Only

Creating a Simple Alert from the Elements menu

To create a new Simple Alert from the Elements menu:

1. In the Management Module Editor window, select Elements > New Alert > New Simple Alert.
2. In the Name field, enter a name for the Simple Alert.

   Note: Use informative names for alerts. When the recipient of an alert notification recieves a notification, in email or in some other way, often the only information they have is the alert name itself. So you should use a name which, as clearly as possible, helps identify the source of an alert.
3. Choose a Management Module to contain the Simple Alert in one of these ways:
   - Select a Management Module from the drop-down list box.
   - Click Choose, select a Management Module from the list, then click Choose again.
4. Click OK.

   The Simple Alert you just created is highlighted in the Management Module Editor tree, and its settings appear in the settings pane.

Creating a Simple Alert From a Metric in the Triage Map Tree

You can create a Simple Alert while viewing metrics displayed in the metrics tree under the triage map tab. (For more on the triage map tab display, see Using the Triage Map Tab (see page 75).)

To create a Simple Alert from a triage map tree metric:

1. Right-click the metric where it appears in the triage map tree.
2. Choose New Simple Alert From Metrics "<Metric_Name>".
3. In the Name field, enter a name for the Simple Alert. The Management Module is the same one that the metric grouping belongs to.

   Note: Use informative names for alerts. When the recipient of an alert notification receives a notification, in email or in some other way, often the only information they have is the alert name itself. So you should use a name which, as clearly as possible, helps identify the source of an alert.
4. Select the Description tab to enter descriptive text and any important information about the alert in the Description Text field.

   This field should contain no more than 64 KB of data. After it is applied, it will be persisted in the management module jar.

CA Technologies Prerelease Documentation
Authorized Use Only.

5. Click OK.

Proceed to the section, Configuring Simple Alert settings (see page 311), to configure specific Simple Alert settings.

Selecting a metric grouping to Supply Data to the Simple Alert

When you create a Simple Alert, a metric grouping with the same name as the Simple Alert is automatically created. You can either:

- Customize the newly created metric grouping with regular expression information (follow the instructions in the section).
- Choose an existing metric grouping to supply data for this Simple Alert (described here).

To select a metric grouping to supply data to the Simple Alert:

1. In the metric grouping area, choose a metric grouping to supply data to the Simple Alert:
   - Select a metric grouping from the drop-down list.
   - Click Choose, select a metric grouping from the list, then click Choose again.
2. If you are finished, click Apply to apply the changes.

   Otherwise, proceed to the next section to configure specific Simple Alert settings.

Configuring Simple Alert settings

After creating a Simple Alert (see Creating Simple Alerts (see page 309)), you define the conditions under which it is triggered.

To configure Simple Alert settings:

1. If the settings for the Simple Alert are not already visible, locate the Simple Alert you just created in the Management Module Editor tree, under the Management Module into which you placed it. Click on the Simple Alert to select it and display its settings.
2. In the settings pane, check the Active checkbox to activate the Simple Alert.
3. Configure and save the Simple Alert settings.

CA Technologies Prerelease Documentation. Authorized Use Only.

Resolution

Select or enter a time period resolution in hours, minutes, or seconds.

A Simple Alert uses input data from a selected metric grouping. For the time resolution you select, Introscope gathers information and summarizes a value for that time period. The resulting value depends on the type of data in the metric. For example, if the metric is a rate, the summarized value is the average rate during that time period. Or if the metric is a counter, it produces the most recent value of the counter.

Note: Time resolution values must be in 15-second increments.

Combination

Choose a value from the drop-down list:

- **any—**a Simple Alert is triggered when one metric goes over a threshold
- **all—**a Simple Alert is triggered only when all metrics go over a threshold

Note: The Combination field is ignored when the Notify by individual metric box is checked.

Comparison Operator

Choose a value from the drop-down list for the condition that triggers the Simple Alert, either Less Than, Greater Than, Equal To, or Not Equal To.

The Comparison Operator, along with Danger and Caution Threshold values, defines the condition that trigger the Simple Alert.

The Comparison Operator relates to the Danger and Caution threshold values. For example, if you want to be notified when an average servlet response time is greater than 5000, you would use the "greater than" operator. The Comparison operator also affects the Caution and Danger Threshold values. If the Comparison operator is set to greater than, the Danger Threshold value must be greater than the Caution Threshold value. Conversely, if the Comparison operator is set to less than, the Danger Threshold value must be less than the Caution Threshold value.

Notify by individual metric

Select to trigger individual metric alerts. You can use the Individual Metric Alert Notification and Resolution Alert together.

The Notify by Individual Metric feature (also called Metric Level alert) configures Introscope to trigger a Simple Alert status when an individual metric crosses a user-defined threshold. This is helpful if you create a Simple Alert on a metric grouping and use this option—you only need to set up one Simple Alert, and you can receive individual Simple Alerts for each individual metric in the metric grouping.

CA Technologies Prerelease Documentation. Authorized Use Only.

Notifications are sent as if there were separate Simple Alerts for each metric, so be aware that multiple Alerts/Resolutions are possible in the same period.

Trigger Alert Notification

Choose an option from the drop-down list:

- Each Period While Problem Exists—produces a problem message every period that the Simple Alert is in Caution or Danger.
- When Severity Increases—produces a problem message on any period when the state of the Simple Alert escalates from Normal to Caution, from Normal to Danger, or from Caution to Danger. This is the default state for the Simple Alert.
- Whenever Severity Changes (Resolution alert)—produces problem and/or resolution messages on any state transition. For example, a change of state of the Simple Alert from Danger to Caution would produce a resolution message (Danger status has been resolved) and a problem message (the Caution status is still a problem). This type of resolution alert produces a resolution message if the state changes from Caution or Danger.
- Report Only Final State Whenever Severity Changes (Resolution alert)—produces a problem or resolution message only for the final state of an alert transition. For example, for a change from Danger to Caution, the Simple Alert would trigger only a problem message for the final state, which is Caution. This type of resolution alert produces a resolution message only if the state goes to Normal.

More information:

About Alert Notification options, messages, and exceptions (see page 322)

Danger Threshold

Danger Thresholds specify when Simple Alerts are to be triggered.

You set Danger Thresholds in conjunction with the Comparison Operator.

To set the threshold for a Danger alert:

1. In the Threshold field, enter the value that triggers a Danger alert.

   The units in the Danger Threshold values correspond to the value used in the metric grouping. For example, if you are making a Simple Alert for Servlet Average Response Time, the value is milliseconds.

2. Set the ratio of excessive periods that must be met for the alert to be triggered.

   To do this, you put one value in the Periods Over Threshold field and another in the Observed periods field. For example, if you enter 8 and 10, then the danger alert will be triggered only if the metric exceeds the danger threshold in 8 of the 10 observed periods.

CA Technologies Preferred Documentation. Authorized Use Only.

Note: If you want to change an existing threshold setting: When you edit the threshold of an active alert to change the threshold to a number which is lower than the current metric being reported, the alert state will change to open, meaning the alert has been triggered, with the resulting actions that you have configured for the alert. Therefore, before making such changes, you should notify alert recipients of possible coming false alerts.

Actions

Add actions as described in Activating an action (see page 316).

Action delay

Enter the delay in hours, minutes, and seconds.

Note: When using the Resolution alert option, Danger action delay is not available.

Danger action delays (also known as SmartTrigger functionality), determine when a Simple Alert action is triggered.

To prevent being flooded by Simple Alert notifications when values remain in or re-enter a danger region, enter delay for Danger action delay. The action is not repeated until the delay time has elapsed. For more information on SmartTrigger functionality, see Alerts and the SmartTrigger feature (see page 322).

Caution Threshold

Caution Thresholds specify when Simple Alerts are to be triggered.

You set Caution Thresholds in conjunction with the Comparison Operator.

To set the threshold for a Danger alert:

1. In the Threshold field, enter the value that triggers a Caution alert.

   The units in the Caution Threshold values correspond to the value used in the metric grouping. For example, if you are making a Simple Alert for Servlet Average Response Time, the value is milliseconds.

2. Set the ratio of excessive periods that must be met for the alert to be triggered.

   To do this, you put one value in the Periods Over Threshold field and another in the Observed periods field. For example, if you enter 8 and 10, then the Caution alert will be triggered only if the metric exceeds the danger threshold in 8 of the 10 observed periods.

CA Technologies Proprietary Documentation. Authorized Use Only.

Note: If you want to change an existing threshold setting: When you edit the threshold of an active alert to change the threshold to a number which is lower than the current metric being reported, the alert state will change to open, meaning the alert has been triggered, with the resulting actions that you have configured for the alert. Therefore, before making such changes, you should notify alert recipients of possible coming false alerts.

Caution action delay

Enter the delay in hours, minutes, and seconds.

Caution action delays (also known as SmartTrigger functionality), determine when a Simple Alert action is triggered.

To prevent being flooded by Simple Alert notifications when values remain in or re-enter a caution region, enter delay for Caution action delay. The action is not repeated until the delay time has elapsed. For more information on SmartTrigger functionality, see Alerts and the SmartTrigger feature (see page 322).

Note: When using the Resolution alert option, Caution action delay is not available.

Adding actions

Add actions to occur when the Alert Comparison condition is met (when a caution or danger threshold value has been exceeded). You can add an action for either or both of the Danger or Caution conditions. You can also create multiple actions for the same condition.

Note: If you define both a Caution and a Danger Threshold action for a Simple Alert, and a Simple Alert status goes directly from normal (green) to danger (red) during a defined time period, only Danger actions are triggered.

To add an action:

1. Under either Danger actions or Caution actions, click Add.
2. Select an action and click Choose.
3. Add another action if appropriate.
4. In the Simple Alert settings pane, click Apply (in the bottom left corner).

   The Simple Alert is complete, and appears in the tree under the Management Module you placed it in.

Note: When creating a Simple Alert from a metric, a metric grouping is created automatically when a Simple Alert is created. It is in the same Management Module as the Simple Alert you just created.

CA Technologies Prerelease Documentation. Authorized Use Only.

Activating an action

The default actions included with Introscope (and any new actions created from this dialog) must be activated before first use.

To active an action:

1. In the Management Module Editor tree, find and select the action you just defined (or created) for the Simple Alert. Notice it is dimmed (faded and not colored) because it is not yet active.



2. In the action's settings pane, check the Active checkbox to activate the action.
3. Click Apply.

About Summary Alerts

A Summary Alert provides a way to show the status of multiple underlying Simple Alerts with one overall status.

A Simple Alert has one of four states—not reporting, green, yellow, and red. Because the Summary Alert state is defined as the worst state among the Simple Alerts it contains, Summary Alerts have no explicit danger or caution thresholds or comparison expressions, as with Simple Alerts.

To understand the relationship between the states of the underlying Simple Alerts and the containing Summary Alert, consider an example Summary Alert, which includes two Simple Alerts—GC Heap Alert, and Connection Pool Alert.

The following table defines alert states:

State Icon	Definition	Numeric Value
Red octagon	Danger	3
Yellow diamond	Caution	2
Green disc	Normal	1
Gray disc	Not Reporting	0
Red with a black circle in the middle	Danger, but in a downtime period, and still be reported.	-3
Yellow with a black circle in the middle	Caution, but in a downtime period, and still be reported.	-2

CA Technologies Pre-Release Documentation. AT&T only. Internal Use Only.

Green with a black circle in the middle	Normal, but in a downtime period, and still be reported.	-1

Summary Alerts and time periods of underlying Simple Alerts

Summary Alerts do not have a user-defined time period for checking the current state of each underlying Simple Alert. The Summary Alert period is automatically configured as the minimum period of any included Simple Alerts. There are two possible cases:

- All Alerts have the same time period—There is no delay between the evaluations of any of the Alerts on the latest metric data.
- All Alerts have different time periods—The Summary Alert still evaluates its own state on its own period, using the most recently calculated state from each of the Alerts it depends on. For example, if Summary Alert A depended upon Alert X (30 second period) and Alert Y (45 second period), then every 30 seconds, Summary Alert A would determine its state based upon the current states of Alert X and Alert Y using whatever states they had most recently calculated for their own periods.

  If the underlying Simple Alerts have different time periods, the minimum period is used as the Summary Alert time period. This works well if the underlying Simple Alerts have time periods that are relatively close together. However, with underlying Simple Alerts with a wide variance of time periods, this can lead to old or stale states in which the user might expect the Summary Alert to be Green, but because of a Simple Alert with a long period, it might have a state reflective of a time far back in the past.

  For example, if a Summary Alert (Application Health) depends upon a Simple Alert (WebServerSlow) which has a time period of one hour, then the Summary Alert could show a state that is triggered by the WebServerSlow state, which could be as old as an hour. If the WebServer was slow one hour ago, but corrected itself 50 minutes ago, the real application state might be apparent as Green/OK, but the WebServerSlow Simple Alert might still be Red, and by extension the Summary Alert, Application Health, is still red.

  The best way to prevent this situation from occurring is to compose Summary Alerts of Simple Alerts whose time periods are the same or at least close together.

Summary Alert notes

Note these caveats about Summary Alerts:

- Summary Alerts can contain Simple Alerts and other Summary Alerts
- Summary Alerts are only viewable with the alert status indicator data viewer

CA Technologies Confidential/Proprietary Information. Authorized Use Only.

- Including a Simple Alert in a Summary Alert does not disable any notification actions the Simple Alert might have. If a notification action is defined at both the level of the Simple Alert and the level of the Summary Alert, it is possible to get multiple notifications with duplicate information for the same problem. Therefore, you might want to disable the actions in the Simple Alert if you don't want duplicate notifications for the same problem.

Summary Alert notifications

Summary Alert notifications do not include any metric data; they include this information:

- timestamps
- the name of the Summary Alert
- the state the Summary Alert is in
- a list of the underlying Simple Alerts that triggered the current state of the Summary Alert

This is an example of the format of a Summary Alert notification message:

```
4/13/04 12:31:45 PM PST
The Summary Alert "Application Health" is in the danger state due to:
SuperDomain/<Acme>|<SimpleAlertName1> is in danger
SuperDomain/<Acme>|<SimpleAlertName2> is in caution
SuperDomain/<Acme>|<SimpleAlertName3> is normal
SuperDomain/<Acme>|<SimpleAlertName4> is not reporting
```

Creating a Summary Alert

To create a Summary Alert:

1. In the Management Module Editor window, select Elements > New Alert > New Summary Alert.
2. In the Name field, enter a name for the Summary Alert.

   Note: Since Summary Alerts and Simple Alerts appear together under the Alerts node, it might be helpful to name the Summary Alert to distinguish it from a Simple Alert.
3. Choose a Management Module to contain the Summary Alert in one of these ways:
   - Select a Management Module from the drop-down list box.
   - Click Choose, select a Management Module from the list, then click Choose again.
4. Click OK.

   The Summary Alert you created is highlighted in the Management Module Editor tree, and appears in the settings pane.

CA Technologies Proprietary Use Only. Authorized Use Only. Documentation.

5. In the settings pane, check the Active checkbox to activate the Summary Alert.
6. Specify the alerts to be included in the Summary alert, by selecting one or more alerts and using the arrows to move them from the Available list to the Included list:

Arrow key	Action
>	moves a single selected alert into the Included list
>>	moves all sample alerts into the Included list
<	removes a single selected alert from the Included list
<<	removes all sample alerts from the Included list

   Both Simple and Summary alerts appear in the list of available alerts. Basing a Summary alert on other Summary alerts enables you to build high-level alerts. For example, you can create a high-level summary alert that incorporates system health alerts into one overall system health alert.

   Note: Do not define two Summary alerts that are inputs for one another. The resulting recursive effect produces unpredictable results.
7. Configure the Summary Alert settings.

Trigger alert notification

Select a Trigger Alert Notification state for Any Alert or All Alerts. The Any Alert option will take the maximum state of all the alerts and the All Alerts option will take the minimum state of all the alerts having a state of more than 0 (not reporting).

Example: You have a summary alert that consists of the following alerts:

- Alert: A; State: 0
- Alert: B; State: 1
- Alert: C; State: 1
- Alert: D; State: 2
- Alert: E; State: 3

In this situation, the Any Alert option will take 3, and the All Alerts option will take 1.

The trigger state determines how the Summary Alert behaves:

- **Each Period While Problem Exists**—produces a problem message every period that the Summary Alert is in Caution or Danger.
- **When Severity Increases**—produces a problem message on any period when the state of the Summary Alert escalates from Normal to Caution, from Normal to Danger, or from Caution to Danger. This is the default state for the Summary Alert.

CA Technologies Prerelease Documentation. Authorized Use Only.

- **Whenever Severity Changes (Resolution Alert)**—produces problem and/or resolution messages on any state transition. For example, a change of state of the Summary Alert from Danger to Caution would produce a resolution message (Danger status has been resolved) and a problem message (the Caution status is still a problem). This type of resolution Alert produces a resolution message if the state changes from Caution or Danger.
- **Report Only Final State Whenever Severity Changes (Resolution Alert)**—produces a problem or resolution message only for the final state of an Alert transition. For example, for a change from Danger to Caution, the Summary Alert would trigger only a problem message for the final state, which is Caution. This type of resolution Alert produces a resolution message only if the state goes to Normal.

Resolution Alert Information

The Resolution Alert can be set to notify you when:

- a Summary Alert status changes to Caution or Danger
- a Summary Alert status changes from Caution or Danger

For more alert notification information, see About Alert Notification options, messages, and exceptions (see page 322).

Adding an action

Add actions to occur when the aggregate Summary Alert status is Danger or Caution. You can add an action for either or both of the Danger or Caution conditions. You can also create multiple actions for the same condition.

Note: If you define both a Caution and a Danger Threshold action for an Alert, and an Alert status goes directly from normal (green) to danger (red) during a defined time period, only Danger actions are triggered.

To add an action:

1. Under either the Danger actions or Caution actions, click Add.
2. Select an action and click Choose.
3. Add another action if needed.
4. In the Summary Alert settings pane, click Apply (in the bottom left corner).

   The Summary Alert is complete, and appears in the tree under the Management Module in which you placed it.

CA Technologies Prerelease Documentation
Authorized Use Only

Activating actions

The default actions included with CA APM (and any new actions created from this dialog) must be activated before they are used for the first time.

To activate an action:

1. In the Investigator tree, select the action you just defined (or created) for the Alert.

   It is dimmed, because it is not yet active.
2. In the action's settings pane, check the Active checkbox to activate the action.
3. Click Apply.

Action delay

You can enter the delay in hours, minutes, and seconds.

Danger action delays (also known as SmartTrigger functionality) determine when a Summary Alert action is triggered.

To prevent being flooded by Summary Alert notifications when values remain in or re-enter a danger region, enter delay for Danger action delay. The action is not repeated until the delay time has elapsed. For more information on SmartTrigger functionality, see Alerts and the SmartTrigger feature (see page 322).

When you use the Resolution Alert option, Danger action delay is not available.

Caution action delay

You can enter the delay in hours, minutes, and seconds.

Caution action delays (also known as SmartTrigger functionality) determine when a Summary Alert action is triggered.

To prevent being flooded by Summary Alert notifications when values remain in or re-enter a caution region, enter delay for Caution action delay. The action is not repeated until the delay time has elapsed. For more information on SmartTrigger functionality, see Alerts and the SmartTrigger feature (see page 322).

Note: When you use the Resolution Alert option, Caution action delay is not available.

CA Technologies Product Documentation. Authorized Use Only.

About Alert Notification options, messages, and exceptions

Alert Notification options determine when a Simple Alert or a Summary Alert notification is triggered and what type of information message is generated by Introscope.

A Simple Alert or a Summary Alert status generates two types of information messages—a problem message and a resolution message.

This information can be output (through a Shell Command Action, for example) to an external enterprise control panel (such as CA Unicenter). The four Alert notification options produce a combination of these messages under different conditions.

Note: An action must be defined for the Simple Alert or Summary Alert in order to output the problem/resolution message information.

These exceptions affect when a resolution alert is generated:

- When you configure a Simple Alert or a Summary Alert to be a resolution Alert, the resolution Alert behavior does not take effect until the next Summary Alert time period.
- If you are editing a Simple Alert or a Summary Alert, resolution Alert notifications based on the new information are not generated until you click Apply.
- If you shut down the Enterprise Manager, Resolution Alert notifications are not generated.
- Resolution Alert notifications are not generated for Simple Alerts or Summary Alerts on metrics/agents that disconnect or stop reporting.

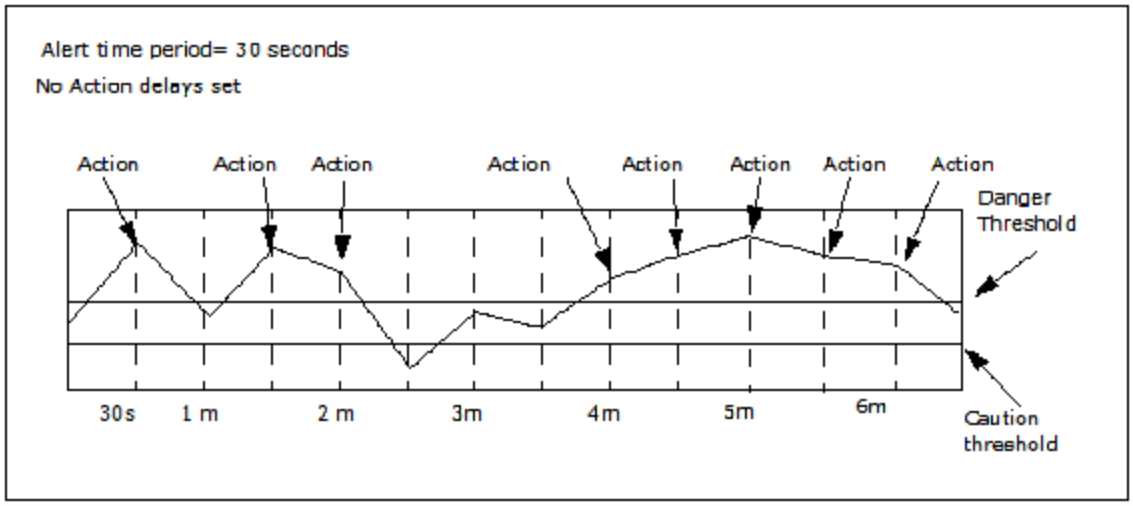
Alerts and the SmartTrigger feature

SmartTrigger functionality (through Danger and Caution action delays) determines the conditions under which danger or caution statuses are reported by a Comparison result in an action. SmartTrigger functionality prevents you from being flooded by Alert notifications. It acts like a snooze button for Alert notifications, enabling you to set a delay between the first Alert notification and subsequent notifications.

CA Technologies Pre-release Documentation. Authorized Use Only

Using SmartTrigger to delay actions

Imagine a situation where you have your Alert time period set for 30 seconds. If the information generates a Danger alert status and you defined an action for it, the action is triggered. Without SmartTrigger set, if the Danger status continues, you are notified every time the Danger threshold is exceeded, as shown in this illustration.

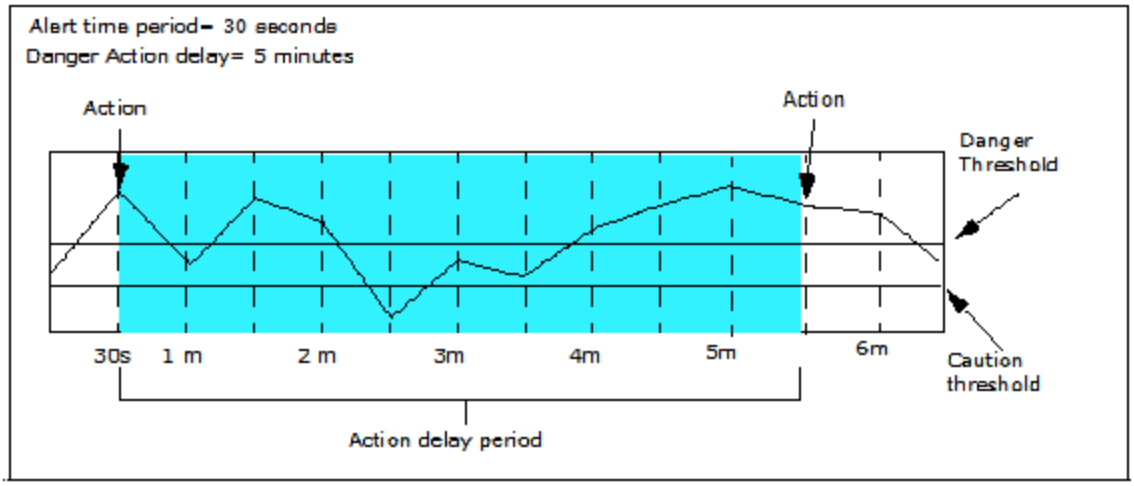


As the illustration shows, you would be notified eight times over this short period.

Because problems usually cannot be resolved in such a short time period as 30 seconds, it makes sense to delay the subsequent actions with an action delay. For example, with the same 30 second Alert time period, if you set a five minute action delay for the Danger status, you receive the first Alert notification at the 30-second mark as usual.

CA Technologies Internal / Authorized Use

However, if the Danger status occurs again during that five-minute action blackout period, and the Danger threshold is still exceeded when the blackout period ends, you are not notified by a second action until five minutes after the first notification, as this illustration shows:

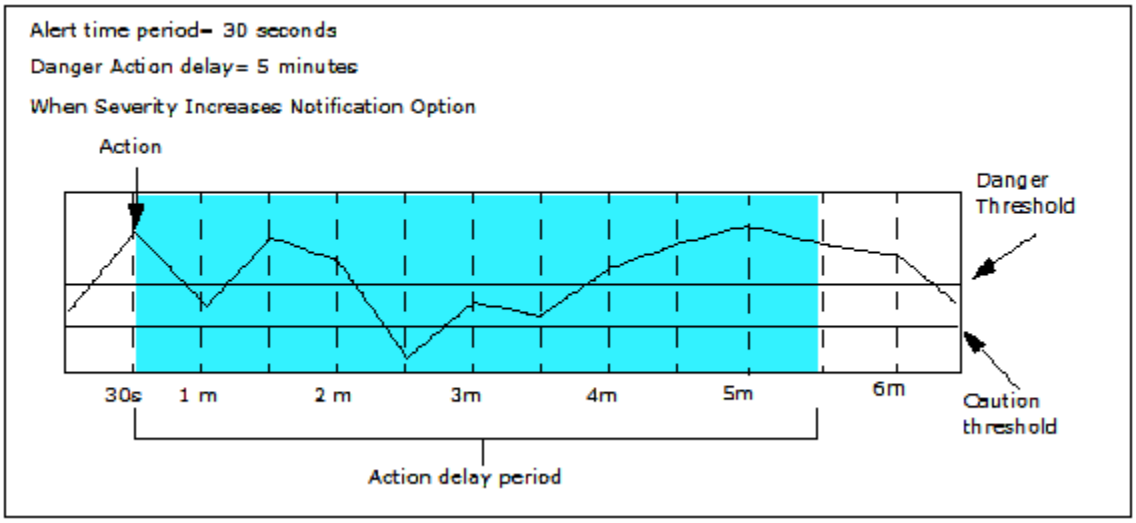


Using SmartTrigger when severity increases option

Sometimes you only want to be notified when an Alert status worsens, such as changing from normal to caution, or from caution to danger. To do this, use the When Severity Increases option in the Trigger Alert Notification field.

Let's go back to the previous example. You might only want to be notified if the status worsens to danger, not if it exceeds the danger threshold and stays there.

This illustration shows what happens with a Danger action delay of five minutes, and the When Severity Increases option selected:



In this example, you wouldn't receive a Danger Alert notification at the 5.5 minute mark as in the previous illustration, because the values are on a decline—the status is improving.

Generating alert state metrics

You can configure Introscope to create metrics that represent the trinary state of all alerts in an Enterprise Manager. This enables you to view live and historical views of alert states in the Workstation and in WebView.

Alert state metrics can be used in custom views, such as:

- a graph of alert state over time, correlated with other graphed metrics
- a chart showing the percentage of time an Alert spent in each Alert state over a period of time

When you delete or rename an Alert, the old metric for it is grayed out.

More information:

Alert state metrics in the Investigator (see page 325)
Alert state metrics in the Management Module Editor (see page 326)

Alert state metrics in the Investigator

In the Investigator, Alert state metrics appear in the virtual agent (calculator agent) for each domain defined in the Enterprise Manager, under the Alerts node. Alert state metrics in the Investigator are classified by numbers. For more information about the states, see About Summary Alerts (see page 316).

Alert-state metrics appear in the Investigator in the virtual agent (calculator agent) for each domain defined in the Enterprise Manager, under the Alerts node:

Alerts|[management module name]:[alert name]

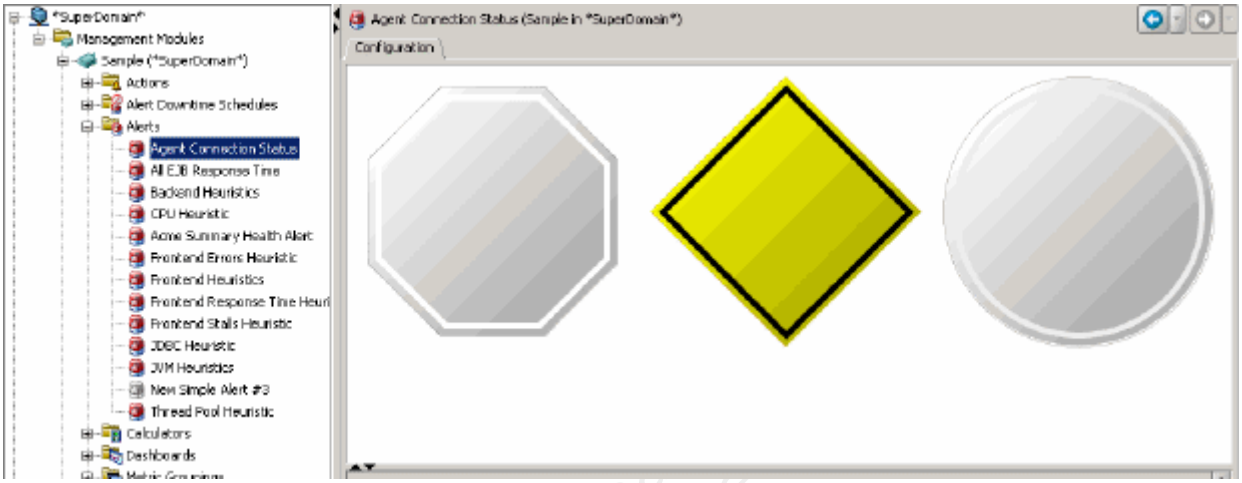
Note: You can configure a different metric name for the node that contains Alert state metrics in the Enterprise Manager properties file, using the property introscope.enterprisemanager.alertstatemetric.prefix.

CA Technologies Prerelease Documentation
Authorized Use Only

Alert state metrics in the Management Module Editor

In the Management Module Editor, Alert state metrics appear as an alert status indicator, showing the states of green (OK), yellow (Caution), and red (Danger).

This example shows a yellow alert status indicator for the Agent Connection Status metric; it is in a state of Caution:



Working with Alert Downtime Schedules

Alert Downtime Schedules let you manage downtime periods from the Management Module Editor.

An Alert Downtime Schedule can be associated with one or more alerts. It also provides a convenient way of associating alerts in one or more management modules. While any action associated to an alert will not occur during a downtime period, the alert status will still be reported. For more information, see About Summary Alerts (see page 316).

This feature is also capable of handling overlapping downtime periods that affect the same alert. For example if you have two downtimes scheduled that affect the same alert, the system is sophisticated enough to know this and maintain the down time.

For example, if there is a downtime scheduled for Monday 8 AM to 10 AM, which affects Alert A, and additionally there is another downtime scheduled for Monday 9:30 AM to 10:30 AM that affects the same Alert A, then the system will continuously maintain the downtime. Thus the system will remain down from 8 AM through 10:30 AM without any glitches.

CA Technologies Prerelease Unauthorized Use

Creating an Alert Downtime Schedule

To create a new alert downtime schedule:

1. From the Management Module Editor, choose Elements > New Alert Downtime Schedule.
2. Enter a name for the alert downtime schedule in the Name field.
3. The Force Uniqueness check box is selected by default to ensure that the Alert Downtime Schedule names are unique within a within a Management Module. If you create a new schedule with a name already in existence, the system will append a number to the name to force uniqueness. Deselect this check box to turn off this option.
4. Choose a Management Module from drop-down menu or click Choose to enter a search string.
5. Click OK.

   The new Alert Downtime Schedule will appear highlighted in the Management Module editor tree and its definitions will appear in the Settings tab in the lower editor pane.

   You can select the Description tab to enter information about the alert or the Settings tab to define the alert time downtime settings. For more information, see Defining Alert Downtime Schedules (see page 327).

Defining Alert Downtime Schedules

Define these settings after you create an alert downtime schedule (see Creating an Alert Downtime Schedule (see page 327)) or from selecting an existing alert downtime schedule from the editor tree and selecting the criteria in the Settings tab:

To define alert downtime schedule:

1. In the Name field enter or rename the existing Alert Downtime Schedule.
2. Select the Active check box to make the make the Alert Downtime Schedule active.
3. Select a Management Module from the drop-down menu or click Choose to enter a search string.
4. Select one of the following scheduling options:
   - **Simple Schedule–**allows you to schedule weekly, monthly, and daily one time or recurring alerts that can be set to start and end at a specific time.
   - **Cron Schedule–**a Unix scheduling tool that uses expressions. While this tool provides a wide range of capabilities, the values entered in the fields need to be precise. For more information, see Using Cron to schedule alert downtimes (see page 328).

CA Technologies Prerelease Documentation Authorized Use Only

5. Select the Alerts or Management Modules option to pick from a list of alerts or Management Modules to which you wish to apply the settings.

   You can apply settings only to alerts in your Management Module. This is done as a safeguard measure, so alerts aren't inadvertently deactivated.

6. Click Apply to or Revert to.

Using Cron to schedule alert downtimes

Cron is a powerful UNIX tool that provides a variety of scheduling capabilities. It uses expressions that can trigger alerts that apply rules such as "8:00am every Monday through Friday" or "1:30am every last Friday of the month".

Use it by selecting the Cron Schedule option when you define alert downtime schedules (For more information, see Defining Alert Downtime Schedules (see page 327)):

Cron expressions can be as simple as: ** * * * * ? **

or, a complex as: *0 0/5 14,18,3-39,52 ? JAN,MAR,SEP MON-FRI 2002-2010*

For more sample expressions, see Cron Sample Expressions (see page 330).

The following table lists the values and special characters that are allowed in the Cron Schedule fields:

Field Name	Required	Values	Special Characters
Minutes	yes	0-59	, - * /
Hours	yes	0-23	, - * /
Day of Month	yes	1-31	, - * ? / L W C
Month	yes	1-12 or JAN-DEC	, - * /
Day of Week	yes	1-7 or SUN-SAT	, - * ? / L C #
Year	no	empty, 1970-2099	, - * /

Cron Special Characters

The following table lists the Cron special characters and their meanings.

Special Character	Definition
* (all values)	Selects all values within a field. For example, "*" in the minute field means "every minute".

CA Technologies Prerelease Documentation. Authorized Use Only

? (no specific value)	Specifies something in one of the two fields in which the character is allowed, but not the other. For example, if I want my trigger to fire on a particular day of the month (say, the 10th), but don't care what day of the week that happens to be, I would put "10" in the day-of-month field, and "?" in the day-of-week field. See the examples below for clarification.
-	Specifies ranges. For example, "10-12" in the hour field means "the hours 10, 11 and 12"
,	Specifies additional values. For example, "MON, WED, FRI" in the day-of-week field means "the days Monday, Wednesday, and Friday".
/	Specifies increments. For example, "0/15" in the seconds field means "the seconds 0, 15, 30, and 45". And "5/15" in the seconds field means "the seconds 5, 20, 35, and 50". You can also specify '/' after the '*' character - in this case '*' is equivalent to having '0' before the '/'. '1/3' in the day-of-month field means "fire every 3 days starting on the first day of the month".
L (last)	Specifies the last of something. This special character has a different meaning in the two fields for which it is allowed: Day of Month and Day of Week. For example, the if you insert an L in the Day of Month field, it means the last day of the month, which would be day 31 for January, day 28 for February on non-leap years. If used in the day-of-week field by itself, it simply means "7" or "SAT". But if used in the day-of-week field after another value, it means "the last xxx day of the month" - for example "6L" means "the last friday of the month". When using the 'L' option, it is important not to specify lists, or ranges of values, as you'll get confusing results. The 'L' and 'W' characters can also be combined in the day-of-month field to yield 'LW', which translates to "last weekday of the month".
W (weekday)	Specifies the weekday (Monday-Friday) nearest the given day. As an example, if you were to specify "15W" as the value for the day-of-month field, the meaning is: "the nearest weekday to the 15th of the month". So if the 15th is a Saturday, the trigger will fire on Friday the 14th. If the 15th is a Sunday, the trigger will fire on Monday the 16th. If the 15th is a Tuesday, then it will fire on Tuesday the 15th. However if you specify "1W" as the value for day-of-month, and the 1st is a Saturday, the trigger will fire on Monday the 3rd, as it will not 'jump' over the boundary of a month's days. The 'W' character can only be specified when the day-of-month is a single day, not a range or list of days.

CA Technologies Proprietary and Confidential. Authorized Use Only. Documentation.

#	Specifies "the *nth*" *XXX* day of the month. For example, the value of "6#3" in the day-of-week field means "the third Friday of the month" (day 6 = Friday and "#3" = the 3rd one in the month). Other examples: "2#1" = the first Monday of the month and "4#5" = the fifth Wednesday of the month. Note that if you specify "#5" and there is not 5 of the given day-of-week in the month, then no firing will occur that month.
c (calendar)	This means values are calculated against the associated calendar, if any. If no calendar is associated, then it is equivalent to having an all-inclusive calendar. A value of "5C" in the day-of-month field means "the first day included by the calendar on or after the 5th". A value of "1C" in the day-of-week field means "the first day included by the calendar on or after Sunday".

Cron Sample Expressions

The following table lists the Cron special characters and their meanings.

Expression	Meaning
0 0 12 * * ?	Trigger an alert at 12pm (noon) every day
0 15 10 ? * *	Trigger an alert at 10:15am every day
0 15 10 * * ?	Trigger an alert at 10:15am every day
0 15 10 * * ?	Trigger an alert at 10:15am every day
0 15 10 * * ? 2005	Trigger an alert at 10:15am every day during the year 2005
0 * 14 * * ?	Trigger an alert every minute starting at 2pm and ending at 2:59pm, every day
0 0/5 14 * * ?	Trigger an alert every 5 minutes starting at 2pm and ending at 2:55pm, every day
0 0/5 14,18 * * ?	Trigger an alert every 5 minutes starting at 2pm and ending at 2:55pm, AND fire every 5 minutes starting at 6pm and ending at 6:55pm, every day
0 0-5 14 * * ?	Trigger an alert every minute starting at 2pm and ending at 2:05pm, every day
0 10,44 14 ? 3 WED	Trigger an alert at 2:10pm and at 2:44pm every Wednesday in the month of March.
0 15 10 ? * MON-FRI	Trigger an alert at 10:15am every Monday, Tuesday, Wednesday, Thursday and Friday
0 15 10 15 * ?	Trigger an alert at 10:15am on the 15th day of every month
0 15 10 L * ?	Trigger an alert at 10:15am on the last day of every month

CA Technologies Proprietary Document. Authorized Use Only.

0 15 10 ? * 6L	Trigger an alert at 10:15am on the last Friday of every month
0 15 10 ? * 6L	Trigger an alert at 10:15am on the last Friday of every month
0 15 10 ? * 6L 2002-2005	Trigger an alert at 10:15am on every last friday of every month during the years 2002, 2003, 2004 and 2005
0 15 10 ? * 6#3	Trigger an alert at 10:15am on the third Friday of every month
0 0 12 1/5 * ?	Trigger an alert at 12pm (noon) every 5 days every month, starting on the first day of the month.
0 11 11 11 11 ?	Trigger an alert every November 11th at 11:11am.

Downtime Schedules for Triage Map Alerts

You can use the Management Module editor to set downtime schedules for Triage Map alerts. A radio button on the Alert Downtime Schedule Configuration tab allows you to select only Triage Map alerts for inclusion in a schedule.

To set downtime schedules for Triage Map alerts:

1. Open the Triage Map Configurations Management Module.
2. Click the Alert Downtime Schedules node.
3. Configure a downtime schedule, using the steps in Creating an Alert Downtime Schedule (see page 327) and Defining Alert Downtime Schedules (see page 327).
4. Select the Triage Map Alerts radio button.

   Selecting this radio button causes the usual three-column “Available Alerts” table to be replaced by a single-column table of available Triage Map Alerts.
5. Select one or more Triage Map Alerts to apply the schedule to, and click right-arrow button to shift them into the Included Alerts list.
6. Click Apply.

Creating actions and notifications

An action is caused by an Alert, and defines what happens when an Alert is triggered.

Introscope provides three standard action types:

- A *Workstation notification action* displays an alert notification on all running Workstations connected to the Enterprise Manager—see Creating a Workstation Notification Action (see page 332).

CA Technologies Internal / Confidential Documentation. Authorized Use Only

- A *shell command action* runs a shell script residing on the Enterprise Manager machine—see Creating a Shell Command Action (see page 332).
- An *SMTP email action* sends an e-mail to the recipient(s) specified in its settings—see Creating an SMTP Email Action (see page 333).

Introscope includes two default actions in the Sample Management Module: SMTP Email Notification and Workstation Notification. They must be configured and activated to be used.

Creating a Workstation Notification Action

The Workstation action type displays an Alert notification on all running Workstations connected to the Enterprise Manager. This is the simplest action type, because it doesn't require any other systems or setup to work.

To create a new Workstation Notification Action:

1. In the Management Module Editor, select Elements > New Action > New Workstation Notification Action.
2. Name the action and select a Management Module to contain it.
3. Check the Active box to activate the action.
4. Click Apply.
5. Click Test Now to see the result of the action (the action must be active).

   This example shows a test of a Workstation Notification Action.

Note: The Test Now button produces a test result only for the last applied action.

Creating a Shell Command Action

The Shell Command Action type runs a shell script residing on the Enterprise Manager machine. The action can pass a short text message to the script, describing the reason why the Alert was triggered—for example:

```
4:05:15 PM PST Introscope Enterprise Manager (aardvark:192.168.0.45) reported: The Alert My App Heap Bytes In Use Alert was triggered because the value 841904 exceeded danger target of 800000 for Acme|c1737019-a|AcmeUSA|AcmeWest|GC Heap:Bytes In Use"
```

To create a new Shell Command Action:

1. In the Management Module Editor, select Elements > New Action > New Shell Command Action.
2. Name the action and select a Management Module to contain it.
3. Select the Force Uniqueness check box to make the name unique within a Management Module.

CA Technologies Prerelease Documentation. Authorized Use Only.

4. Choose a Management Module from the drop-down menu or click Choose to enter a search string and narrow the list of options.
5. Click OK.

   The Shell Command Action Settings pane appears.
6. Choose a Management Module to contain the Shell Command Action:
   - Select a Management Module from the drop-down list box, or
   - Click Choose, select a Management Module from the list, then click Choose again.
7. Check the Active checkbox to activate the action.
8. Enter the name of the shell command in the Shell Command field.
9. Enter an optional parameter that needs to be passed to the shell command In the User-defined parameter field.
10. Click Test Now to see the result of the action.

    Note: The Test Now button produces a test result only for the last applied action.
11. Select the Command Parameters option to pick from a list of command parameters which you want to be included during shell script execution.

    The command parameters are added by selecting the parameter from the Available Command Parameters list and clicking the > button. All the command parameters can be added by clicking the >> button.The command parameters are removed by selecting the command parameters from the Included Command Parameters list and clicking the < button. All the command parameters can be removed by clicking the << button.
12. Click Apply to apply the changes or Revert to revert back to the original values.

Creating an SMTP Email Action

The SMTP Email Action type sends an e-mail to the recipient(s) specified in its settings. This action type requires access to an SMTP e-mail server that the Enterprise Manger can connect to.

This action type could send an e-mail to several places:

- a regular e-mail address
- an already-defined mail list
- a pager gateway, which can trigger a person's pager
- a management system that can take the e-mail text as input and trigger an action

CA Technologies Proprietary Documentation
Authorized Use Only

To create a new SMTP Email Action:

1. In the Management Module Editor, select Elements > New Action > New SMTP Email Action.
2. Name the action and select a Management Module to contain it.
3. Check the Active checkbox to activate the action.
4. Enter the return address of the e-mail message in the From: field.
5. Enter the recipient name in the To: field (can be a single e-mail address, or multiple addresses separated by commas)
6. Enter the name of the SMTP host in the SMTP Host field.
7. Select Send Short Message to send an abbreviated version of the notification message, for bandwidth-sensitive channels such as a pager.
8. Click Apply.
9. Click Test Now to see the result of the action.

   Note: The Test Now button produces a test result only for the last applied action.

About the SNMP Alert Action Plugin

The SNMP Alert Action Plugin allows the APM Catalyst Connector to get Introscope alert data and supply it to other CA applications.

Note: Both simple and summary alerts are imported, and are displayed as triage map alerts. In the case of legacy alerts, which are Introscope alerts created through the Management Module editor and not through the triage map, summary alerts will not be imported.

SNMP Plugin Configuration

Plugin configuration settings specify which alerts forward data to the APM Catalyst Connector. You configure these settings in Introscope Management Module objects.

To configure Management Module objects so the SNMP Alert Action Plugin sends Introscope data to APM Catalyst Connector, you:

1. Create one or more alerts.

   Your SNMP Alert Action requires at least one alert to reference.
2. Create a new SNMP Alert action which references the alerts you created.

CA Technologies Pre-release Documentation. Authorized Use Only

Creating and configuring alerts

Your SNMP Alert Action requires at least one alert to reference. An Introscope alert is simply a holder for Caution and Danger threshold settings. For the procedure on creating an alert for this purpose, see the topic Creating Simple Alerts in the CA APM Workstation User Guide, part of the CA APM documentation.

Important: When creating an alert for your SNMP Alert Action to reference, make sure **Notify by individual metric** is checked. If you do not check this option, false or misleading alerts will be raised in the CA Service Operations Insight Manager.

Once you have created an alert, you can create an SNMP Alert action which references it.

To configure the trap action for an alert:

1. Identify the Management Module which is the source of the Introscope alerts you want to transform to Alert CIs.
2. Create a new SNMP Alert action:
   a. From the Elements menu, select Elements > New Action > New SNMP Alert Action
   b. Type a name for the new action.
   c. Confirm that the correct Management Module is shown. If not, select the correct one from the dropdown.
   d. Check the Active check box.
3. Configure the following information in the **SNMP Destination** section:

   Host IP

   The IP address of the host machine where APM Catalyst Connector is installed.

   NOTE: Only IPv4 is currently supported.

   Trap Port

   The SNMP Trap port configured on the APM connector host machine. Default is 162.

   Community

   The SNMP community string defines the relationship between an SNMP server system and the client systems. This string acts like a password to control client access to the server.

   This should be set to the same value as the *EMSNMPCommunity* property in the CA Service Operations Insight configuration for CA APM. Configure the Catalyst Connector for APM

CA Technologies Proprietary Documentation
Authorized Use Only.

4. Configure the following information in the **Introscope WebView** section:

   Protocol

   Connection protocol. Select one from the dropdown:

   - http
   - https

   Host IP

   The IP address of the host machine where Enterprise Manager is installed.

   NOTE: Only IPv4 is currently supported.

   Port

   WebView port number. The default is 8080.

   Management Module

   Name of the Management Module where the action resides.

   Dashboard Name

   Name of the Dashboard where the alert appears.

Click Apply.

Using Calculators

Calculators take the values from a metric grouping as input, average or sum the values, and output the resulting value as a custom metric in the Investigator tree. Metrics generated by calculators appear under a virtual process, Custom Metric Process, running on a virtual host, Custom Metric Host.

Creating Calculators

To create a calculator:

1. In the Management Module Editor, select Elements > New Calculator.
2. Name the calculator, and choose a Management Module to contain the calculator:
   - Select a Management Module from the drop-down list box.
   - Click Choose, select a Management Module from the list, then click Choose again.

CA Technologies Prerelease Documentation. Authorized Use Only.

3. Click OK.

   The calculator you created is highlighted in the Management Module Editor tree, with its settings shown in the settings pane.

   You must specify a metric grouping to supply data to the calculator. When the calculator was created, a metric grouping was automatically created with the same name as the calculator. However, the metric grouping must be customized before it can supply data to the calculator—see Configuring metric groupings (see page 273).

4. To choose another metric grouping:

   - Select a metric grouping from the drop-down list box.
   - Click Choose, select a metric grouping from the list, then click Choose again.

   Note: Choose a metric grouping that provides integer values—calculators cannot accept non-integer values as input. Mixed types produce unexpected results.

5. From the Operation menu, choose average or sum to determine the action to perform on the input from the metric grouping.

6. If you are creating a Sum calculator, from the Metric Type menu select the metric type for the calculator, either counter or interval counter. Use an interval counter when the calculator is to create a sum of interval counts; otherwise, use counter.

7. In the Destination field, specify a name for the metric to label the calculator's output:

   - To have the metric appear in a Resource folder instead of directly under Custom Metrics, specify the Resource name, followed by a colon (:), followed by the metric name.
   - To specify a chain of nested Resource folders, separate the Resource names by pipe symbols (|). You do not need to use an escape character with these pipe symbols.

     name of resource|name of resource:name of metric

8. Click Apply.

Tip: As you add calculators, you can use the "Calculators:Total Number of Evaluated Metrics" metric, which appears in the Investigator tree under Enterprise Manager | Internal | Calculators, to verify that metrics produced by new calculators don't cause your Manager of Managers (MOM) to exceed the maximum number of metric subscriptions, as suggested in the *APM Sizing and Performance Guide.* Note that you will not see this metric in the Investigator's Browse tree until a calculator is created.

CA Technologies Pre-release Documentation. Authorized Use Only.

Calculators and weighted averages

Introscope calculators that produce metrics based on averages are based on weighted averages, not straight averages. This is especially useful when monitoring the performance of your application in a clustered environment, so you have an accurate response time between multiple servers that are likely to have different load levels.

For example, if you have a calculator that generates a metric from the average response time for five servlets, a straight average would add up the response time for a defined time period, and divide by five. A weighted average would give more weight to the servlets that were called more often, giving a more accurate average.

Changing operation types in Management Module calculators

When you edit a calculator in a Management Module, changing the operation type (for example from MIN to MAX) redefines the meaning of calculator's output metric. If you keep the calculator output metric name the same, viewing this metric will juxtapose the old values in history (calculated, for example, by the MIN) with the new values (for example, the MAX), with no indication as where the alteration in processing occurred. If you are concerned that users might get confused by this, simply rename the output metric of the calculator when you change the operation type.

Using JavaScript calculators

A JavaScript calculator reads input metrics and produces output metrics according to calculations specified in a user-created JavaScript text file. The new calculated metrics can appear in the Investigator tree under the Virtual Custom Agent, or in any node of the Investigator tree, according to the output metric that is specified in the calculator script.A calculated metric can be shut off, but the calculator producing it does not know about the shutoff state.

The Enterprise Manager JavaScript engine allows you hot-deploy JavaScript calculators to a running Enterprise Manager.

Writing JavaScript calculators

JavaScript calculator files must end with a .js extension and reside in the Enterprise Manager's scripts directory. Sample JavaScript calculator files are provided in the *examples/scripts/* directory of your Enterprise Manager installation.

JavaScript calculators specify input metrics and produce one or more output metrics.

CA Technologies Pre-release Documentation
Unauthorized Use Prohibited

The execute( ) function

Each calculator must have an execute() function that takes two arguments. Additionally, helper functions are available to help build metrics to send back to the Enterprise Manager. The syntax is:

```
function execute(metricData,javascriptResultSetHelper)
```

where

- metricData—is an array of metric Data supplied to the function when it is called at every 15 seconds before execute() intervals
- javascriptResultSetHelper—is an object that collects the new metric data produced by the script and sends them back to the EM
  - kDefaultFrequency - is used as input to the frequency argument of the addMetric() helper function
  - kIntegerConstant - maps to the integer constant metric type
  - kIntegerFluctuatingCounter - maps to the integer fluctuating counter metric type
  - kLongConstant - maps to the long constant metric type
  - kLongFluctuatingCounter - maps to the long fluctuating counter metric type
  - kLongTimestamp - maps to the long timestamp metric type
  - kLongTimestampConstant - maps to the long timestamp constant metric type
  - kIntegerPercentage - maps to the integer percent metric type
  - kIntegerDuration - maps to the integer duration metric type
  - kLongDuration - maps to the long duration metric type \
  - kLongIntervalCounter - maps to the long interval counter metric type
  - kStringIndividualEvents - maps to the string metric type
  - addMetric(metricName, count, value, min, max, metricType, frequency) - supports setting the count/value/min/max of a metric value, which is needed for the rate and interval count metric types, where the "value" of the metric is based on its "count"
  - getCustomMetricAgentMetric(agentMetric) - helps build a fully qualified metric name using the agent metric supplied and filling in the rest based on the name of the SuperDomain custom metric agent

The execute() function is called every 15 seconds by the scripting engine.

CA Technologies Prerelease Documentation
Authorized Use Only

Specifying input metrics

The calculator script can specify the input metrics that it receives in one of two ways:

- The easiest way to specify input metrics is with a pair of methods: function getAgentRegex() which returns a string containing a regular expression to match agents and function getMetricRegex() which returns a string containing a regular expression to match metrics.
- You can also use the method function getMetricSpecifier() which returns a metric Specifier.

Note: The regular expressions created as strings in function getAgentRegex() and function getMetricRegex() must use character escaping differently than other regular expressions you use in Introscope—for example, in metric groupings or in the Search view. Any Java escape characters that are returned from these JavaScript functions must also be escaped in the JavaScript. So, for example, '\|' must be escaped as '\\|' in the JavaScript.

Global variable log

All JavaScript calculator functions have access to a global variable log, which is of type IModuleFeedbackChannel. For example:

```
function execute(metricData,javascriptResultSetHelper) {
log.info("message");
log.error("message");
log.debug("message");
}
```

Note: If you want to use advanced JavaScript features or are concerned with ECMA compliance, note that the script engine embeds the Mozilla Rhino JavaScript library, version 1.6_R1.

Creating output metric data

Creating output metric data requires:

- **Metric name**—consisting of the agent plus the full path to the appropriate node in the metric tree.
  - The metric name can be created based on the incoming data, in which case the new calculated data appears along with that agent's other metric data, or
  - A new calculator metric name can be specified, to show the calculated metric data in its own node in the metric tree.

CA Technologies Pre-release Documentation
Authorized Use Only.

- **Data value**—calculated by the script.
- **Result data type**—specified by a constant value from the class com.wily.introscope.spec.metric.MetricTypes.
- **Reporting Frequency**—The frequency at which the new metric data is reported to the Enterprise Manager, which can be obtained from the incoming data, or specified explicitly. You can change this to a multiple of the Enterprise Manager's default frequency (15 seconds).

A typical calculated value from a script looks like this:

```
javascriptResultSetHelper.addMetric(metricName,
heapUsedValue,Packages.com.wily.introscope.spec.metric.MetricTypes.kIntegerFluctu
atingCounter,frequency)
```

Note: Specify regular expressions with care, as they can potentially match any metrics you produce. For instance, a regular expression of "EJB.*Time.*" could insert a new value under EJB. (that is, inserting a new value under "EJB" when you have a regex on "EJB.*Time.*"). You can either change your regular expression to do this, or remove metric data from your own metrics.

Adding a JavaScript calculator

To install a new JavaScript calculator, copy the JavaScript text file into the *<EM_Home>/scripts* directory of your Enterprise Manager installation. You can use another directory for scripts; if you do, specify the directory using the introscope.enterprisemanager.javascript.dir property.

Scripts are automatically deployed from this scripts directory at the frequency specified by the introscope.enterprisemanager.javascript.refresh property, which by default is 60 seconds.

After successful deployment, the new metrics appear in the Metric Browser tree.

Running JavaScript Calculators on the MOM

You can run a JavaScript calculator on the MOM to produce metrics for the MOM's Custom Metric Agent. While it cannot produce metrics for agents that are connected to a Collector, it can see input metrics from agents in the Collectors.

If a calculator is added, modified, or deleted in a clustered environment, the MOM will automatically propagate the change to all Collectors unless you turn off the automatic update for collectors. For more information, see Turning off the automatic update for Collectors (see page 342).

CA Technologies Pre-release Documentation. Authorized Use Only.

The runOnMOM function

A JavaScript calculator that should not run on the MOM must implement a runOnMOM function that returns false, such as in the following example:

```
// return false if the script should not run on the MOM
// default is true.
function runOnMOM()
{
   return false;
}
```

If the runOnMOM function returns true or is not implemented, the JavaScript calculator will run on the MOM.

Reducing the number of logged metric creation errors

If a calculator runs on the MOM and creates a metric for an agent that exists in the Collectors, there is a one-time logging of the event at the WARN level, such as in the following example:

```
5/15/07 02:32:20 PM PDT [WARN] [Manager.MetricCalculatorBean] Calculator
Registered Metric <ID=7, JavaScript calculator
C:\workspaces\workspaceKrakatau\com.wily.introscope.em.feature\rootFilesMOM\.
\scripts\HeapUsedPercentage.js>. A JavaScript calculator in the MOM cannot output
metric data to an agent that exists in a Collector:
SuperDomain|rhart-dt1|EPAgentProcess1|EPAgent15|GC Heap:Heap Used (%) 5/15/07
02:32:20 PM PDT [WARN] [Manager.MetricCalculatorBean]
```

Subsequent events are logged only at the DEBUG level.

Turning off the automatic update for Collectors

Clustered environments are automatically set to propagate an added, modified, or deleted JavaScript calculator to all Collectors. However, you could turn this feature off if you do not want the calculators propagated.

To turn off the automatic update for Collectors:

1. Open the property file on the MOM Enterprise Manager.
2. Edit the property *introscope.enterprisemanager.javascript.hotdeploy.collectors.enable* which has the default value of *true*. Change the value to *false*.
3. Verify the change has been applied by viewing the *JavaScriptCalculatorsMOM.properties* file in the *<EM_Home>\config\internal\server\scripts* directory on the Collector.
4. Save and close your changes.
5. Restart the MOM Enterprise Manager.

CA Technologies Proprietary/Trade Secret Documentation. Authorized Use Only.

Deploying Management Modules

This section provides information about adding new or updated Management Modules to an Enterprise Manager. You can deploy Management Modules without restarting the Enterprise Manager using the Hot Deploy Service (see Using the Management Module Hot Deploy Service (see page 343)). However, it is recommended that you do not use this mechanism on production Collectors or MOMs. For more information, see Management Module Hot Deployments to Avoid (see page 344).

To deploy new or updated management modules, you place them in a deploy directory, which the Enterprise Manager polls periodically. When the Enterprise Manager detects new Management Module files in the deploy directory, it automatically deploys them to the Introscope SuperDomain.

If you have multiple Introscope domains and want to deploy Management Modules selectively, create a subdirectory for each target domain within the deploy directory. The Enterprise Manager polls any subdirectories of the deploy directory, and deploys Management Modules it finds to the domain corresponding to the subdirectory name.

When you deploy a Management Module that includes links to elements in another Management Module, you must deploy the Management Module that contains the target elements as well.

Note: By default, if a Management Module contains references to another management module which has not been deployed you will receive a startup warning notification. The warning will be logged in the IntroscopeEnterpriseManager.log file located in *<EM_home>/logs* directory and will contain detailed information to help you troubleshoot and resolve the issue. To change the default system behavior, you can update the property in the *IntroscopeEnterpriseManager.properties* file. For more information, see the *CA APM Configuration and Administration Guide.*

CA Technologies Product Documentation. Authorized Use Only

Updating deployed Management Modules

Before deploying an update to a Management Module, delete the existing Management Module before deploying the updated version of the Management Module *.jar* file.

Using the Management Module Hot Deploy Service

Use the Hot Deploy Service with certain Management Modules, but do not perform Management Module hot deployments on production Collectors or MOMs. For more information, see Management Module Hot Deployments to Avoid (see page 344).

To use the Management Module hot deployment, copy the Management Module .jar file or files into the *<EM_Home>/deploy* directory, or a domain-specific subdirectory of the deploy directory, as appropriate.

The Management Modules are deployed at the next polling interval which, by default, occurs within 60 seconds. The Management Module .jar file is:

- copied to the *config/modules* folder
- backed up in the *config/modules-backup* folder
- deleted from the *deploy* directory

The deployment is logged in the Enterprise Manager log.

Management Module Hot Deployments to Avoid

Important: Do not perform Management Module hot deployments on production Collectors or MOMs. It will lock the system and prevent the metric data from being reported. Because the hot deployment of virtual agents and Management Modules is CPU intensive, it can lock up Collectors for a couple of minutes during which metric harvest does not happen. This can happen if you change the virtual agent definitions or redeploy Management Modules in the MOM or Collector; the consequence can be that the cluster stops responding to Workstation users for extended periods. Based on this, we strongly recommend not doing Management Module hot deployments in production environments.

You may perform a hot deployment during development and when developing Management Modules. However, if you are working with a large fully loaded Enterprise Manager or a large cluster, avoid performing a Management Module hot deployment, as it is likely that the system will stop responding.

For more information about Virtual Agents, see the *CA APM Configuration and Administration Guide.*

CA Technologies Proprietary & Confidential Information.
Authorized Use Only

Appendix A: Introscope Metrics

Introscope, and its extensions and add-ons, displays application performance data collected from remote and local systems as metrics. This appendix is a guide to understanding these metrics.

This section contains the following topics:

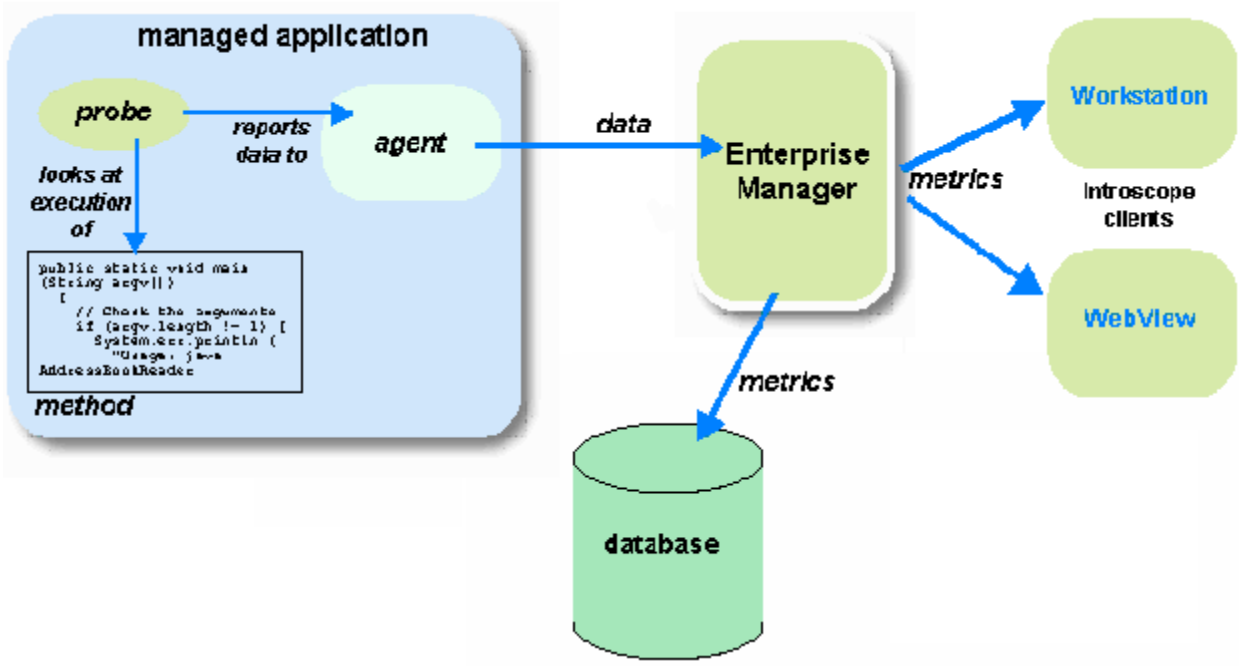
How Introscope monitors application performance (see page 346)
Viewing metrics (see page 349)
The Five Basic Introscope Metrics (see page 349)
Other common metrics (see page 358)
Other metrics (see page 376)
Data About Machines (see page 388)

CA Technologies Prerelease Documentation. Authorized Use Only.

How Introscope monitors application performance

Introscope monitors application performance by measuring the performance of individual methods as they are executed by various application components. Probes inserted into application component bytecode report data to agents, which in turn report data to the Introscope Enterprise Manager (EM). Other subsystems, like JMX and PMI, also report data which is collected by agents.

The Enterprise Manager complies this data into metrics—application performance as measured at many points in the application subsystems—and displays the metrics in Introscope Workstation or WebView. You can also export the metrics to an external database.



Common terms

To understand Introscope metrics, you should understand how Introscope uses some common terms.

A complete glossary is available in the *APM glossary.*

backend

An external system, such as a database, a mail server, a transaction processing system (such as CICS or Tuxedo), or a messaging system (such as WebSphere MQ).

concurrency and concurrent invocations

Concurrent methods are methods that started during an interval without finishing during the same interval. Since you want methods to complete quickly, an unusually high value for concurrent invocations is undesirable.

errors

Errors generated by the application or system being monitored.

frontend

The component of an application that first handles an incoming request. It may be a Servlet, a JSP, a management DB, an EJB or some other component.

harvest

The process in which Introscope gathers data from Collectors.

interval

A user-defined time slice used to define and average metrics. In Introscope this is usually 7.5 seconds, though the way some of the monitored systems capture data sometimes necessitates a different interval.

response

Response always refers to method execution. Repsonse is measured as:

- count, referring to the number of transactions finished during that interval.
- time, referring to the time it took to execute a method, in milliseconds.

Responses Per Interval is the standard Introscope throughput metric.

response time

The time it took to execute a method. May be measured as:

- average response time (ms)—The average time, in milliseconds, it took to execute the method during the interval.
- response time, min and max—The lowest and highest response times during the interval.

rate

The number of method executions per second or time interval.

stall

An instance where a method's invocation time has exceeded a threshold defined by an administrator.

CA Technologies Prerelease Documentation. Authorized Use Only.

Types of metrics

Types of metrics include:

- Count metrics (see page 348)
- Heuristic metrics (see page 348)
- Percentage metrics (see page 348)
- String data (see page 349)

Count metrics

Count is an integer. It may represent, for example:

- The number of data points which were averaged to compute a metric.
- The number of events since a certain point in time.
- The number of threads in use.

Examples of count metrics are errors and stall count.

Heuristic metrics

Heuristic metrics are used to evaluate and report status. They are integers, but the integers are symbols of status and do not measure anything. For example, a dashboard alert may be based on a heuristic metric with these values:

0 = green = normal

1 = yellow = caution

2 = red = danger

Note: These values are only examples. Your system may be configured with different values.

For more information, see:

- Heuristics and metric baselines (see page 138)
- How alerts are defined using heuristic metrics (see page 306)

Percentage metrics

Percentages are used to measure resource use against the maximum available resources. Examples are:

- CPU utilization
- Percentage of time spent in Garbage Collections during the last 15 minutes.

CA Technologies Pre-Release Documentation. Authorized Use Only.

String data

In addition to measurements and status, Introscope collects information that identifies monitored applications and systems. Examples of this type of data are system component names such as the name of a database, JVM versions, or IP address.

Viewing metrics

Introscope provides two tools to view the metrics that Introscope gathers:

- Workstation
- WebView.

To run these tools, see:

- The chapters in this manual on the Workstation Investigator (see page 30) and Console (see page 29) interfaces.
- The *Introscope WebView Guide*.

The rest of this chapter discusses the various kinds of metrics you view through these tools.

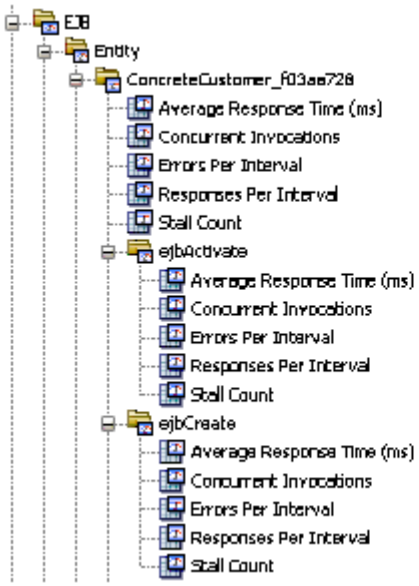
The Five Basic Introscope Metrics

Most instrumented methods report five metrics, as shown in the illustration below:

- Average Response Time (see page 351)
- Concurrent Invocations (see page 353)
- Errors Per Interval (see page 354)
- Responses Per Interval (see page 355)
- Stall Count (see page 357)

CA Technologies Prerelease Documentation. Authorized Use Only.

These are known as Blame metrics.



The illustration shows a portion of the Workstation Investigator tree—the EJB node, with an Entity subnode, and an entity bean, ConcreteCustomer_f03aa728. Under the bean we can see two of its methods, each of which reports the same five metrics. In this configuration, the metrics at the higher level (in this example, the entity bean level) *aggregate* those of the same name below.

You can see a similar arrangement under many of the nodes in the Investigator tree, particularly nodes which—like EJB—correspond to high-level J2EE APIs or their .NET equivalent.

CA Technologies Proprietary Use Only
Authorized Use Only

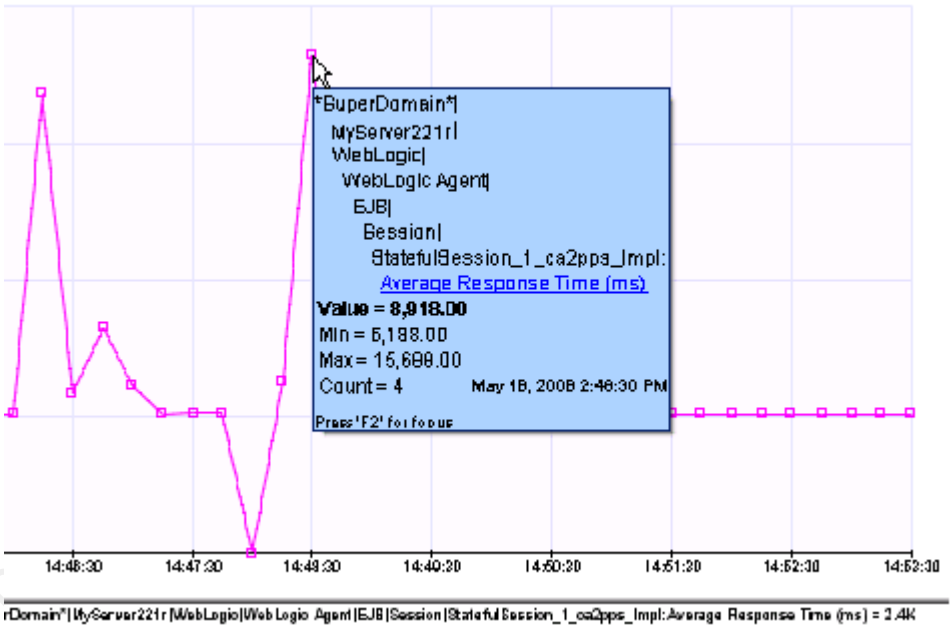
Average Response Time (ms)

Response Time is the time it takes for a request to complete; this provides a basic measurement of application response speed. Therefore:

- Low response times are desirable.
- High response times suggest a problem.

The Average Response Time metric averages the response times of all requests that were completed during an interval.

Note: The *count* for Average Response Time is identical to the value of Responses Per Interval.



The illustration above shows an Average Response Time graph for an EJB session, as displayed in Introscope Workstation. Things to notice:

- Mouse over a data point to see a tooltip with more information about the data point.
- In the example above:
  - the *value* of the data point, 8919 ms, is the average response time of the requests completed during the interval.
  - the count, 4, means that four requests were completed during the interval selected.

- In addition to value and count, each data point has min and max data.
  - Min is the lowest single value of the requests represented in the count—in this example, the request that took the least time to be completed.
  - Max is the highest single value of the requests represented in the count—in this case, the request that took the most time to be completed.

Triaging using Average Response Time

You can use trends in Average Response Time, coupled with changes in other metrics, to identify and diagnose problems. (See the index to find information on the other metrics mentioned in this section.)

Consistent problems

When accompanied by a low Available Thread count, consistently high Average Response Times may indicate the following problems:

- Inefficient code
- Overuse of external system
- Slow backend
- Too many layers

Periodic problems

Periodically high Average Response Times are shown by a graph which periodically spikes, then returns to normal.

When accompanied by a low Available Thread Count, periodically high Average Response Times may indicate:

- Frequent GC leaks
- Load-related backend bottleneck

When accompanied by a low CPU Utilization reading, periodically high Average Response Times may indicate:

- Internal chokepoint

Progressive problems

A steady increase in Average Response Time over a long period, when accompanied by a low Responses Per Interval reading, may indicate

- memory leak

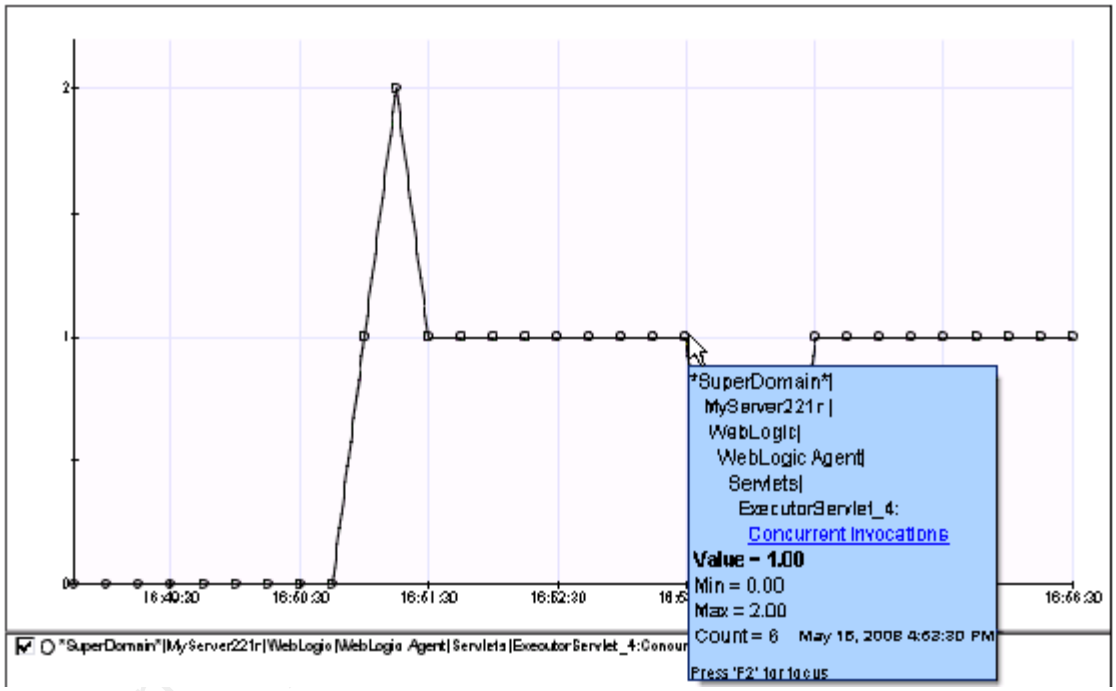
CA Technologies Prerelease Documentation. Authorized Use Only.

Concurrent Invocations

Invocations are requests handled by the application and its various parts; concurrent invocations are the requests being handled at a given time.

Introscope calculates the Concurrent Invocations metric by counting the number of requests which were still in flight (that is, which were still being handled) at the end of a particular interval.

- A low Concurrent Invocations value is desirable.
- A high Concurrent Invocations value suggests a problem.



In the illustration above, the *value* of 1 indicates that one request was still in flight, or being handled, at the end of the selected interval.

Requests that were still in flight at the close of the selected interval will likely be completed during subsequent intervals. Those which are not completed before the end of a specified threshold are called stalls (see Stall Count (see page 357)).

Triaging using Concurrent Invocations

You can use trends in Concurrent Invocations, coupled with changes in other metrics, to identify and diagnose problems. Note that the following guidelines refer to the Concurrent Invocations value, not count.

Consistent problems

Consistently high Concurrent Invocation values may indicate the following problems:

- Overuse of external system
- Slow backend

When accompanied by a low Responses Per Interval reading, consistently high Concurrent Invocation values may indicate:

- Inefficient code
- Too many layers

Periodic problems

Periodically high Concurrent Invocation values are shown by a graph which periodically spikes, then returns to normal. This may be an indication of:

- load-related backend bottleneck

When accompanied by a low Available Connections reading, periodically high Concurrent Invocation values may indicate:

- frequently collected garbage leaks

When accompanied by a low Available Thread Count reading, periodically high Concurrent Invocation values may indicate:

- internal choke point

Progressive problems

A steady increase in Concurrent Invocations over a long period, especially when accompanied by a low Responses Per Interval reading, may indicate:

- thread leak

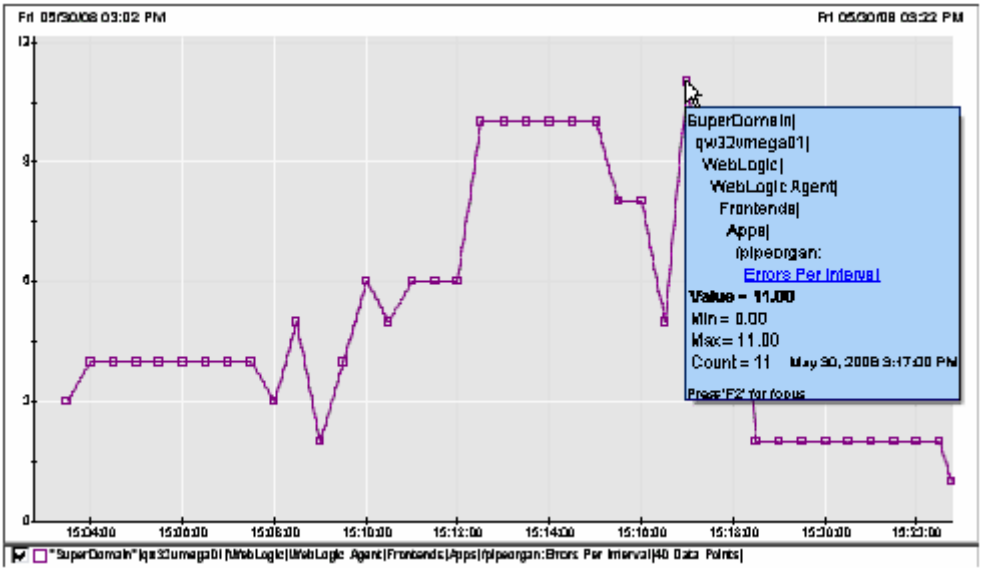
Errors Per Interval

Errors are the number of exceptions reported by JVM and HTTP error codes. Examples of errors include:

- a 404 Page Not Found status reported by the HTTP server
- a SQL exception
- a Java exception

Obviously, a low error count is desirable.

CA Technologies Pre-release Documentation. Authorized Use Only.



The metric is a simple count of errors reported during the interval. The illustration above shows one data point selected with a value of 11, meaning 11 errors were reported during that timeslice. Since this is a simple count metric, the value and Max value will always be the same.

The metric path beneath the graph identifies the application reporting the exception. To find more information about the errors shown in a graph, check the logs for that application.

Error snapshots

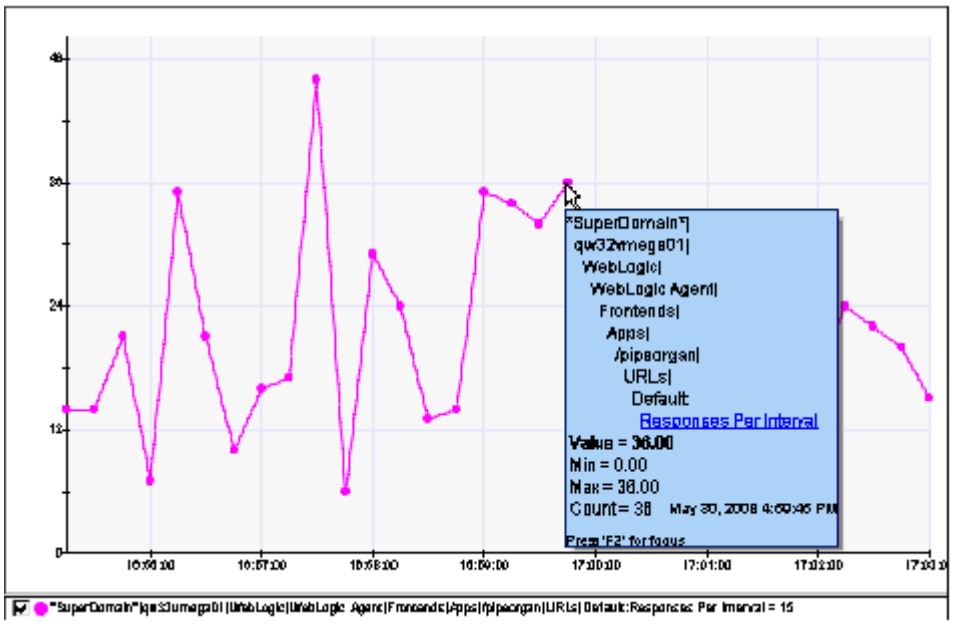
For systems with ErrorDetector enabled, errors also generate error snapshots—detailed information about what was happening when an error occured—which are stored in the Transaction Events database. A large number of errors will generate a large amount of documentary information, and preventing this is another reason to minimize errors.

Responses Per Interval

Responses Per Interval reflects number of invocations finished that interval; it is a measure of data throughput and thus of application performance. Generally:

- A high number is desirable.
- A low number is undesirable.

- Of course, an unexpected spike in responses could indicate overuse of the external system, such as a denial of service attack on a website.



The metric is a simple count of requests completed during an interval.

In the illustration above, the tooltip shows the value of the selected data point. Since this is a simple count metric, the value and the Max value of the metric will always be the same.

Please note:

- The value of the Responses Per Interval metric is always the same as the count for the Average Response Time metric.
- Responses Per Interval is a metric of type IntCounter. It is not an average of the number of responses; it is always the Max value of the number of responses during the interval.

Triaging with Repsonses Per Interval

You can use trends in Repsonses Per Interval, coupled with changes in other metrics, to identify and diagnose problems. (See the index to find information on the other metrics mentioned in this section.)

Consistent problems

Consistently high Responses Per Interval values may indicate:

- Over-usage of external system

Stall Count

Stalled requests are those which have not completed within a specified time threshold. If a request is counted as stalled, that does not mean it is hung and will never be completed, but that its execution exceeded the stall threshold.

- A low count is desirable.
- A high count is undesirable.

The default stall threshold is 30 seconds.

Information on stall events is stored in the Transaction Events database.

How stall count is measured

Occasionally, a Transaction Trace will show several requests that were not completed during the specified time threshold, i.e. stalls, but the Investigator will display a different number as the stall count.

This is because stall count is recorded as a point value (at a point in time during an interval) and not as a range value (for a time period). This means that while there could be several stall values representing long transactions that are completed during an interval, only the count available during a single moment is used as the data point.

Triaging with Stall Count

You can use trends in Stall Count, coupled with changes in other metrics, to identify and diagnose problems. (See the index to find information on the other metrics mentioned in this section.)

Consistent problems

Consistently high Stall Count values may indicate:

- Slow backend system

Periodic problems

Periodically high Stall Count values may indicate:

- Load-related backend bottleneck

Progressive problems

A steady increase in Stall Count values over a long period, espeically when accompanied by a low Available Threads count, may indicate:

- Resource leak—threads

CA Technologies Preferred Documentation. Authorized Use Only.

Other common metrics

In addition to the five metrics that commonly appear with instrumented methods, you can see other common metrics at various places in the Investigator tree.

Memory-Related Metrics

Garbage Collection is the process of freeing memory taken up by objects no longer in use; once memory is freed up it is useable by other objects.

- GC Heap metrics (see page 359): Introscope reports GC Heap metrics by default. These metrics use bytes as a unit of measure.
- GC Monitor metrics (see page 359): These provide additional information about memory use. These metrics are not collected or reported until they are enabled by an administrator.

In addition, File system, UDP and Sockets metrics (see page 365) are measures of data throughput.

Garbage collection concepts

Garbage collection is the automatic reclamation of memory devoted to objects that are no longer used by an application. When the process encounters an object that is unused, the memory is reclaimed; when the process encounters an object that is still live, it is copied to a later-generation memory pool. As young generation memory pools fill, minor garbage collection takes place, and live objects are copied to the second survivor space memory pool. When this second survivor space is not sufficient to hold all objects, live objects are also copied to the tenured memory pool spaces.

Conceivably, garbage collection could take place extremely often, so that the amount of reclaimed memory is maximized, but this would require too much overhead devoted to the process. Conversely, garbage collection that did not occur often enough would leave too little memory, and when the process did occur, it would also require significant overhead to execute. Therefore, garbage collection is most efficient when the right amount of time elapses between minor garbage collections to balance the number of objects cleaned up with the amount of overhead required to clean them.

In an efficient garbage collection process, young generation memory pools are the right size. If they are too small, automatic garbage collection takes place too often. If they are too large, too many unused objects accumulate and cause the less-frequent GC process to use too much overhead when it runs; this would cause a spike in percentage of time spent in garbage collection.

CA Technologies Confidential/Restricted Use Only
Product Documentation

GC Heap Metrics

These metrics are enabled by default.

GC Heap|Bytes In Use

GC Heap|Bytes In Use reports the amount of memory being currently used by objects.

GC Heap|Bytes Total

GC Heap|Bytes Total reports the total amount of memory allocated by the JVM.

Contrast this with the metric Current Capacity (bytes) (which is available if you have GC Monitor enabled). Current Capacity gives information about amount of memory committed for all JVM memory segments, whereas Bytes Total tells the amount of memory committed to the JVM in total.

GC Monitor Metrics

GC Monitor metrics report information about Garbage Collectors and Memory Pools, helping you detect GC issues which are adversely affecting performance.

GC Monitor metrics appear in the Metric Browser tree directly below the GC Heap node, in the hierarchy shown below.

CA Technologies Prerelease Documentation. Authorized Use Only

The metrics are enabled by default.

- GCMonitor
  - GC Policy
  - JVM Type
  - Percentage of Java Heap Used
  - Garbage Collectors
    - *<Collector1>*
      - GC Algorithm
      - GC Invocation Per Interval
      - GC Invocation Total Count
      - GC Time Per Interval (ms)
      - Percentage of time spent in GC during last 15 mins
      - Total GC Time (ms)
    - *<Collector2>*
      - (Metrics same as above)
  - Memory Pools
    - *<Memory Pool1>*
      - Amount of Space Used (bytes)
      - Current Capacity (bytes)
      - Growth Rate
      - Maximum Capacity (bytes)
      - Memory Type
      - Percentage of Maximum Capacity Currently Used
    - *<Memory Pool2>*
      - (Metrics same as above)

Limitations

Note the following limitations:

- These metrics can be returned only by version 9.x and above agents.
- These metrics can only be collected from supported JVMs, listed below.

CA Technologies Auth

Supported JVMs and their Garbage Collectors and Memory Pools

GC Monitor metrics are available only for the following JVMs:

- Sun JVM 1.5 and 1.6
- IBM JVM 1.5 and 1.6

Generic metrics

Some of the metrics below have preset thresholds which trigger alert indicator displays in the GC Monitor Overview tab (see page 136). For more information about how to use these metrics to tune memory allocation, see Monitoring Performance with the GC Metrics (see page 164).

GC Policy

Names of the Garbage Collector used in the JVM.

JVM Type

Identifies the JVM being monitored.

Percentage of Java Heap Used

Percentage of the available heap memory currently used on the machine where the agent is deployed.

- Caution threshold: 60%
- Danger threshold: 80%

By default, the virtual machine grows or shrinks the heap at each collection to try to keep the proportion of free space to live objects at each collection within a specific range. This target range is set as a percentage by the parameters -XX:MinHeapFreeRatio=<minimum> and -XX:MaxHeapFreeRatio=<maximum>, and the total size is bounded below by -Xms and above by -Xmx .

The default size may often be too small.

This metric should remain under 60%. If it goes over 80%, you should tune the JVM heap size by granting sufficient and affordable memory to the virtual machine using -Xms and -Xmx parameters.

The default values of target range is set to 30% for minimum and 70% to maximum. Larger applications often experience problems with the default values. One problem may be slow startup, which occurs when the initial heap is small and must be resized over many collections. Setting -Xms and -Xmx to the same value increases predictability by removing the most important sizing decision from the virtual machine. On the other hand, the virtual machine can't compensate if you make a poor choice.

Be sure to increase the memory as you increase the number of processors, since allocation can be parallelized.

CA Technologies Prerelease Documentation. Authorized Use Only.

Garbage Collector metrics

GC Algorithm

Name of the garbage collection algorithm used by the corresponding memory manager.

GC Invocation Per Interval

A Count metric reporting the number of Garbage Collections which occurred in each 15 second interval. This is an aggregated metric calculated from GC Invocation Total Count by tracking the difference in GC invocation count between the current and the most recent interval.

This metric indicates per-interval collection done on the memory pool. If it increases over a period of time, it indicates that there are frequent collections happening on a memory pool and it is not the right size. Increasing the memory pool size helps reduce frequent garbage collection.

GC Invocation Total Count

Total number of Garbage Collections that have occurred since the JVM was started.

This metric indicates number of collections since server startup time. It should grow slowly during garbage collections at regular intervals.

Metric spikes indicate frequent collections, affecting overall application throughput. To reduce GC frequency and raise throughput, increase the memory pool size.

GC Time Per Interval (ms)

The amount of time Garbage Collection took during the 15 second interval. This is an aggregated metric calculated from GC total time by tracking the difference in GC time between the current and the most recent interval.

Under normal behavior, this metric remains steady, or grows slowly as the time taken for garbage collection increases.

Drastic increases indicate slow application execution time by increasing garbage collection pause times. To avoid this, configure max memory using the -Xmx flag to an optimal value. Proper adjustment will cause GC pause times to fall and improve GC throughput. If you set memory too high, GC frequency will fall and GC throughput/efficiency will improve but the application will experience long pause times as the system tries to maintain a too-large heap space. An optimal heap size ensures low pause times and garbage collection times.

CA Technologies Proprietary Documentation
Authorized Use Only

Percentage of time spent in GC during last 15 mins

This metric is an aggregated metric calculated using an Enterprise Manager calculator. First, the *Sum* of values for total GC time for 15 mins duration (60 values with heartbeat interval of 15 secs) is calculated. The first data point value (firstValue) is subtracted from this *Sum* to get the actual time spent in GC during this 15 min interval. The percentage of this value is then calculated using this formula:

```
((Sum - firstValue) / (15 * 60 * 1000)) * 100
```

This metric should grow slowly over time.

- Drastic increases indicate slow application execution time by increasing garbage collection pause times. Configure max memory using the -Xmx flag to an optimal value.
- If the metric is steady and then shows a sudden spike, this signifies a one-off garbage collection that took more than the usual time. After this spike, the metric returns to normal, and no action is required.

Total GC Time (ms)

The cumulative total of time used by the garbage collection process, in milliseconds.

Under normal behavior, this metric increases gradually.

Drastic increases indicate slow application execution time by increasing garbage collection pause times. To avoid this, configure max memory using the -Xmx flag to an optimal value. Proper adjustment will cause GC pause times to fall and improve GC throughput.

Memory Pool metrics

Amount of Space Used (bytes)

The amount of used memory in bytes. It includes the memory occupied by all objects in the pool including both reachable and unreachable objects.

Under normal behavior, this metric increases gradually. It may fall when garbage collection is finished and memory is reclaimed.

A temporary spike followed by a return to normal could be a warning of memory problems.

In a rapid increase, the metric may reach maximum memory limit and cause Out of memory exceptions. To avoid this, set the maximum size of the memory pool to a more affordable value.

CA Technologies Product Documentation. Authorized Use Only.

Current Capacity (bytes)

The amount of memory in bytes committed for use for this pool for all JVM memory segments. This amount of memory is guaranteed for the JVM to use.

NOTE: Adding all Current Capacity metrics from individual memory segments should approximately equal the Bytes Total metric (see GC Heap Metrics (see page 359)).

If the amount of space used reaches the current capacity, it would throw memory exceptions. To avoid this, take into account the need to handle not just day-to-day operations but also unexpected peaks.

Growth Rate

Average growth rate, expressed in bytes/second, of used memory in a memory pool, in bytes per second over the past minute. This aggregated metric is calculated by finding the difference between the last data point value in bytes (lastValue) and the first data point value in bytes (firstValue) of memory space used in the most recent 1 minute interval. The rate is calculated using this formula:

```
(lastValue - firstValue) / 60
```

This metric may grow slowly, remain steady, or may fall if the unused memory is returned to the pool.

Drastic increases over an interval of 15 minutes or more indicates that memory is not being recycled after garbage collection. This indicates a possible memory leak. Further investigation is called for.

Maximum Capacity (bytes)

The maximum amount of memory in bytes that can be used for memory management. This amount of memory is not guaranteed to be available for memory management if it is greater than the Current Capacity (amount of committed memory).

This metric remains steady over time.

Memory Type

Type of memory; one of:

- Heap
- Non-Heap

Percentage of Maximum Capacity Currently Used

Percentage representation of the current memory usage over maximum amount of memory. This metric indicates the percentage of memory used over time.

This metric may grow slowly, remain steady, or may fall if the unused memory is returned to the pool.

If the metric exceeds 70-80%, set maximum memory to a higher and optimal value.

CA Technologies Pre-release Documentation
Authorized Use Only.

File system, Sockets, UDP

Like Responses Per Interval (see page 355), these are measures of data throughput. They are measured in Bytes Per Second:

File system

- File output rate (bytes per second)
- File input rate (bytes per second)

UDP (User Datagram Protocol)

- Output bandwidth (bytes per second)
- Input bandwidth (bytes per second)

Sockets (total as well as host/port specific information)

- Output bandwidth (bytes per second)
- Input bandwidth (bytes per second)

A large number of port-related metrics indicates socket rate metrics should be turned off, because this is possibly a metric explosion problem.

For other socket metrics, see Socket metrics (see page 366).

Utilization metrics

Utilization metrics measure the percentage of available resources being used. The most common is CPU Utilization.

CPU Utilization

CPU utilization is measured by Introscope's platform monitor, and measures the amount of CPUs being used. There are two different measurements:

CPU:Utilization % (process)

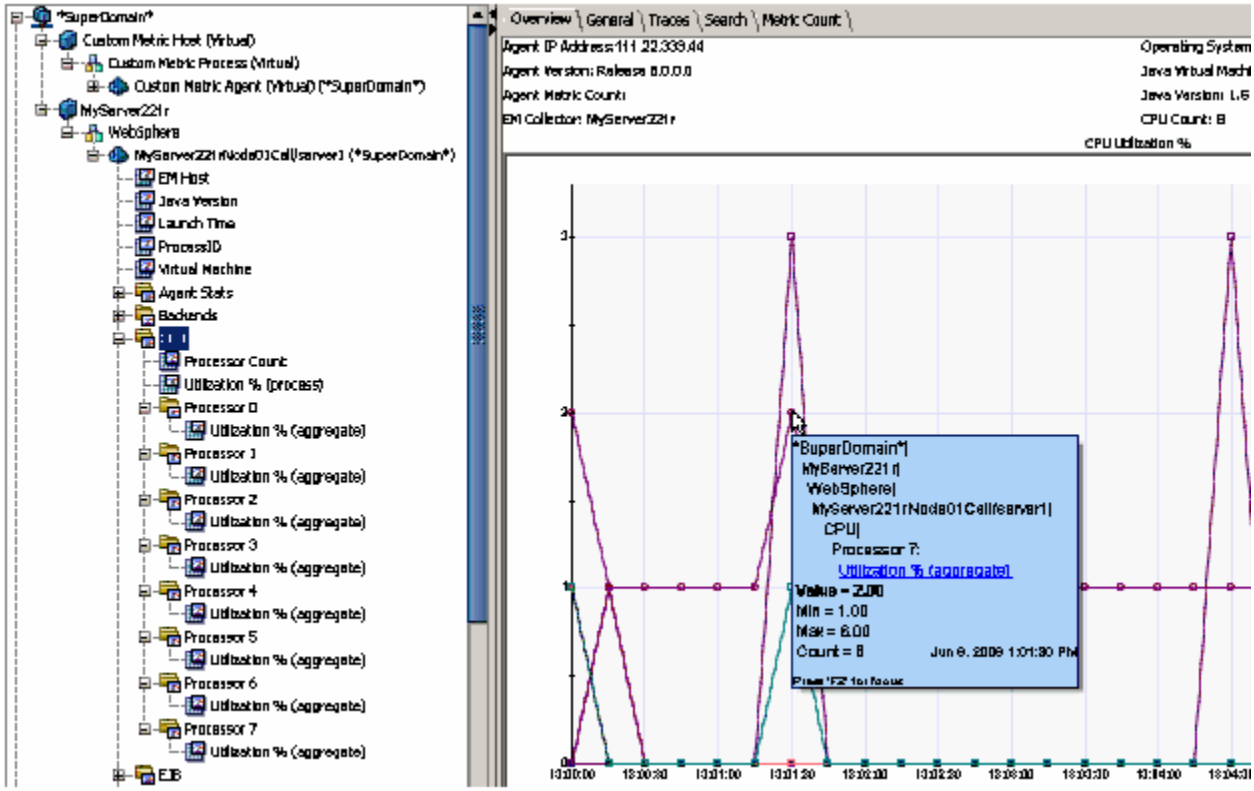
Percentage of the total computing power of the Introscope host, but limited to the percentage utilized by the JVM process that Introscope monitors.

CPU:Utilization % (aggregate)

Utilization of an individual processor.

CA Technologies Prerelease Documentation. Authorized Use Only.

The illustration below shows CPU utilization metrics for an 8 processor host. One of the data points is selected.



Socket metrics

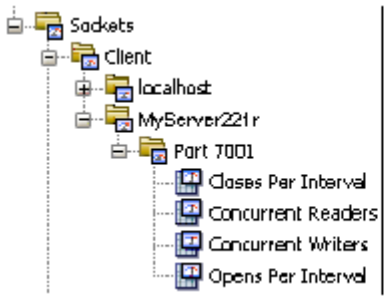
Socket metrics are reported by port by type:

- Client sockets
- Server sockets

CA Technologies Proprietary Authorized

They are displayed at the following location in the Investigator tree:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Sockets | [Client|Server] | Enterprise Manager | Port`



Accepts Per Interval

The number of Accepts per interval.

Closes Per Interval

The number of sockets per interval that were closed.

Concurrent Readers

The number of threads being read by this port, per interval.

Concurrent Writers

The number of threads being written using this port, per interval.

Opens Per Interval

The number of sockets per interval that were opened.

Input Bandwidth (Bytes Per Second)

Input bandwidth for the port, measured in bytes per second.

Output Bandwidth (Bytes Per Second)

Output bandwidth for the port, measured in bytes per second.

CA Technologies Prerelease Documentation
Authorized Use Only.

Thread Dump Metrics

Deadlock count metric

The current number of deadlocked threads. This metric displays here in the metric browser tree:

<Agent name> | Threads | Deadlock Count metric

This metric is not enabled by default. To enable the Deadlock Count metric, see the *CA APM Java Agent Implementation Guide*.

Thread pool metrics

The Threads metric shows the number of active, or currently in use, instrumented threads created from classes that have had probes added by Introscope. The metrics are typically gathered from JMX (on Java applications) or PMI (on WebSphere applications).

The metrics are divided into:

- I/O threads
- Worker threads

For both of these types, you can view the following metrics:

Active Threads

Number of active threads.

Available Threads

The total number of threads available.

Maximum Idle Threads

The maximum number of threads that can be idle.

Minimum Idle Threads

The minimum number of threads that can be idle.

Threads in Use

The number of threads in use.

CA Technologies Prerelease Documentation. Authorized Use Only

Thread Creates

The number of created threads during the interval.

Thread Destroys

The number of destroyed threads during the interval.

OpenSessionsCurrentCount

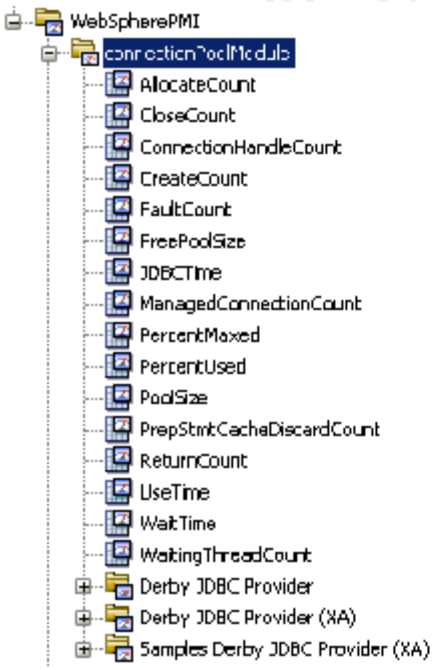
The number of currently open sessions.

Connection pool metrics

Connection pool metrics are typically gathered from JMX (on Java applications) or PMI (on WebSphere applications). The metrics are typically divided into:

- Count metrics for individual connection types
- Percent metrics
- Time metrics

The illustration below shows all three kinds of connection pool metrics configured for a WebSphere application.



CA Techr Release Documentation. Use Only.

Connection pool count metrics

Counts of various kinds of connections, depending on what has been configured for the application. These usually include:

PoolSize

The number of total connections in the connection pool.

FreePoolSize

The number of free connections in the connection pool.

avgUseTime

Average Use Time.

avgWaitTime

Average wait time.

concurrentWaiters

Number of waiting threads.

faults

Number of faults.

jdbcOperationTimer

numAllocates

numConnectionHandles

numCreates

numDestroys

numManagedConnections

numReturns

prepStmtCacheDiscards

CA Technologies Prerelease Documentation. Authorized Use Only.

Connection Pool percent metrics

PercentMaxed

The percentage of connections in the connection pool that are maxed out.

PercentUsed

The percentage of connections in the connection pool that are active.

Event metrics

Event metrics are recordced by Introscope in specific situations. They include:

- stalls (see Stall Count (see page 357))
- system logs

  This Metric type monitors the application system out and system error output. It is typically turned off. See System logs (see page 371).
- exception

  Captures throwing/catching exceptions. Provides the ability to trace all locations where exceptions are thrown and caught.

Note: Exception catching should be turned off in production as it can cause significant performance degradation.

System logs

Standard error

Prints the stderr log in text format.

Standard output

Prints the stdout log in text format.

CA Technologies Prerelease Documentation. Authorized Use Only

Resource Metrics

Resource metrics are available for all locations. They provide health information on that location's resources. In the Metric Browser tree they appear under the agent node as follows:

```
<Agent_Name>
|
|--Agent Stats
   |
   |--Resources
      |
      |--<metrics>
```

Resource metrics are based on metric paths you specify in the ResourceMetricMap.properties file. See information about that file in the *APM Configuration and Administration Guide.*

Note: Depending on the application server, not all these metrics are available for all resources.

% CPU Utilization (Host)

% of total available CPU (central processing unit) resources being utilized on the host

% Time Spent in GC

% of the CPU time during the interval that the JVM was occupied by the GC (garbage collection) process. (Also see Memory-Related Metrics (see page 358))

Threads in Use

Total number of threads in use at the end of the interval. (Also see Thread Pool Metrics (see page 368))

Threads Available

Total number of threads available during the interval.

JDBC Connections in Use

Total number of JDBC connections in use at the end of the interval.

JDBC Connections Available

Total number of JDBC connections available during the interval.

CA Technologies Prerelease Documentation. Authorized Use Only

Customer Experience Metrics

These metrics are available when a TIM has been configured to report metrics on a business service. These are distinct from standard Introscope health metrics reported on a business transaction component (BTC), but may be used in comparison with BTC health metrics.

In the Map tree they appear under the Customer Experience node:

```
By Business Service
|
|--<Business Service Name>
   |
   |--<Business Transaction Name>
      |
      |--<Business Transaction Component Name>
      |--Customer Experience
         |
         |-<metrics>
```

In the Browse tree they appear under the CEM node:

```
Domain|<host>|CEM|Tess Agent|Business Process|<metric>
```

Average Response Time (ms)

Average Response Time for the business transaction during the interval.

Total Transactions Per Interval

Total transactions for the Business Transaction across all reporting TIMs, per interval.

Total Defects Per Interval

Number of defects per interval. Defects are defined in the CE interface based on certain events captured by the TIM.

Customer Experience Transaction Metrics

Customer experience transaction metrics are collected by deployed TIMs. They are formerly known as BT Stats metrics because they provide metrics on Business Transactions, and also (formerly) as real-time transaction metrics or RTTM.

To configure these metrics, see the information about configuring real-time transaction metrics integration in the *CA APM Integration for CA Infrastructure Management Guide.* That section gives you information on administering and configuring these metrics.

CA Technologies Prerelease Documentation. Authorized Use Only.

How the metrics are calculated

Customer experience transaction metrics are calculated using Javascript calculators on the Enterprise Manager which is collecting metrics from the TIM.

NOTE: Aggregated metrics will be calculated only on a collector Enterprise Manager with a running TIM Collection Service and BTstats processor. These calculators will not run on a MOM Enterprise Manager.

Defect types

Real-Time Transaction Metrics (RTTM metrics) are grouped into several defect types. RTTM metrics can appear under any of these types, which are the default names of defects before users customize them.

Defect metrics will be collected for each defect type associated with the business transaction, including user-named transactions -- such as "Slow time for *<BT_Name>".*

Following are the default values for each defect type, where s = second.

Slow Time

Transaction Time > 5.000 s

Fast Time

Transaction Time < 0.005 s

High Throughput

Throughput > 100.0KB / s

Low Throughput

Throughput < 1.0KB/ s

Large Size

Transaction Size > 100.0KB

Small Size

Transaction Size < 0.1KB

Missing Transaction

Component Timeout = 10.000 s

CA Technologies Prerelease Documentation. Authorized Use Only

Metrics Aggregated Across TIMs

Metrics aggregated across TIMs are values that are calculated on the Enterprise Manager, not directly sent by the TIM.

Metrics per Business Transaction for a given Business Service

Total Defects Per Interval

Number of defects for all defect types for a business transaction, aggregated across TIMs.

Average Response Time (ms)

For each interval, the average time it took to execute the business transaction, in milliseconds.

Total Transactions Per Interval

Total number of transactions for the business transaction, aggregated across TIMs, per interval.

Metrics per defect type for a Business Transaction for a given Business Service

Defects Per Interval

A count metric showing the number of failed transaction opportunities over the most recent interval.

Metrics Reported by an Individual TIM

Each of these metrics is reported by a single TIM (Transaction Impact Monitor) monitoring transactions on a single machine.

NOTE: The percentage metrics are calculated values; the others listed in this section are reported directly by a TIM.

Metrics per Business Transaction for a given Business Service

Total Defect Ratio (%)

The number of total defects for all transactions for a Business Transaction, aggregated across TIMs.

Total Defects Per Interval

Number of defects for all defect types for a business transaction, aggregated across TIMs.

Average Response Time (ms)

For each interval, the average time it took to execute the business transaction, in milliseconds.

CA Technologies Prerelease Documentation
Authorized Use Only.

Total Transactions Per Interval

Total number of transactions for the business transaction, aggregated across TIMs, per interval.

Metrics Per Defect Type for a Business Transaction for a given Business Service

Defect %

A defect is a single transaction opportunity that failed. Defect percentage is calculated using the following formula, where BT stands for Business Transaction:

Rounded Integer Value for (Defects Per Interval for a given defect type / Total Transactions Per Interval Count for the BT) * 100

Defects Per Interval

A count metric showing the number of failed transaction opportunities over the most recent interval.

Using perflog.txt

The Enterprise Manager records performance time for system events in a performance log file, *<EM_home>/logs/perflog.txt*. As an alternative to the metrics displayed in the Investigator, this file may contain useful information. For information on reading and understanding this file, see the *APM Sizing and Performance Guide*.

Other metrics

Depending on your system architecture, the following metrics may also appear in the Workstation Investigator. Most appear in the Metric Browser tree, but the section Application Triage Map metrics has metrics which appear in the Triage Map tree.

Application triage map metrics

The following metrics in the Triage Map tree show specific component and segment information for MOM and Collector Enterprise Managers. These numbers allow you to see the Enterprise Manager and APM database usage overhead associated with the application triage map. The new metrics are:

ApplicationTriageMap|TransactionSegmentsReceived

The aggregate number of segments recorded by the agent and sent to the Collector in the last harvest period. This metric implies the proportional amount of time devoted to APM database update or insert during the last harvest period. However, it does not add time to the harvest period itself, but only indicates the load on the APM database and the system as a whole.

CA Technologies Product Documentation
Authorized Use Only.

ApplicationTriageMap|ProcessingTime|TransactionSegment(ms)

The aggregate time, in ms, it takes the system to record transaction segment information, send it to the Enterprise Manager, and store it in the database, during the last harvest period. This metric applies only to collector Enterprise Managers.

ApplicationTriageMap|ProcessingTime|TransactionComponent(ms)

The aggregate time, in ms, it takes the system to record transaction component information, send it to the Enterprise Manager, and store it in the APM database, during the last harvest period. This metric applies only to collector Enterprise Managers.

ApplicationTriageMap|TransactionComponentsReceived

The aggregate number of components recorded by the agent and sent to the Collector in the last harvest period. This metric implies the proportional amount of time devoted to APM database update or insert during the last harvest period.

EJB

Where Enterprise Java Beans (EJBs) are part of your architecture, they may be of the following types:

- EJB entity
- EJB session
- EJB message driven

For each of these types, the following two metrics appear:

- Average Method Invocation Time (ms)
- Method Invocations Per Interval

For each EJB implementation (class or method), which appears as a child node under EJB types, Enterprise Manager reports the five basic Introscope metrics:

- Average Response Time (ms)
- Concurrent Invocations
- Errors Per Interval
- Responses Per Interval
- Stall Count

More information:

The Five Basic Introscope Metrics (see page 349)

CA Technologies Prerelease Documentation. Authorized Use Only.

Servlets

The Servlets node commonly displays the five basic Introscope metrics for each of the servlets invoked by the application being monitored by Enterprise Manager:

- Average Response Time (ms)
- Concurrent Invocations
- Errors Per Interval
- Responses Per Interval
- Stall Count

More information:

The Five Basic Introscope Metrics (see page 349)

JDBC

The JDBC node commonly displays these metrics for JDBC calls invoked by the application being monitored by Enterprise Manager:

Average Result Processing Time (ms)

The average number of milliseconds it takes for a round-trip query, averaged over an interval.

Queries Per Second

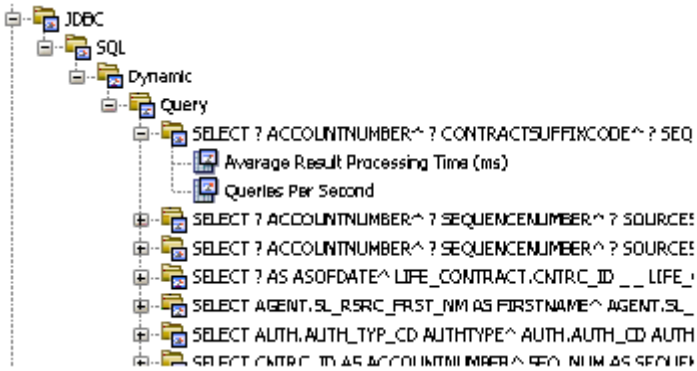
The number of times this query was issued during each interval.

CA Technologies Pre-release Documentation. Authorized Use Only.

For example, you can configure the agent to monitor the performance of individual SQL series using the JDBC protocol. The metric path would be something like:

```
*SuperDomain* | <Host_Name> | <Process_Name> | <Agent_Name> | JDBC | SQL |
[Dynamic|Static] | Query | <SQL_Query>
```

The illustration below shows how these metrics are reported for each query executed:



Connection Count

Number of JDBC connections per interval.

Commit Count

Cumulative count of commits since agent start.

Commits Per Second

Average number of commits per second over the 15-second interval.

JSP (Java Server Pages)

Average Response Time (ms)

Average response time of the *_jspService()* methods of all the JSPs executing in the JVM. The Response Times of all the individual JSPs are averaged to calculate this value.

Responses Per Interval

Number of completed invocations of the *_jspService* methods of all the JSPs executing in the JVM in the past 15 second time period.

CA Technologies Prerelease Documentation
Authorized Use Only

Average Response Time (ms) by class name

Average response time in milliseconds of the JSP identified by the class name. Each invocation of the *_jspService()* method is timed and averaged to arrive at this value.

Responses Per Interval

Number of completed invocations of the *_jspService()* method of the JSP identified by the class name in the most recent 15 second interval.

Responses Per Second

Rate at which the *_jspService()* methods of all the JSPs executing in the JVM are being completed.

Responses Per Second by class name

Rate at which invocations of the *_jspService()* method of JSP identified by a particular class name are being completed.

Stalled Methods by class name and by method name

The number of JSPs that are taking longer than a defined threshold to complete the execution of the *_jspService()* method.

Concurrent Invocations

The number of threads executing the *_jspService()* method.

JSP tag libraries (JSP TagLib)

Tag libraries are collections of custom tags used in JSP pages to invoke custom actions. A custom action is any action not included in the set of six standard actions provided for in the JSP specification. Examples of tasks invoked by custom actions are form control, accessing external systems like databases and email, and flow control.

The following metrics are available for JSP tag libraries:

- Average Method Invocation Time (ms)
- Method Invocations Per Interval
- Average Method Invocation Time (ms) by class name and method name
- Method Invocations Per Interval by class name
- Method Invocations Per Interval by class name and method name
- Method Invocations Per Second
- Method Invocations Per Second by class name
- Method Invocations Per Second by class name and method name

CA Technologies Prerelease Documentation. Authorized Use Only.

- Concurrent Method Invocations
- Concurrent Method Invocations by class name
- Concurrent Method Invocations by class name and method name
- Stalled Methods over 30 seconds by class name and method name
- Average Method Invocation Time (ms)

JSP IO TagLibrary

- Warning Count
- Exception Count

RMI (Remote method invocations)

Remote method invocations are invocations of methods of distributed Java objects—that is, Java objects which may exist on more than one host.

The following metrics are available for both RMI clients and RMI servers.

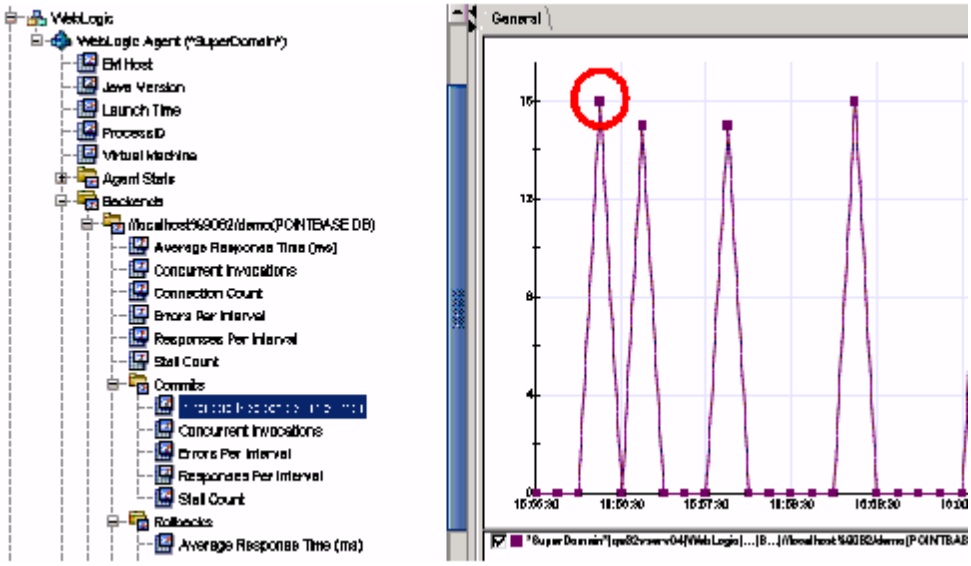
- Average Method Invocation Time (ms)
- Method Invocations Per Interval
- Average Method Invocation Time (ms) by class name
- Method Invocations Per Interval by class name
- Method Invocations Per Second
- Method Invocations Per Second by class name
- Stalled Methods over 30 seconds
- Concurrent Method Invocations
- Concurrent Method Invocations by class name

CA Technologies Pre-release Documentation. Authorized Use Only

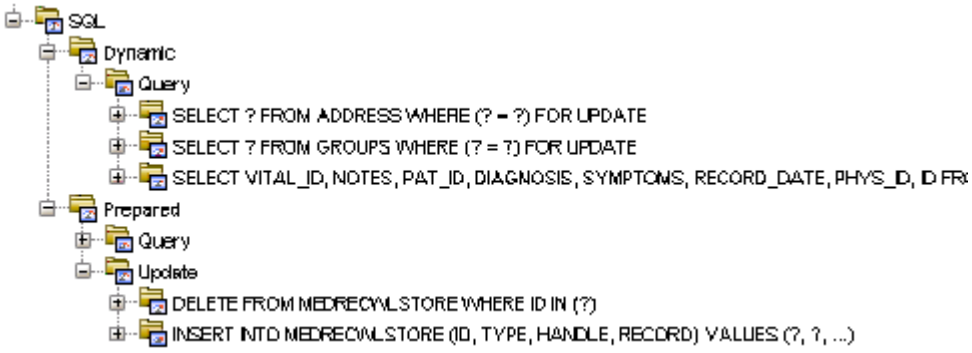
Database metrics (SQL)

Each database backend can be configured to report the following metrics:

- Commits—Each completed query-and-response transaction is known as a *commit.* The five standard metrics are collected and displayed for all the transactions that commit in a given interval. For example, in the screenshot below, the circled datapoint shows the average response time for all the committed database transactions in that interval.



- Rollbacks—A *rollback* is an unsuccessfully completed query-and-response transaction. The five standard metrics are collected and displayed for all the rolled back transactions in a given interval.
- SQL:



For each of the statements processed by the database during a given interval, six metrics are reported:

- Average Response Time (ms)
- Concurrent Invocations
- Errors Per Interval
- Connection Count
- Responses Per Interval
- Stall Count

Things to notice:

- The statements are separated by subnode according to whether they are Prepared or Dynamic.
- Each type of SQL statement—for example, *GRANT, UPDATE, QUERY, REVOKE, DROP*—is listed under a subnode for that statement type.

XML (Extensible Markup Language)

XML metrics can be of the following types.

SAX

- SAX:Average Method Invocation Time (ms)
- SAX:Method Invocations Per Interval
- SAX:Average Method Invocation Time (ms) by class name
- SAX:Method Invocations Per Interval by class name
- SAX:Method Invocations Per Second
- SAX:Method Invocations Per Second by class name
- SAX:Stalled Methods over 30 seconds by class name and method name
- SAX:Concurrent Method Invocations
- SAX:Concurrent Method Invocations by class name

XSLT

- XSLT:Average Method Invocation Time (ms)
- XSLT:Method Invocations Per Interval
- XSLT:Average Method Invocation Time (ms) by class name
- XSLT:Method Invocations Per Interval by class name
- XSLT:Method Invocations Per Second

CA Technologies Prerelease Documentation. Authorized Use Only.

- XSLT:Method Invocations Per Second by class name
- XSLT:Stalled Methods over 30 seconds by class name and method name
- XSLT:Concurrent Method Invocations
- XSLT:Concurrent Method Invocations by class name

JAXM

- JAXM|Listener:Average Method Invocation Time (ms)
- JAXM|Listener:Method Invocations Per Interval
- JAXM|Listener:Average Method Invocation Time (ms) by class name
- JAXM|Listener:Method Invocations Per Interval by class name
- JAXM|Listener:Method Invocations Per Second
- JAXM|Listener:Method Invocations Per Second by class name
- JAXM|Listener: Stalled Methods over 30 seconds by class name and method name
- JAXM|Listener:Concurrent Method Invocations
- JAXM|Listener:Concurrent Method Invocations by class name

J2EE Connector

- Average Method Invocation Time (ms)
- Method Invocations Per Interval
- Average Method Invocation Time (ms) by class name
- Method Invocations Per Interval
- Method Invocations Per Second
- Method Invocations Per Second by class name
- Stalled Method count over 30 seconds by class name and method name
- Concurrent Method Invocations
- Concurrent Method Invocations by class name

JTA (Java Transaction API)

- Average Method Invocation Time (ms)
- Method Invocations Per Interval
- Average Method Invocation Time (ms) by class name
- Method Invocations Per Interval by class name

CA Technologies Prerelease Documentation. Authorized Use Only.

- Method Invocations Per Second
- Method Invocations Per Second by class name
- Stalled Methods over 30 seconds by class name and method name
- Concurrent Method Invocations

JNDI (Java Naming and Directory Interface)

JNDI metrics include:

- JNDI lookup (see page 385)
- JNDI lookupLink (see page 385)
- JNDI search (see page 386)
- JNDI called metrics (see page 386)

JNDI Lookup

- Lookup:Context Average Method Invocation Time (ms)
- Lookup:Context Method Invocations Per Interval
- Lookup:Context Average Method Invocation Time (ms) by class name
- Lookup:Context Method Invocations Per Interval by class name
- Lookup:Context Method Invocations Per Second
- Lookup:Context Method Invocations Per Second by class name
- Lookup:Context Stalled Methods over 30 seconds by class name and method name
- Lookup:Context Concurrent Method Invocations
- Lookup:Context Concurrent Method Invocations by class name

JNDI lookupLink

- lookupLink:Context Average Method Invocation Time (ms)
- lookupLink:Context Method Invocations Per Interval
- lookupLink:Context Average Method Invocation Time (ms) by class name
- lookupLink:Context Method Invocations Per Interval by class name
- lookupLink:Context Method Invocations Per Second

CA Technologies Prerelease Documentation. Authorized use only.

- lookupLink:Context Method Invocations Per Second by class name
- lookupLink:Context Stalled Methods over 30 seconds by class name and method name
- lookupLink:Context Concurrent Method Invocations
- lookupLink:Context Concurrent Method Invocations by class name

JNDI search

- Search:Context Average Method Invocation Time (ms)
- Search:Context Method Invocations Per Interval
- Search:Context Average Method Invocation Time (ms) by class name
- Search:Context Method Invocations Per Interval by class name
- Search:Context Method Invocations Per Second
- Search:Context Method Invocations Per Second by class name
- Search:Context Stalled Methods over 30 seconds by class name and method name
- Search:Context Concurrent Method Invocations
- Search:Context Concurrent Method Invocations by class name

JNDI called metrics

- File system I/O

JMS (Java Messaging Service)

JMS has four sub-nodes:

- message listener
- message consumer
- topic publisher
- queue sender

The following metrics can appear under any of the sub-nodes:

- Average Method Invocation Time (ms)
- Method Invocations Per Interval
- Average Method Invocation Time (ms) by class name
- Method Invocations Per Interval by class name
- Method Invocations Per Second

CA Technologies Prerelease Documentation. Authorized Use Only

- Method Invocations Per Second by class name
- Stalled Methods over 30 seconds by class name and method name
- Concurrent Method Invocations
- Concurrent Method Invocations by class name

Java Mail

Java mail has two sub-nodes:

- Java Mail (Send)
- Java Mail (sendMessage)

The following metrics can appear under either the Send or sendMessage sub-nodes:

- Transport:Average Method Invocation Time (ms)
- Transport:Method Invocations Per Interval
- Transport:Average Method Invocation Time (ms) by class name
- Transport:Method Invocations Per Interval by class name
- Transport:Method Invocations Per Second
- Transport:Method Invocations Per Second by class name
- Transport:Stalled Methods over 30 seconds by class name and method name
- Transport:Concurrent Method Invocations
- Transport:Concurrent Method Invocations by class name

CORBA

- Average Method Invocation Time (ms)
- Method Invocations Per Interval
- Average Method Invocation Time (ms) by class name
- Method Invocations Per Interval by class name
- Method Invocations Per Second
- Stalled methods in any class over 30 seconds
- Concurrent Method Invocations
- Concurrent Method Invocations by class name

CA Technologies Pre-release Documentation. Authorized Use Only.

Struts

- Average Method Invocation Time (ms)
- Method Invocations Per Interval
- Average Method Invocation Time (ms) by class name and method name
- Method Invocations Per Interval by class name
- Method Invocations Per Second
- Method Invocations Per Second by class name
- Stalled Methods over 30 seconds by class name and method name
- Concurrent Method Invocations
- Concurrent Method Invocations by class name

Instance Counts

Instance counts metrics measure the number of object instances of a particular class on the heap.

- Approximate Instance Count by package and class name

Data About Machines

The following data are reported for the machine hosting the Enterprise Manager as well as each machine with instrumented methods.

- Java Version
- Virtual machine
- Launch time
- Process ID
- Host IP address
- Host operating system
- Host wall clock time

CA Technologies Prerelease Documentation. Authorized Use Only.

Supportability metrics

Supportability metrics display information about the Enterprise Manager rather than the application it is monitoring. They appear in the Investigator tree, under:

Custom Metric Host (Virtual)

Custom Metric Process (Virtual)

Custom Metric Agent (Virtual)(SuperDomain)

Beneath this level, supportability metrics are arranged in the following hierarchy. Definitions for some of these metrics follow the list.

See also: Memory-related metrics for GC Heap (see page 359) and GC Monitor (see page 359).

Agent node

Agents have the following hierarchy:

```
<Host_Name>
    <Agent_Type> (for example, DatabaseAgent)
        <Agent_Name>
```

For each <Agent_Name>, the following metrics are available:

- ConnectionStatus, one of:
  - 3 = disconnected
  - 2 = connected, slowly or no data
  - 1 = connected
  - 0 = unmounted
- Historical Metric Count
- IsClamped, one of:
  - 1 = Clamped
  - 0 = Not clamped

  For more information about metric clamping, see Clamped transactions (see page 214).
- Metric Count
- Raw Metric Count

CA Technologies Prerelease Documentation. Authorized Use Only.

Agent metrics

<Agent name> | Transaction Tracing Events Per Interval

The total number of Transaction Trace events per agent per interval.

<Agent name> | Transaction Tracing Events Limit Exceeded

The number of times the clamp limit defined in the introscope.enterprisemanager.agent.trace.limit property was exceeded for a particular interval. This metric is displayed for each agent.

Note: You can define the clamp limit for introscope.enterprisemanager.agent.trace.limit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the <*EM_Home*>\config directory.

<Agent name> | Error Snapshot Events Per Interval

The total number of error snapshot events per agent per interval.

<Agent name> | Error Snapshot Events Limit Exceeded

The number of times the clamp limit defined in the introscope.enterprisemanager.agent.error.limit property was exceeded for a particular interval. This metric is displayed for each agent.

Note: You can define the clamp limit for introscope.enterprisemanager.agent.error.limit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the <*EM_Home*>\config directory.

CA Technologies Pre-release Documentation. Authorized Use Only.

Enterprise Manager node

Under the Enterprise Manager node, the following supportability metrics are available:

- Host
- Name
- Overall Capacity (%)
- Port
- CPU
  - EM CPU Used (%)

- Configuration
  - Agent Clusters Metric Load
  - Number of Agent Clusters
  - Number of Metric Groupings
- Connections
  - Max Number of Agent Connection Limit Exceeded Per Interval
    - Indicates whether the maximum number of agents allowed to connect to the given MOM, Collector, or Standalone Enterprise Manager been exceeded.
  - Max Number of Disallowed Agent Connection Limit Exceeded Per Interval
    - Indicates whether the maximum number of disallowed agents allowed to connect to the given MOM, Collector, or Standalone Enterprise Manager has been exceeded.
  - Metrics Queued (%)
  - Number of Agents
  - Number of Applications—The number of agent applications currently reporting data.
  - Number of Disallowed Agents
    - The number of disallowed agents connected to the given MOM, Collector, or Standalone Enterprise Manager.
  - Number of Historical Metrics
  - Number of Metrics
  - Number of Metrics Handled
  - Number of Workstations

Data Store node

Under the Data Store node, the following metrics are available:

- SmartStor
  - Metrics Appended To Query Per Interval
  - Metrics Converted From Spool to Query Per Interval
  - SmartStor Disk Usage (mb)

CA Technologies Prerelease Documentation. Authorized Use Only.

- MetaData
  - Agents with Data
  - Agents without Data
  - Metrics with Data
  - Partial Metrics with Data
  - Partial Metrics without Data
  - Write Duration (ms)
- Tasks
  - Converting Spool To Data
  - Data Append
  - Reperiodizing

- Transactions
  - Number of Dropped Per Interval
  - Number of Inserts Per Interval
  - Number of Queries Per Interval
  - Number of Traces in Database
  - Number of Traces in Insert Queue
  - TT Database Disk usage (mb)
  - Total Data Insertion Duration Per Interval (ms)
  - Total Index Insertion Duration Per Interval (ms)
  - Total Query Duration Per Interval (ms)
- Volume Space Free
  - Baseline Volume Free (mb)
  - Log Volume Free (mb)
  - Smartstor Archive Volume Free (mb)
  - Traces Volume Free (mb)

Database sub-node

- Metric Data Points Sent per Interval
- Queued Metric Data Points

CA Technologies Pre-release Documentation. Authorized Use Only.

Health Sub-node

- CPU Capacity (%)
- GC Capacity (%)
- Harvest Capacity (%)
- Heap Capacity (%)
- Incoming Data Capacity (%)
- SmartStor Capacity (%)

Internal Sub-node

The following metrics appear under the Internal sub-node.

Alerts | Total Number of Evaluated Metrics

The total number of metrics evaluated for all alerts.

Calculators | Total Number of Evaluated Metrics

The total number of metrics evaluated for all calculators.

Note: This metric appears in the Investigator only when a calculator is defined.

<calculator name> | Total Number of Evaluated Metrics

The total number of metrics evaluated for an individual calculator.

In addition, the following metrics appear under the Internal sub-node:

- Number of Connection Tickets
- Number of metric Data Queries per Interval
- Number of Queued Async Data Queries
- Number of Registered Async Data Queries
- Number of Registered Async MG Queries
- Number of Registered Async Path Queries
- Number of Transaction Trace Action Sessions
- Number of Transaction Trace Session Clients
- Number of Virtual Metrics
- AlertID
- Query memory in transit (bytes)

CA Technologies Prerelease Documentation. Authorized Use Only.

- Alerts
  - <Management_Module_Name>
    - Agent Connection Status
      - Number of Evaluated Metrics
    - Backend Heuristics
    - CPU Heuristic
    - Console Summary Alert
    - Frontend Errors Heuristic
    - Frontend Heuristics
    - Frontend Response Time Heuristic
    - JDBC Heuristic
    - JVM Heuristics
    - Thread Pool Heuristic
- GC Heap
  - Collectors
    - <Collector_Name>
      - Collection Count Per Interval
      - GC Duration (ms)
  - Pools
- Messaging
  - Active Incoming Threads
  - Active Outgoing Threads
  - Corrupted Messages Per Interval
  - Post Offices
    - <Post_Office_Name>
      - Number of Mailboxes
      - Queued Messages
- Query
  - Cache Queries Duration (ms)
  - Cache Queries Per Interval
  - Smartstor Queries Duration (ms)
  - Smartstor Queries Per Interval

CA Technologies Prerelease Documentation. Authorized Use Only.

- Threads
  - <Thread_name>
    - Blocked Count
    - Blocked Time (ms)
    - CPU Time (ms)
    - User Time (ms)
    - Wait Count
    - Wait Time(ms)

Problems sub-node

- Management Modules
  - Warning Count

Tasks sub-node

- Harvest Duration (ms)
- Smartstor Duration (ms)

Harvest metrics

Harvest Capacity

The Harvest Capacity metric displays the percent of time needed for the data harvest in a 15-second time slice. For example, if the data harvest takes 15 seconds, the metric value would be 100. The Investigator displays this metric at the location `Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Health | Harvest Capacity (ms)`

Harvest Duration

The Harvest Duration metric shows the time in milliseconds (during a 15-second time slice) spent harvesting data. It is generally a good indicator in determining whether or not the Enterprise Manager is keeping up with the current workload. You can find this metric at the following location in the Investigator tree. `Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Tasks | Harvest Duration (ms)`

For more information about this metric, see the *APM Sizing and Performance Guide*.

CA Technologies Prerelease Documentation. Authorized Use Only.

Incoming Data Capacity (%)

The capacity of the Enterprise Manager to handle incoming data. The metric is calculated by multiplying the total metric capcity by 2. For example, if 150,000 metrics are in queue waiting to be processed, and the Enterprise Manager has a capacity to handle 300,000 metrics, Incoming Data Capacity is 25%.

You can find this metric at the following location in the Metric Browser tree:
`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Health | Incoming Data Capcity (%)`

For more information about this metric, see the *APM Sizing and Performance Guide*.

Collector metrics

The following metrics are Collector metrics.

Collector Metrics Received Per Interval

The Collector Metrics Received Per Interval metric is an extremely simple way of gauging how much load metric data queries are placing on the cluster. This metric is the total sum of Collector metric data points that the MOM has received each 15-second time period, including data queries. You can find the Collector Metrics Received Per Interval metric here in the Investigator tree:
`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | MOM | Collector Metrics Received Per Interval`

A large Collector Metrics Received Per Interval metric value, coupled with degradation of the cluster, indicates that the MOM has been asked to read too much metric data from the Collectors.

For more information about this metric, see the *APM Sizing and Performance Guide*.

Collector | <Collector name> | Skew Time (ms)

Indicates the clock skew for a specific Collector.

Collector Number of Async Queries per Interval

The total number of asynchronous queries for all Collectors.

Collector Async Queries Duration (ms)

The total time taken for all asynchronous queries for all Collectors.

Collector Number of Sync Queries per Interval

The total number of synchronous queries for all Collectors.

Collector Sync Queries Duration (ms)

The total time taken for all synchronous queries for all Collectors.

CA Technologies Internal Documentation
Authorized Use Only.

Collector Number of Sync Queries by CLW per Interval

The total number of synchronous queries for all Collectors requested from all Command Line Workstations.

EM Live Metric Clamped

Indicates if the number of live metrics handled by Enterprise Manager is less than or greater than the maximum limit specified in the introscope.enterprisemanager.metrics.live.limit property for Enterprise Manager clamps. The metric value is 0 if the number of live metrics for the Enterprise Manager is less than the specified limit. The metric value is 1 if the number of live metrics for the Enterprise Manager is greater than the specified limit.

Note: You can define the clamp limit for introscope.enterprisemanager.metrics.live.limit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the <*EM_Home*>\config directory.

EM Historical Metric Clamped

Indicates if the number of live metrics handled by Enterprise Manager is less than or greater than the maximum limit specified in the introscope.enterprisemanager.metrics.historical.limit property for Enterprise Manager clamps. The metric value is 0 if the number of live metrics for the Enterprise Manager is less than the specified limit. The metric value is 1 if the number of live metrics for the Enterprise Manager is greater than the specified limit.

Note: You can define the clamp limit for introscope.enterprisemanager.metrics.historical.limit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the <*EM_Home*>\config directory.

Max Number of Agent Connection Limit Exceeded Per Interval

The number of times the clamp limit defined in the introscope.enterprisemanager.agent.connection.limit property was exceeded for a particular interval.

Note: You can define the clamp limit for introscope.enterprisemanager.agent.connection.limit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the <*EM_Home*>\config directory.

Number of Events Processed

Indicates the total number of all events such as Transaction Traces and errors that the Enterprise Manager processes in each interval.

CA Technologies Product Documentation. Authorized Use Only.

Number of Events Processed Limit Exceeded

The number of times the clamp limit defined in the introscope.enterprisemanager.events.limit property was exceeded for a particular interval.

Note: You can define the clamp limit for introscope.enterprisemanager.events.limit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the <*EM_Home*>\config directory.

Number of Collector Metrics

The Number of Collector Metrics metric shows the total number of metrics currently being tracked in the cluster. You can find the Number of Collector Metrics metric here in the Investigator tree:
`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | MOM | Number of Collector Metrics.`

For more information about this metric, see the *APM Sizing and Performance Guide.*

Query metrics

Data Points Retrieved From Disk Per Interval

The number of data points retrieved from SmartStor per interval.

Data Points Returned Per Interval

The number of data points that the Enterprise Manager returned to clients per interval.

Metrics Read From Disk Per Interval

Number of metrics read from SmartStor per interval.

Metrics Returned Per Interval

The number of unique metrics that the Enterprise Manager returned to clients.

Queries Exceeding Max Data Points Read From Disk Limit Per Interval

Indicates if the maximum number of metric data points specified in the introscope.enterprisemanager.query.datapointlimit property that an Enterprise Manager can retrieve for a particular CLW or JDBC query was exceeded. The metric value is 0 if the number of metric data points returned by the Enterprise Manager is less than the specified limit. The metric value is 1 if the number of metric data points returned by the Enterprise Manager is greater than the specified limit.

CA Technologies Prerelease Documentation. Authorized Use Only

Note: You can define the clamp limit for introscope.enterprisemanager.query.datapointlimit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the <EM_Home>\config directory.

Queries Exceeding Max Data Points Returned Limit Per Interval

Indicates if the maximum number of metric data points specified in the queryintroscope.enterprisemanager.query.returneddatapointlimit property that an Enterprise Manager can return for a particular CLW or JDBC query was exceeded. The metric value is 0 if the number of metric data points returned by the Enterprise Manager is less than the specified limit. The metric value is 1 if the number of metric data points returned by the Enterprise Manager is greater than the specified limit.

Note: You can define the clamp limit for introscope.enterprisemanager.query.returneddatapointlimit property in the IntroscopeEnterpriseManager.properties file. The IntroscopeEnterpriseManager.properties file is located in the *<EM_Home>*\config directory.

Converting Spool to Data metric

The Converting Spool to Data metric tracks whether or not the spool to data conversion task is running. You can find this metric at the following location in the Investigator tree: `Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Data Store | SmartStor | Tasks | Converting Spool to Data`

If this metric stays at a value of 1 for more than 10 minutes per hour, this indicates that reorganizing the SmartStor spool file is taking too long. For more information on this metric, see the *APM Sizing and Performance Guide*.

Overall Capacity (%) metric

The Enterprise Manager Overall Capacity (%) metric estimates the percentage of the Enterprise Manager's capacity that is consumed. You can find it at this location in the Investigator tree:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager: Overall Capacity (%)`

For more information on this metric, see the *APM Sizing and Performance Guide*.

CA Technologies Internal/Prerelease Documentation Use Only

SmartStor Capacity (%) metric

The SmartStor Capacity (%) metric displays the percent of time needed for the SmartStor write process in a 15-second time slice, where 15 seconds equals 100%.

You can find it at this location in the Investigator tree:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Health | SmartStor Capacity (%)`

For more information on this metric and on SmartStor, see the *APM Sizing and Performance Guide*.

Heap Capacity (%) metric

The Heap Capacity (%) metric is determined by what percentage of heap the JVM is currently using (based on the GC Heap: In Use Post GC (mb) metric).

For more information on this metric, see the *APM Sizing and Performance Guide.*

Write Duration (ms) metric

The Write Duration (ms) metric diplays the duration, in milliseconds, of the SmartStor write process. This is the integer version of the SmartStor Capacity metric (see above). You can find it in this location in the Investigator tree:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Data Store | SmartStor | MetaData | Write Duration (ms)`

Number of Agents metric

This metric displays the number of currently connected agents. It is located in:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Connections`

Number of Metrics metrics

This metric displays the total metric load on the Enterprise Manager. It is located in:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Connections`

CA Technologies Pre-release Documentation. Author Use Only.

Historical Metric Count metric

The Historical Metric Count metric shows the total number of metrics from an agent that are live or recently active. It is located in the Investigator tree in:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)| Agent | Historical Metric Count`

Number of Historical Metrics metric

The Number of Historical Metrics metric displays the total number of metrics an Enterprise Manager is tracking across all agents. You can find this metric at the following location in the Investigator tree:

`Custom Metric Host (Virtual)| Custom Metric Process (Virtual)| Custom Metric Agent (Virtual)(*SuperDomain*)| Enterprise Manager | Connections | Number of Historical Metrics.`

For more information on this metric, see the *APM Sizing and Performance Guide*.

CA Technologies Prerelease Documentation. Authorized Use Only

Appendix B: Introscope Extensions

The following CA APM extensions, formerly separate products, are installed automatically by the CA APM installer.

Each extension must be enabled and configured before you can use it.

This section contains the following topics:

SNMP Adapter (see page 403)
ErrorDetector (see page 406)

SNMP Adapter

SNMP (Simple Network Management Protocol) is a standard protocol for monitoring and controlling network components. The SNMP Adapter enables Introscope to report metrics to external SNMP frameworks.

SNMP enables Introscope users to collect metric data and report it to an external SNMP Manager, where it can be viewed using the SNMP Manager Console.

The Introscope SNMP Adapter works out-of-the box with HP OpenView, BMC Patrol, and HP LoadRunner.

SNMP Adapter is installed automatically with Introscope. After installation, you must enable SNMP in the *IntroscopeEnterpriseManager.properties* file. See the *APM Configuration and Administration Guide* for information on configuring SNMP functionality.

Creating an SNMP collection

An SNMP collection defines which metrics will be published to the MIB.

An SNMP collection is an Introscope Workstation element and is stored in a Management Module. You create an SNMP collection using the Management Module Editor in the Introscope Workstation.

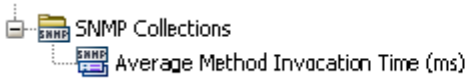
You can create and manipulate SNMP collections any time, after performing the configuration steps described in the *CA APM Configuration and Administration Guide.*

CA Technologies - Prerelease Documentation. Authorized Use Only.

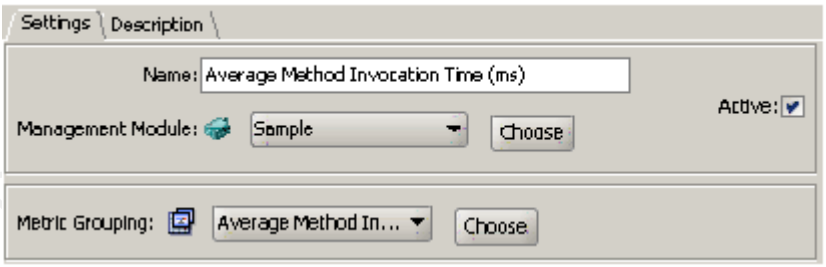
To create an SNMP collection for a metric:

1. Start the Workstation.
2. Open a Management Module Editor or Investigator window, and expand the tree to show metrics.
3. Select a metric grouping in the Management Module Editor, or a metric in the Investigator tree.
4. Right-click on the metric grouping or metric for which to create an SNMP collection.
5. Select New SNMP Collection from... from the drop-down menu.
6. Accept the default name for the SNMP collection in one of these ways:
   - Select a Management Module for the SNMP collection from the drop-down list.
   - Click Choose, choose a Management Module from the list, then click Choose.
7. Click OK.

   Now you activate the SNMP Collection.
8. Locate your new SNMP Collection in the Management Module Editor, under the Management Module where you saved it.



9. In the Settings tab for the SNMP Collection, check the Active box to activate the collection.



   Note: A metric grouping is automatically created when you create an SNMP collection from a metric. It belongs to the same Management Module as the one in which you saved the SNMP collection.
10. Click Apply.

Repeat the above steps to define additional SNMP collections.

Publishing a MIB

After you define SNMP collections, you publish a MIB from the Workstation to capture the metric data stored in the SNMP Collections. The Publish MIB command takes a snapshot only of current SNMP Collections, and publishes only metrics that are currently reporting. If you select an SNMP collection in the Management Module Editor tree, you can view the metrics that currently match. When you use the Publish MIB command, those metrics are published.

Metrics published to the MIB are assigned a unique Object Identifier (OID). SNMP Managers use OIDs to reference individual metrics in the MIB. A metric's OID does not change when the MIB is re-published, as long as the metric appears in an SNMP Collection each time.

If you publish a MIB without the metric, however, the guarantee of OID stability is lost. To ensure that the metric OIDs are stable, ensure that all metrics are reporting before you publish the MIB.

Ensuring that all metrics are reporting is especially important after restarting the Enterprise Manager. You must wait for the agents to connect to the Enterprise Manager and for the metrics to appear in the Workstation before you re-publish the MIB.

If an SNMP collection has no matching metrics when a MIB is published, but has matching metrics after publishing that you want published using SNMP, you must republish the MIB.

You can publish a MIB at any time, but be aware that:

- the old MIB file in the Introscope directory will be overwritten.
- you must load the new MIB file into your SNMP Manager, because the old MIB no longer contains current information.

Note: You will not be able to publish a MIB without installing the *SNMPAdapter.jar* file and enabling SNMP in the *IntroscopeEnterpriseManager.properties* file.

To publish a MIB:

1. Log into the Workstation as a user with publish_mib access to the server.
2. In the Workstation, open the Management Module Editor.
3. Select Manager > Publish MIB.
4. In the MIB Type area, check the box of the MIB type to publish. You can publish multiple MIB types.
5. Select a time period from the drop-down menu, to specify how often values reported by the SNMP agent are updated. The default period is one minute.
6. Click Publish.

CA Technologies Pre-Release Documentation. Authorized Use Only.

About MIB files

MIB files are named according to the host, and are saved in this directory:

<EM_Home>/snmp/

An MIB file is published and uniquely named using this syntax:

- HP LoadRunner: Introscope-*<hostname>.lr.mib*
- HP OpenView: Introscope-*<hostname>.mib*
- BMC Patrol
  - Introscope-*<hostname>.bmc.mib*
  - Introscope-*<hostname>.bmc.smi*
- Other: Introscope-*<hostname>.other.mib*

Any illegal characters in the host name are removed; periods in the host name are converted to dashes.

ErrorDetector

ErrorDetector is a formerly separate Introscope extension which is fully integrated with Introscope.

Introscope includes a ProbeBuilder Directive (PBD) file called *errors.pbd* with the agent installation, and collects error information as defined in the *errors.pbd* file. To use this file to enable and configure error metrics, see the *CA APM Java Agent Implementation Guide*.

When ErrorDetector is configured and enabled, Introscope allows application support personnel to detect and diagnose the cause of serious errors as they occur in live applications, determine the frequency and nature of the errors, and deliver specific information about the root cause to developers.

Some examples of common errors are:

- HTTP errors (404, 500, etc.)
- SQL statement errors
- network connectivity errors (timeout errors)
- backend errors (e.g., can't send a message through JMS, can't write a message to the message queue).

CA Technologies Prerelease Documentation. Authorized Use Only.

Introscope identifies these "serious" errors based on information contained in the J2EE and .NET specifications.

Note: Occasionally, HTTP 404 errors originate in a web server instead of an application server. If this occurs, ErrorDetector will not detect the web server error through the agent.

Introscope considers both errors and exceptions to be errors. The most common type of error is a thrown Java exception.

Reading and understanding error metrics

From the Workstation:

- You can view error metric data in the Investigator.
- You can view live errors in the Live Error Viewer.
- You can view error details in the Error Snapshot, which shows component-level information on how the error occurred.

Viewing error metrics in the Investigator

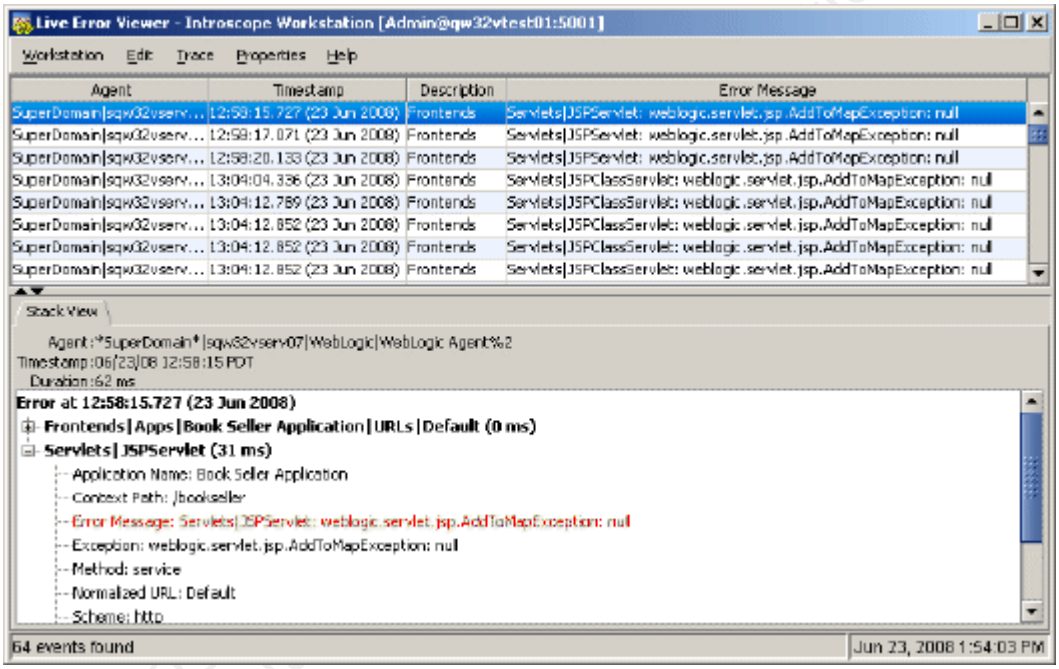
The *errors.pbd* file generates *Errors Per Interval* metrics that appear under several of the default resources. Local Product produces Errors Per Interval metrics for J2EE resources such as J2EE connectors, servlets, JTA, and HTTP; as well as for .NET framework resources such as ASP.NET pages, ADO.NET data sources, messaging queues, web mail, and enterprise services transactions.

CA Technologies Internal/Prerelease Documentation.
Authorized Use Only

Viewing error data in the Live Error Viewer

To view live, currently occurring errors, select Workstation > New Live Error Viewer from a Workstation Console or Investigator window. The Live Error Viewer has two parts:

- The Error Data Table, in the top part of the Live Error Viewer, lists errors currently occurring.
- The Error Snapshot, in the lower part of the Live Error View shows the details for the error currently selected in the Error Data Table.



Note: Due to the way the Oracle application server handles non-JSP requests, such as requests to .html .gif and other non-active resources, these requests are not traced by blamepoint tracers, so no error snapshot will be created for these errors.

Error data table

The Error Data Table displays information about current errors. This table shows the information displayed for each error.

Column Name	Information
Agent	Agent name
Timestamp	Start time (based on the system's clock) of the invocation of the root component

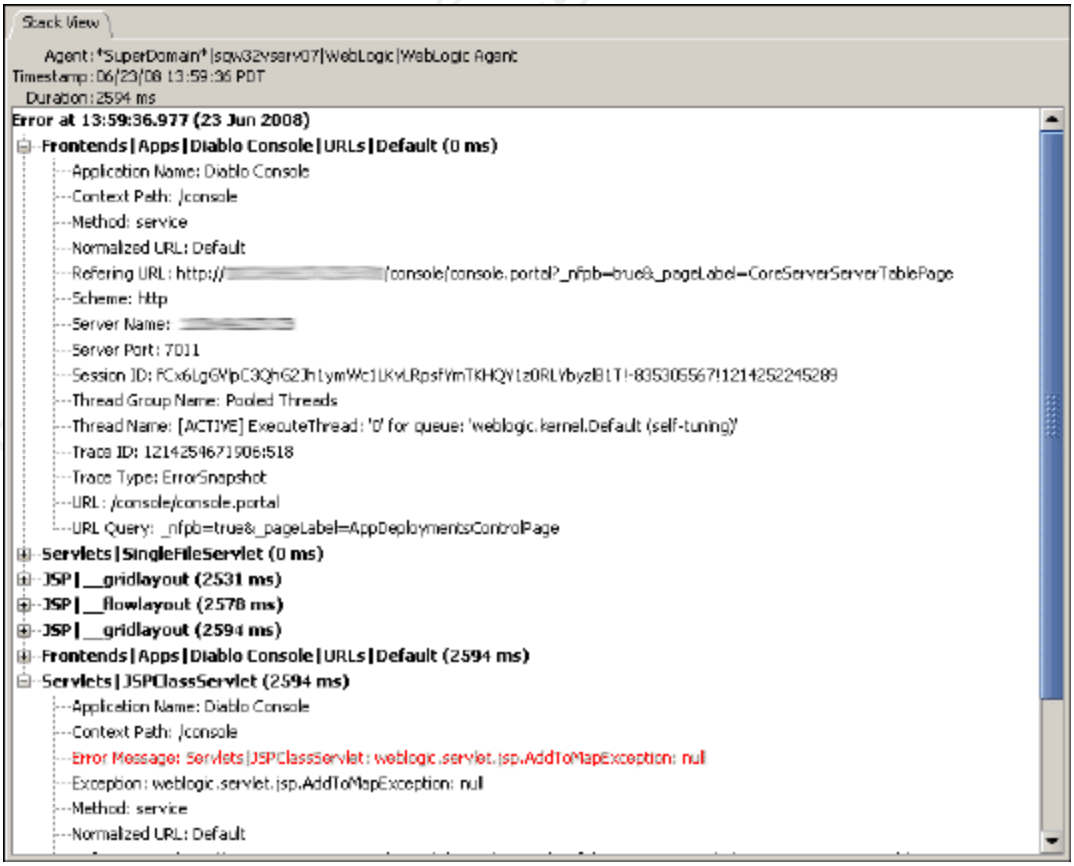
Description	Type of component of the error. This maps to the first segment of the component's resource name: for standard J2EE Blamed metrics, examples include Servlets, JSP, EJB, JNDI, etc.; for .NET components examples include Messaging and WebMail. For custom tracer implementations, the category matches the first segment in the blamed method's metric resource. If the metric resource has zero segments, the Description maps to "Custom Tracer".
Error Message	Exact error message captured.

- Click on a row to display an Error Snapshot for that error in the lower pane.
- Click on a column header to sort the rows by contents of that column. As new errors occur, they appear in sort order.

Error Stack View pane

Selecting an error in the Error Table pane causes its Error Stack View to appear in the lower pane.

The error message appears in red. The tree shows where the error occurred in the component trace. Components are in shown in bold, followed by component data.



Note: You can copy a line of the Error Snapshot to include in an e-mail, report, or text message. To copy an error, click to select it, then use Ctrl + C to copy it.

Viewing and analyzing historical error data

The Transaction Event Database contains error and transaction trace data captured by the agent. You can view and analyze error information in the Transaction Event Database by querying for errors based on error attributes and text. You can expand your analysis by querying for errors that are similar to, or correlated to, a selected error. To query the Transaction Event Database, see Querying stored events (see page 229).

View Errors Using the Transaction Tracer

Transaction Traces that contain errors appear in the Transaction Trace Viewer window.

Note: Occasionally, errors that have been tagged to be ignored appear in the Transaction Trace Viewer, but not in the Live Error Viewer. This can occur when errors in the application appear before the agent has processed the ignore command in the *IntroscopeAgent.profile*.

You view errors that occur within a transaction using the Transaction Trace Viewer. (For general information about using the Transaction Trace Viewer, see Using the Transaction Trace Viewer (see page 209))

To view errors within Transaction Trace Viewer:

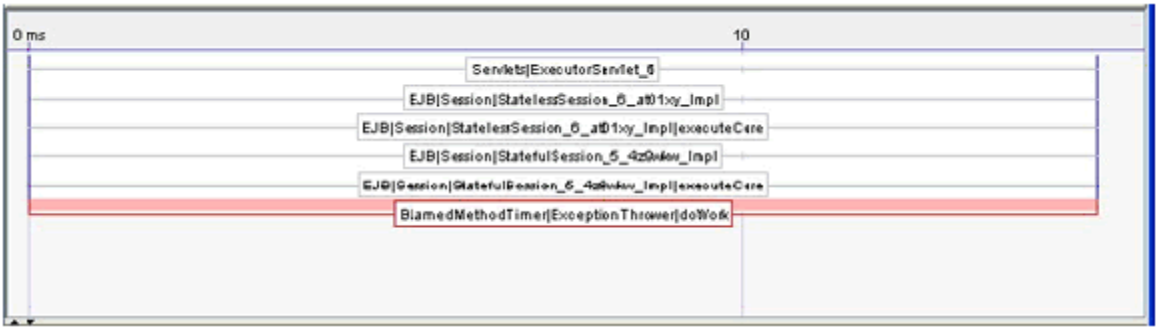
1. Select Workstation > New Transaction Trace Session.
2. In the Minimum transaction duration dialog, enter the length of the Transaction Trace session.
3. In the Trace Agents section, select an option to define which agents to trace during the Transaction Trace session.
4. In addition to other filtering options, select *Errors* and enter any error information to search on in the field.
5. Click OK.

   Transaction Traces that contain errors appear in the Transaction Trace Viewer window.

   Note: Occasionally, errors that have been tagged to be ignored appear in the Transaction Trace Viewer, but not in the Live Error Viewer. This can occur when errors in the application appear before the agent has processed the ignore command in the *IntroscopeAgent.profile*.

CA Technologies Proprietary Documentation
Authorized Use Only

When you select a transaction in the table, the Transaction Snapshot appears. If an error occurred within the transaction, the background of the particular transaction snapshot appears in red in the Transaction Trace window



6. Click on the error in the snapshot to display its details in the properties pane below the Transaction Snapshot.

Note: You can select the text of any field in the properties details and copy it using keyboard commands.

The default *errors.pbd* is designed to report serious errors, while minimizing overhead as much as possible. Overuse of error tracing, for instance, applying ExceptionErrorReporter to every monitored method can result in a high volume of "false positives." For example, if a user enters "California" in a numeric field, this may cause a *NumberFormatException*, which you would not want to report as a serious problem.

CA Technologies Pre-release Documentation
Authorized Use Only

Index

A

Absolute max • 243
Absolute min • 243
Action - as an element in alert objects • 314
Action - creating • 331, 332, 333
Alert indicators - on sample dashboards • 53
Alert indicators - on the map display • 81
Alert state metrics • 325
Alerts
  Alert notification • 173
  Alert notification - as an element in alert objects • 322
  Alert notification - responding • 176
  Alerts - creating and configuring • 309, 318
Application triage map • 84, 87, 90, 176
Average Respsonse Time • 351

B

By Frontend tree • 77, 90

C

calculators • 336, 338
Concurrent Invocations • 353
Console lens • 48
Console, about • 45

D

Dashboards, creating • 278
Dashboards, sample • 50
Data, agent-centric view • 72
Data, application centric and business centric views • 75

E

ErrorDetector • 406
Errors • 142, 354, 371
Errors Per Interval • 354

F

Filtering - in the Console • 48
Frontends map • 90
Frontends vs. backends • 90, 114

G

Graphs, changing and manipulating • 147

H

Historical data - in the Console • 64
Historical events - querying in Transaction Tracer • 230
Hyperlinks • 302

I

Instrumentation - see also Transaction Tracer • 217
  Instrumentation, converting temporary to permanent • 222
  Instrumentation, removing • 224
  Instrumentation, temporary • 219
Investigator
  Investigator - application triage map • 87
  Investigator - Map tab • 75
  Investigator, about • 30
  Investigator, agent-centric view • 72
  Investigator, application-centric and business-centric view • 75
  Investigator, general features • 69
  Investigator, using • 69

J

Java Web Start • 16
JavaScript calculators • 338

L

Launching using Java Web Start vs. WebView • 16
LeakHunter • 145

M

Management Module editor • 265
Management Modules
  Management Modules - about • 263
  Management Modules - deploying • 343
  Mangement Modules - creating • 264
Metric groupings • 273
Metrics
  Metrics - about • 346, 348, 349
  Metrics - frontends • 77, 170, 377, 378

CA Technologies Prerelease Documentation. Authorized Use Only.

Metrics - memory related • 358
MIB • 403

P

Performance - nominal • 163
Permissions • 130

R

Regular expressions • 231, 254
Reports • 239
Report templates • 239
Responses Per Interval • 355

S

Sample Console dashboards • 50
Search • 140, 230
SmartTrigger • 322
SNMP Adapter • 403
Stall Count • 357
Super Domain • 73
Supportability metrics • 74

T

Templates - for reports • 239, 259, 261
Thread dumps
tab • 144
viewing • 144, 183
Transaction Tracer
Transaction Tracer - about • 203
Transaction Tracer - starting/stopping a trace • 205
Transaction tracer - viewer • 209
Transactions, clamped - in a transaction trace • 214
Transactions, business • 79
Triaging • 176, 179

U

User permissions • 130

W

Workstation, as part of APM • 16

CA Technologies Prerelease Documentation. Authorized Use Only.