CA Application Performance Management

CA Technologies Prerelease Documentation.
Authorized Use Only.

Configuration and Administration Guide

Release 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 Documentation. Limited Use Only.

CA Technologies Product References

This document references the following CA Technologies products and features:

- CA Application Performance Management (CA APM)
- CA Application Performance Management ChangeDetector (CA APM ChangeDetector)
- CA Application Performance Management ErrorDetector (CA APM ErrorDetector)
- CA Application Performance Management for CA Database Performance (CA APM for CA Database Performance)
- CA Application Performance Management for CA SiteMinder (CA APM for CA SiteMinder)
- CA Application Performance Management for CA SiteMinder Web Access Manager (CA APM for CA SiteMinder Web Access Manager)
- CA Application Performance Management for CA SYSVIEW® (CA APM for CA SYSVIEW)
- CA Application Performance Management for IBM CICS Transaction Gateway (CA APM for IBM CICS Transaction Gateway)
- CA Application Performance Management for IBM WebSphere Application Server (CA APM for IBM WebSphere Application Server)
- CA Application Performance Management for IBM WebSphere Distributed Environments (CA APM for IBM WebSphere Distributed Environments)
- CA Application Performance Management for IBM WebSphere MQ (CA APM for IBM WebSphere MQ)
- CA Application Performance Management for IBM WebSphere Portal (CA APM for IBM WebSphere Portal)
- CA Application Performance Management for IBM WebSphere Process Server (CA APM for IBM WebSphere Process Server)
- CA Application Performance Management for IBM z/OS® (CA APM for IBM z/OS)
- CA Application Performance Management for Microsoft SharePoint (CA APM for Microsoft SharePoint)
- CA Application Performance Management for Oracle Databases (CA APM for Oracle Databases)
- CA Application Performance Management for Oracle Service Bus (CA APM for Oracle Service Bus)
- CA Application Performance Management for Oracle WebLogic Portal (CA APM for Oracle WebLogic Portal)

CA Technologies Pre-release Documentation. Limited to internal use only.

- CA Application Performance Management for Oracle WebLogic Server (CA APM for Oracle WebLogic Server)
- CA Application Performance Management for SOA (CA APM for SOA)
- CA Application Performance Management for TIBCO BusinessWorks (CA APM for TIBCO BusinessWorks)
- CA Application Performance Management for TIBCO Enterprise Message Service (CA APM for TIBCO Enterprise Message Service)
- CA Application Performance Management for Web Servers (CA APM for Web Servers)
- CA Application Performance Management for webMethods Broker (CA APM for webMethods Broker)
- CA Application Performance Management for webMethods Integration Server (CA APM for webMethods Integration Server)
- CA Application Performance Management Integration for CA CMDB (CA APM Integration for CA CMDB)
- CA Application Performance Management Integration for CA NSM (CA APM Integration for CA NSM)
- CA Application Performance Management LeakHunter (LeakHunter)
- CA Application Performance Management Transaction Generator (CA APM TG)
- CA Customer Experience Manager (CA CEM)
- CA Embedded Entitlements Manager (CA EEM)
- CA eHealth® Performance Manager (CA eHealth)
- CA Insight™ Database Performance Monitor for DB2 for z/OS
- CA Introscope® (CA Introscope)
- CA SiteMinder® (SiteMinder)
- CA Spectrum® Infrastructure Manager (CA Spectrum)
- CA SYSVIEW® Performance Management (CA SYSVIEW)

CA Technologies Prerelease Documentation. Authorized Use Only.

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.

Documentation Changes

The following documentation updates have been made since the last release of this documentation:

The contents of this guide was formerly in two guides:

- *CEM Configuration and Administration Guide*
- *Introscope Configuration and Administration Guide*

In addition, these chapters have been moved from the deprecated *CA Wily CEM Integration Guide*:

  - Configuring CA Introscope to work with CA CEM (see page 349)
  - Integrating CA SiteMinder with CA CEM (see page 387)
  - Integrating CA Service Desk with CA CEM (see page 393)
  - CA Introscope agent profile properties

The chapter about Using Virtual Agents to Aggregate Metrics (see page 97), which has information about configuring and using virtual agents, was moved to this guide. This content was previously available in the *Introscope Java Agent Guide*.

Content about agent load balancing formerly in other chapters and the APM Sizing and Performance Guide is now consolidated in a new chapter: Configure Agent Load Balancing.

You can configure the loadbalancing.xml file using the latched property to assign agents to preferred Collectors.

You can plan and carry out how agent connections are carried out across your CA APM environment. This capability is named Configuring Agent to Enterprise Manager Network Topology.

A new apm-events-thresholds.xml (see page 447) file contains new and existing CA Introscope and CA CEM clamp and new CA APM Important Event properties that you can configure (see page 217). You can use these clamps to limit key Enterprise Manager workload factors.

- When these clamps meet or exceed the configured threshold values, the Enterprise Manager reports alerts to the APM Status console Clamps tab.
- When the important events meet or exceed the configured threshold values, the Enterprise Manager reports alerts to the APM Status console Important Events table view.

CA Technologies Internal Use Documentation. Authorized Use Only.

These Enterprise Manager clamp properties have been moved from the IntroscopeEnterpriseManager.properties (see page 403) file to the apm-events-thresholds.xml (see page 447) file:

- introscope.enterprisemanager.agent.connection.limit (see page 448)
- introscope.enterprisemanager.agent.error.limit (see page 449)
- introscope.enterprisemanager.agent.metrics.limit (see page 449)
- introscope.enterprisemanager.agent.trace.limit (see page 450)
- introscope.enterprisemanager.metrics.live.limit (see page 453)
- introscope.enterprisemanager.events.limit (see page 452)
- introscope.enterprisemanager.metrics.historical.limit (see page 453)
- introscope.enterprisemanager.transactionevents.storage.max.disk.usage (see page 453)

These CA CEM clamp properties have been moved from the tess-defaults.properties file to the apm-events-thresholds.xml (see page 447) file and renamed:

- maxNumberDomainConfigurationChanges was renamed as introscope.enterprisemanager.max.number.domain.configuration.changes (see page 456)
- Autogen.Max.NonIdentifying.Components was renamed as introscope.enterprisemanager.transaction.discovery.max.nonidentifying.components (see page 457)
- maxUserGroups was renamed as introscope.enterprisemanager.max.transaction.user.groups (see page 456)
- maxUserRows was renamed as introscope.enterprisemanager.max.application.user.rows (see page 455)

The introscope.enterprisemanager.query.datapointlimit (see page 414) and introscope.enterprisemanager.query.returneddatapointlimit (see page 415) properties, found in the IntroscopeEnterpriseManager.properties (see page 403) file, are now only for CLW and JDBC query performance tuning. They are no longer used as general metric data clamps due to Enterprise Manager query improvements and clamps in the apm-events-thresholds.xml (see page 447) file.

These Command Line Workstation (CLW) commands have been updated or added:

- Triage Map commands
- APM Status console commands (see page 148)
-

The following sections provide new or changed features in CA Application Performance Management (CA APM) 9.1.0 that affect the *CA APM Configuration and Administration Guide*:

CA Technologies Pre-release Documentation. Authorized Use Only

Faster processing of CA CEM defects, logins, and statistics

The processing rate of CA CEM defects, logins, and statistics has been improved. This means that in CA APM, you can:

- Process more defects without the Enterprise Manager falling behind or truncating data
- Monitor more transactions in an CA APM cluster because CA CEM can process data from more TIMs
- Aggregate statistics for larger numbers of users and groups of users

Processing is done as collection tasks, separated into processing pools. There are three processing pools, one for transaction discovery and customer experience metrics, one for defects and logins, and one for events and recordings. This separation prevents the processing of one type of data from starving the processing of another type.

You can use *IntroscopeEnterpriseManager.properties* properties to enable and configure thread dumps. See Thread Dumps (see page 439).

You can export CA CEM data from the Export Data tab. This tab lets you export the business statistics, defects, and incidents data for a specified time range.

Import the resulting CSV (comma-separated value) data into any spreadsheet or reporting tool such as Microsoft Excel, SAP Business Objects, SAS, or Crystal Reports.

Now, you can now export CA CEM data (see page 329) in the following ways:

- Use the Export Data tab in the CEM console.
- Access the CEM data export tool from the command line, as in previous releases.

CA Technologies Pre-release Documentation. Authorized Use Only.

Contents

Chapter 1: Introducing CA APM Configuration and Administration Guide 25

About APM 25
About this guide and other CA Wily guides 26
What you need to know 26
Product Documentation 27
CA APM documentation 29
Introscope Documentation 30
CA CEM Documentation 31
Introscope PowerPack documentation 31

Chapter 2: Starting and Stopping the Enterprise Manager 33

Starting the Enterprise Manager on Windows 33
Starting the Enterprise Manager Manually on Windows 33
Running the Enterprise Manager as a Windows Service 34
Starting the Enterprise Manager on UNIX 35
Starting the Enterprise Manager Manually on UNIX 36
Running Enterprise Manager control scripts on UNIX 36
Running Enterprise Manager in nohup Mode on UNIX 37
Starting the Enterprise Manager on OS/400 37
Running the OS/400 start script 37
Using control scripts for OS/400 and Z/OS 38
Running the Enterprise Manager as a Java program 38
Starting and configuring the Enterprise Manager on z/OS 39
Prerequisites for running Enterprise Manager as a z/OS batch job 39
Configuring Enterprise Manager to run as a z/OS batch job 42
Running the Enterprise Manager on z/OS 44
Verifying the Enterprise Manager is running 46
Starting the Enterprise Manager automatically 46
Configuring the Enterprise Manager to use a different JVM 47
About the Enterprise Manager log file 47
Rolling Enterprise Manager logs by date or size 48
Configuring the Enterprise Manager config directory 49
Stopping the Enterprise Manager 50
Stopping the Enterprise Manager from the terminal window 51
Stopping the Enterprise Manager from the Workstation 51
Automatic shutdown of Enterprise Manager 51

CA Technologies Pre-release Documentation
Authorized Use Only.

Restarting after an unplanned shutdown 52
Enterprise Manager shell and terminal window commands 53

Chapter 3: Configuring Enterprise Manager Communications 55

The default Enterprise Manager communications channel 55
Rebinding a communication channel to a port in use 56
Configuring SSL 56
Connecting to the Enterprise Manager across a firewall 58
Enabling SSL for Enterprise Managers on OS/400 59
Configuring the Enterprise Manager web server for HTTPS 59
Understanding how HTTPS works 60
Default Jetty configuration file 60

Chapter 4: Configuring Enterprise Manager databases 63

Administering the APM database 63
Backing up or restoring the database 63
Exporting or importing database configuration 69
Moving the APM database to another operating system 75
Migrating data from a PostgreSQL database to an Oracle database 75
Configuring SmartStor data storage 77
Determining SmartStor system requirements 79
Specifying SmartStor data tier values 79
Configuring Transaction Event database location and aging 80
Configuring location of baseline storage 81
Backing up SmartStor, baselines, and Transaction Events data 81
Creating a query-only duplicate cluster 83
Using SmartStor tools to tune SmartStor data 83
SmartStor tools commands 84
merge 85
remove_metrics 86
remove_agents 87
keep_agents 88
prune 88
test_regex 89
list_agents 90
Pruning SmartStor metadata 90
Trimming SmartStor data and metadata 92
Example: Using the test_regex command to find a SmartStor data problem 94

CA Technologies Prerelease Documentation
Authorized Use Only

Chapter 5: Using Virtual Agents to Aggregate Metrics 97

Understanding Virtual Agents 97
Virtual Agent requirements....97
Configuring Virtual Agents 98

Chapter 6: Working with Enterprise Manager Clustering 101

Understanding Enterprise Manager clustering 101
When to cluster Enterprise Managers 101
Understanding how MOMs manage Collectors 102
Clustering and Enterprise Manager clock synchronization 102
Clustering and agent naming 103
Clustering and Management Modules....103
Using Enterprise Manager hot failover 104
Enterprise Manager hot failover rules 105
Configuring Enterprise Manager failover to work on a single host 107
About MOM failover and CA CEM TIMs....107
Configuring a cluster of Enterprise Managers....109
Configuring the MOM 109
Configuring a Collector....110

Chapter 7: Configure MOM Agent Load Balancing 111

How the MOM Assigns Agents to Collectors....112
How the MOM Rebalances the Cluster 112
Set the Agent Load Balance Metric Weighting Factor 113
Set the Agent Load Balance Metric Threshold 114
Set the Agent Load Balance Interval Property 115
Agent Load Balancing Usage Scenarios 115
Agent Load Balancing When a New Collector is Added to the Cluster 115
Agent Load Balancing When a Collector with Agents is Added 116
Agent Load Balancing When a Collector Fails 116
Agent Load Balancing When a Collector Restarts 117
Agent Load Balancing with Weights....117
Load Balancing When Agents Prefer a Specific Collector 117
Configure the Assignment of Agents to Collectors 118
Agent Load Balancing Configuration Examples....120

Chapter 8: Configure Agent - Enterprise Manager Network Topology 125

What Happens When an Agent is Disconnected from an Enterprise Manager 127
What Happens When an Agent is in a Denied State 128

CA Technologies Prerelease Documentation. Authorized Use Only

How Agent Connection Information Propagates Across CA Introscope 128
Process of Configuring Agent - Enterprise Manager Network Topology 129
Plan the Agent - Enterprise Manager Network Topology 130
Configure if New Agent Connections are Allowed When No loadbalancing.xml Rules 132
Configure Agent Reconnection Wait Time 133
Configure Disallowed Agent Connection Clamp Limit 133
Configure Agent Continual Reconnect to Any Allowed Enterprise Manager 134
Configure loadbalancing.xml for Allowed and Disallowed Agents by Enterprise Manager 136
Agent Connection Scenarios 136
Troubleshoot Agent - Enterprise Manager Network Topology 137
How Enterprise Manager Handles Disallowed Agents 137
When Collector Uses introscope.apm.agentcontrol.agent.allowed Property 137
Agent Does Not Match loadbalancing.xml Regular Expression as Allowed or Disallowed 138
All Collectors Assigned to an Agent are Unavailable and Other Collectors are Available 138
How Collector Handles a Disallowed Agent with a Direct Connection 138
Default Agent Connection Mode When introscope.apm.agentcontrol.agent.allowed Property is Missing 139
How CA Introscope Determines Where Agents Connect When No MOM Available 139

Chapter 9: Using the Command-Line Workstation 141

About the Command Line Workstation (CLW) 141
Running the CLW 142
JVM requirements 142
Where to run CLW commands 142
Using CLW interactive mode 143
Invoking CLW 143
CLW is case-sensitive 144
CLW command output 144
CLW and regular expressions 144
Backslashes and quotes in CLW commands 144
Using CLW in different environments 144
Configuring the Command Line Workstation Log 145
CLW command reference 145
Enterprise Manager commands 145
Introscope agent commands 147
APM Status Console Commands 148
Management Module commands 153
Action Commands 155
Alert commands 157
Alert Downtime Schedule commands 159
Calculator commands 161

CA Technologies Prerelease Documentation
Authorized Use Only

Dashboard commands 163
Metric Grouping commands 165
Report Templates Commands 166
SNMP Collections commands 169
Configuring Transaction Trace Options 172
Starting a Transaction Trace session 172
Application Triage Map Commands 179
Sample scripts 186
Transaction Trace Session 186
Transaction Trace session triggered by Alert 190

Chapter 10: Using the WatchDog Command-Line Tool 191

About WatchDog 191
Running WatchDog 192
Running WatchDog on SAP 193
WatchDog commands 193
WatchDog start command 194
WatchDog stop command 195
WatchDog status command 196
WatchDog command options 197
WatchDog -watch command option 197
WatchDog -port command option 198
WatchDog -startcmd command option 199
WatchDog -interval command option 199
WatchDog -startuptime command option 200
WatchDog -jre command option 200
WatchDog -emport command option 201
WatchDog -emuser command option 201
WatchDog -empwd command option 202
Viewing WatchDog log files 202
Sample scripts 202

Chapter 11: Configuring WebView Options 203

Configuring WebView for Non-ISO-8859-1 characters 203
Configuring the WebView application context path 203
Running WebView on z/OS 204
Before you begin 204
Configuring WebView to run as a z/OS batch job 206
Running WebView from a shell script 210

CA Technologies Prerelease Documentation
Authorized Use Only.

Chapter 12: Configuring the Workstation 211

Running Workstation in verbose mode 211
Redirecting Workstation output to a file 211
Configuring Workstation to supply login values 212
Configuring the Workstation for Asian-language reports 212

Chapter 13: System Configurations 213

Configuring the clock skew tolerance 213
Setting the Maximum Data Point Retrieval Property 215
Limiting the data points returned by a query 216
Increasing the number of query results returned 217
Configure APM Status Console Clamps and Important Events 217

Chapter 14: Configuring the Application Triage Map 219

The application triage map 219
Transaction sampling 220
Map element aging 220
Application triage map data clamping 223
Application triage map data pruning 224

Chapter 15: Configuring APM to Load Catalyst Data 227

Configure CA APM to Load CA Catalyst Data 228
Configure the CatalystPolicy.xml File 229

Chapter 16: Introducing CA CEM 231

Accessing the CEM console and setup pages 231
Access the CEM Console 231
Troubleshooting Login to CA CEM 233
Access the TIM System Setup Page 233
Access the TIM from the CEM Console 234
CA CEM administrator responsibilities 235
Daily tasks 235
Weekly tasks 235
Monthly tasks 236
Quarterly tasks 236

Chapter 17: Networking Alternatives for CA CEM 239

Process for CA CEM networking 239

CA Technologies Prerelease Documentation. Authorized Use Only.

About CA CEM network options 239
About monitoring network traffic 240
Using network taps 240
Using mirrored ports 241
About CA CEM insertion locations 241
Between the Internet and the Internet router 243
Between the Internet router and the firewall 243
Between the firewall and the aggregation switch 244
Between the aggregation switch and the network load balancer 244
Between the network load balancer and the web server 245
Connections from browser to the Enterprise Manager 245
IPv6 and CA CEM 245
About IPv6 246
IPv6 in CA CEM 246
Monitoring IPv6 network traffic 246
Running CA CEM on dual-stack (IPv6 and IPv4) hosts 247
Allowed formats for IPv6 addresses 248

Chapter 18: Administering CA CEM 251

Process for CA CEM administration 251
Distributing Enterprise Manager services 252
Typical distributions of the Enterprise Manager services 254
Distributing Enterprise Manager services in a clustered environment 255
Selecting the TIMs for real-time transaction metrics 257
Selecting the TIMs for automatic transaction discovery 257
Modifying the Enterprise Manager configuration 258
Removing an Enterprise Manager 259
Troubleshooting Enterprise Manager services 260
Defining your CA CEM domain 262
Incident cost versus business impact 268
Defining incident-related settings 270
Configuring incident management settings 274
Setting incident generating rules 275
Setting business impact and collecting evidence for incidents 276
Arguments for use in evidence collection scripts 278
Configuring incident age-out rules 279
Sending emails when incidents open or change severity 280
Configuring SMTP server settings 281
Setting the security groups for new business services 283
Enabling monitors 287
Synchronizing monitors 288

CA Technologies Prerelease Documentation
Authorized Use Only

Viewing configuration changes since last synchronization 289

Chapter 19: Managing users and groups on CA CEM 291

About user groups 291
The maximum number of user groups 293
Automatically assigning new users by request attributes 294
Automatically assigning new users to subnet user groups 295
Creating subnet user groups for new users 295
Identifying users behind proxy servers 296
Managing the new users group 298
Manually assigning new users to user groups 298
Configuring user groups 299
Configuring user settings with bulk editing 301
Configuring correlational SLAs 302
Viewing correlational SLAs 303
Adding correlational SLAs 303
Updating existing correlational SLAs 304
Deleting correlational SLAs 305
Deleting a user group 305
Deactivating or deleting a user 306
Reactivating a user 306

Chapter 20: Managing system operations on CA CEM 307

About CA CEM events 307
Severity levels 307
List of possible events 308
Viewing CA CEM events 311
Managing CA CEM events 311
Receiving email notification of events 312
Reviewing TIM machine health 314
Appliance Health pages reference table 315
Viewing CA CEM Audit Trails 316

Chapter 21: Reporting with CA CEM 319

Capturing CA CEM reports 319
About CA CEM Report Aggregation 319
Report delivery 320
Securing report data 320
Report formats 321
Report time frames 322

CA Technologies Prerelease Documentation. Authorized Use Only.

Exploring reports 323
PDF reports 323
Email reports 324
Saved or scheduled reports 324
Scheduling reports 325
Schedule a report using My Reports 325
Generate and schedule a report 328
Export CA CEM data 329
Export data 329
Save or schedule a report 332

Chapter 22: Monitoring Secure Web Applications With CA CEM 333

Importing and managing SSL private keys 333
Apache or OpenSSL-based web servers 334
Microsoft Internet Information Service (IIS) 334
SunONE, iPlanet, or Netscape enterprise server 335
Configuring CA CEM for SSL 338
About multiple keys for HTTPS servers 339
Verifying the functionality of CA CEM with SSL 339

Chapter 23: Using a Napatech Adapter with CA CEM 341

About using the Napatech adapter with CA CEM 341
Obtaining the Napatech adapter and software 342
Installing the Napatech adapter and Napatech software 342
Enabling the Napatech adapter and monitoring 343
(Optional) Creating and enabling Napatech filters 344
Other Napatech adapter tasks 345
Troubleshooting a Napatech adapter 347
The Napatech adapter is not receiving traffic 347
The Napatech adapter as an Ethernet port 348
Not seeing the traffic you expect 348
Serial number and other Napatech adapter details 348

Chapter 24: Configuring CA Introscope to work with CA CEM 349

Getting started with CA APM 349
CA APM Customer Experience Metrics Overview 350
CA APM problem resolution triage overview 350
Checklist for configuring CA CEM 352
Deciding which integration features you need 352
Configuring Introscope Settings 353

CA Technologies Prerelease Documentation
Authorized Use Only.

Configuring Customer Experience Metrics Integration 356
Verifying the CA CEM Configuration 362
Checklist for configuring Introscope 362
Deciding which integration features you need 363
Configuring the Enterprise Managers 363
Integrating with Java agents 365
Integrating with .NET agents 367
Verifying the Introscope configuration 371
Troubleshooting the CA Introscope to CA CEM Integration 371
Verifying CA CEM integration on CA Introscope 371
Verifying CA Introscope integration on CA CEM 373
Viewing transaction trace session information 374
Symptoms and solutions 374
Upgrade considerations 382
CA CEM domain configuration information for 4.0 through 4.5 releases 382
Setting up the business application 385

Chapter 25: Integrating CA SiteMinder with CA CEM 387

Checklist for CA SiteMinder 387
Configuring CA SiteMinder on CA CEM 388
Creating custom CA SiteMinder agents 389
Creating a custom agent type for CA CEM on CA SiteMinder 389
Defining policy servers on CA CEM 389
Defining agents on CA CEM 390
Enabling CA SiteMinder 390
Troubleshooting the CA SiteMinder to CA CEM integration 391

Chapter 26: Integrating CA Service Desk with CA CEM 393

About CA Service Desk with CA CEM 393
CA Service Desk integration overview 393
Configuring CA Service Desk for CA CEM 394
Configuring the CA Service Desk plug-in 394
Configuring a CA Service Desk request area 397
Using CA Service Desk with CA CEM 397
Troubleshooting the CA Service Desk to CA CEM integration 398
Symptoms and solutions 399
CA Service Desk to CA CEM integration terminology 400
CA Service Desk terms 400
CA Service Desk to CA CEM term mapping 401

CA Technologies Prerelease Documentation
Authorized Use Only

Appendix A: Properties Files Reference 403

IntroscopeEnterpriseManager.properties 403
Available Processors 404
Alert Behaviors 405
Communications 405
Hot Configuration Service 410
Application Heuristics 410
Logging 411
Persistent collections—recording data to a database 412
Flat File Properties 413
Enterprise Manager Metric Clamping 414
SmartStor 415
Transaction Event Database 419
Transaction Trace Sampling 420
Baseline Data Storage 421
SNMP Collections 421
WebView 422
Management Module Bootstrapping 422
Management Module Deployment 423
Clustering Settings 423
Agent interactions 427
Enterprise Manager hot failover 431
Configuration directory 432
Application triage map data 432
Location Map data 437
Web Services 437
Outgoing message queue 438
Thread dumps 439
SNMP Triage Map Alert Actions 441
Application Map Objects 443
APMEnterpriseManager.properties 443
c3p0 properties 443
Hibernate specific properties 445
apm-events-thresholds.xml 447
Introscope Clamp Properties 448
introscope.clw.max.users 448
introscope.enterprisemanager.agent.connection.limit 448
introscope.enterprisemanager.agent.error.limit 449
introscope.enterprisemanager.agent.metrics.limit 449
introscope.enterprisemanager.agent.trace.limit 450
introscope.cdv.max.users 450

CA Technologies Prerelease Documentation. Authorized Use Only

introscope.enterprisemanager.disconnected.historical.agent.limit 451
introscope.enterprisemanager.events.limit 452
introscope.enterprisemanager.metrics.historical.limit 453
introscope.enterprisemanager.metrics.live.limit 453
introscope.enterprisemanager.transactionevents.storage.max.disk.usage 453
introscope.workstation.max.users 454
Customer Experience Clamp Properties 455
Setting Important Event Thresholds 458
ResourceMetricMap.properties 461
CPU Utilization 463
Time Spent in GC 463
Threads in Use 463
Threads Available 463
JDBC Connections in Use 463
JDBC Connections Available 464
Catalyst.properties 464
catalyst.host 465
catalyst.protocol 465
catalyst.port 465
catalyst.login 465
catalyst.password 465
catalyst.rest.entrypoint 466
catalyst.rest.alerts.pollingperiodsec 466
catalyst.rest.virtualhosts.pollingperiodsec 466
catalyst.product.baseurl 467
catalyst.alert.filter.total 467
catalyst.alert.filter.time 467
catalyst.alert.filter.products 467
catalyst.alert.filter.types 468
catalyst.alert.filter.severities 469
catalyst.rest.maxFailures 469
catalyst.rest.maxAlertsPerPage 470
catalyst.rest.alerts.threads 470
catalyst.rest.virtualhosts.threads 470
catalyst.entity.cache.size 471
catalyst.entity.cache.expirationsec 471
catalyst.database.alerts.enabled 471
catalyst.database.map.name 472
EMservice.conf 472
wrapper.java.command 472
wrapper.java.mainclass 472
wrapper.java.classpath.x 472

CA Technologies Prerelease Documentation
Authorized Use Only

wrapper.java.library.path.1 473
wrapper.java.additional.x= 473
wrapper.java.initmemory 473
wrapper.java.maxmemory 473
wrapper.app.parameter.x 474
wrapper.startup.timeout 474
wrapper.console.format 474
wrapper.console.loglevel 474
wrapper.logfile 475
wrapper.logfile.format 475
wrapper.logfile.loglevel 475
wrapper.logfile.maxsize 475
wrapper.logfile.maxfiles 475
wrapper.syslog.loglevel 476
wrapper.ntservice.name 476
wrapper.ntservice.displayname 476
wrapper.ntservice.description 476
wrapper.ntservice.dependency.x 477
wrapper.ntservice.starttype 477
wrapper.ntservice.interactive 477
Introscope Enterprise Manager.lax 477
lax.application.name 477
lax.class.path 478
lax.command.line.args 478
lax.dir 478
lax.main.class 479
lax.main.method 479
lax.nl.current.vm 479
lax.nl.java.launcher.main.class 480
lax.nl.java.launcher.main.method 480
lax.nl.java.option.additional 480
lax.nl.message.vm.not.loaded 481
lax.nl.valid.vm.list 481
lax.nl.win32.microsoftvm.min.version 481
lax.root.install.dir 481
lax.stderr.redirect 482
lax.stdin.redirect 482
lax.stdout.redirect 483
lax.user.dir 483
lax.version 483
IntroscopeWorkstation.properties 483
log4j.logger.Workstation 484

CA Technologies Prerelease Documentation
Authorized Use Only

log4j.appender.logfile.File 484
introscope.workstation.graph.antialiased 484
transport.http.proxy.host 485
transport.http.proxy.port 485
transport.http.proxy.username 485
transport.http.proxy.password 485
transport.tcp.truststore 486
transport.tcp.trustpassword 486
transport.tcp.truststore 486
transport.tcp.keystore 486
transport.tcp.keypassword 487
transport.tcp.ciphersuites 487
Introscope Workstation.lax 487
lax.application.name 488
lax.class.path 488
lax.command.line.args 488
lax.dir 488
lax.main.class 488
lax.main.method 489
lax.nl.current.vm 489
lax.nl.java.launcher.main.class 489
lax.nl.java.launcher.main.method 489
lax.nl.java.option.additional 490
lax.nl.message.vm.not.loaded 490
lax.nl.valid.vm.list 490
lax.nl.win32.microsoftvm.min.version 490
lax.root.install.dir 491
lax.stderr.redirect 491
lax.stdin.redirect 491
lax.stdout.redirect 492
lax.user.dir 492
lax.version 492
IntroscopeWebView.properties 492
introscope.webview.jetty.configurationFile 493
introscope.webview.tcp.ipaddress 493
introscope.webview.tcp.port 493
introscope.webview.context.path 494
introscope.webview.enterprisemanager.tcp.host 494
introscope.webview.enterprisemanager.tcp.port 494
log4j.logger.WebView 495
log4j.logger.WebServer 495
log4j.logger 496

CA Technologies Prerelease Documentation
Authorized Use Only

log4j.logger.com.wily 496
log4j.logger.com.isomorphic 496
log4j.logger.org.apache.myfaces 497
log4j.logger.org.apache.jasper 497
log4j.logger.org.apache.struts 497
log4j.logger.org.mortbay 497
Logging Configuration Properties 498
introscope.workstation.graph.antialiased 498
introscope.explorer.refresh.interval 499
introscope.webview.svg.text.as.images 499
log4j.appender.console 499
log4j.appender.console.encoding 499
log4j.appender.logfile 499
log4j.appender.logfile.encoding 500
log4j.appender.logfile.File 500
Introscope WebView.lax 500
lax.application.name 500
lax.class.path 500
lax.command.line.args 501
lax.dir 501
lax.main.class 501
lax.main.method 501
lax.nl.current.vm 501
lax.nl.java.launcher.main.class 502
lax.nl.java.launcher.main.method 502
lax.nl.java.option.additional 502
lax.nl.message.vm.not.loaded 503
lax.nl.valid.vm.list 503
lax.nl.win32.microsoftvm.min.version 503
lax.root.installation.dir 503
lax.stderr.redirect 504
lax.stdin.redirect 504
lax.stdout.redirect 504
lax.user.dir 505
lax.version 505

Appendix B: JDBC API **507**

Configuring the JDBC API 507
Using the JDBC API 507
Metric queries 508
Transaction trace queries 509

CA Technologies Prerelease Documentation
Authorized Use Only

Appendix C: Configuring Introscope Extensions 511

Configuring the SNMP Adapter 511
SNMP Adapter limitations 511
SNMP Adapter components 511
Alerts 512
Configuring notifications and alerts 513
Configuring SNMP functionality 514
Uploading to HP LoadRunner 516
Uploading to HP OpenView 518
Converting Introscope metrics to SNMP metrics 519
Configuring the Introscope alerts web services extension 520

Appendix D: The Default tess-default.properties File 523

Modifying the tess-default.properties Property Default File 523
Setting the Perl or Python path for evidence collection 524
perl.path 524
python.path 525
Setting the directory path for the Additional Defect Information link 525
additionalDefectDetailFileLocation 526
Improving defect processing performance 527
defects.affectedUserGroupUpdateInterval_Secs 528
defects.affectedUserGroupAgeoutInterval_Mins 528
Changing the time of the daily APM database purge 529
stats.purgeStatsAt 529
Changing the directory locations for startup, download and bad files 530
collectors.downloadDataDirectory 530
collectors.startupDataDirectory 531
collectors.badDataDirectory 531
Other CA CEM Default Properties 532
baseline.durationInDays 532
BTStatsMetricName.backwardCompatible 533
downloadfiles.minutes.collect 533
email.defaultFromAddress 534
invalid.importexport.filename.chars 535
timTessCommunication.useSsl 535

Index 537

CA Technologies Internal Use Only
Pre-release Documentation
Authorized Use Only

Chapter 1: Introducing CA APM Configuration and Administration Guide

Welcome to the *CA APM Configuration and Administration Guide*.

To use this guide, you need a working knowledge of web servers, browsers, supported platforms, and your network environment. You should be familiar with the applications that will be monitored by CA APM. You should also be aware of other products that will integrate with CA APM.

This guide assumes that:

- CA APM or CA Introscope has been installed according to the information in the *CA APM Installation and Upgrade Guide*.
- You have a basic understanding of CA APM, CA Introscope, and CA CEM as described in the *CA APM Overview Guide*.

This section contains the following topics:

About APM (see page 25)
About this guide and other CA Wily guides (see page 26)

About APM

CA Introscope and CA CEM are the core products in the CA APM solution.

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, as well as 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 SLAs on business services.
- Gain 100 percent transaction visibility.
- Determine the source of problems quickly.
- Conduct triage, identify stakeholders, and perform root-cause analyses.
- Prioritize incidents based on true business impact.

CA Technologies Prerelease Documentation. Authorized Use Only

- Provide proactive and predictive application monitoring.
- Increase reporting and enable continuous improvement.

About this guide and other CA Wily guides

This guide is intended for the CA CEM administrator. For more information, see CA CEM administrator responsibilities (see page 235).

The purpose of this guide is to help administrators configure CA CEM to define, measure, and analyze user transactions in IT production environments.

What you need to know

To use this guide, you need a working knowledge of web servers, browsers, supported platforms, and your network environment. You should be familiar with the applications that will be monitored by CA CEM. You should also be aware of other products that will integrate with CA CEM.

This guide assumes that:

- CA CEM has been installed according to the information in the *CA APM Installation and Upgrade Guide*.
- You have a basic understanding of CA Introscope, CA APM, and CA CEM as described in the *CA APM Overview Guide*.

Each product comes with its own documentation and unique set of resources. This section gives you an overview of other tools available to you.

CA Technologies Prerelease Documentation. Authorized Use Only

Product Documentation

CA APM documentation can be grouped into the following broad categories:

- CA APM
- CA Introscope
- CA CEM
- Extensions and integrations

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.

The following list shows the documentations specific to CA APM.

- *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 CA 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 Installation and Upgrade Guide* — Installation requirements; installing CA APM, including installing Enterprise Manager, APM database, Workstation, WebView, CA CEM, 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 CEM. Introscope properties and TESS default properties references are documented in appendixes.
- *CA APM Security Guide* — Choosing and configuring CA APM, CA Introscope, and CA CEM 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.

The following list shows the documentations specific to CA Introscope. CA APM documentation is also pertinent for CA Introscope.

- *CA APM ChangeDetector User Guide* — Using CA APM ChangeDetector to monitor and report changes in application files and configuration.

CA Technologies Product Release Documentation. Not for Release Only

- *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.
- *CA Cross-Enterprise Application Performance Management 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 APM Java 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.
- *CA APM WebView User Guide* — Using WebView to view Introscope data in dashboards and Investigator.
- *CA Introscope WebView User Guide for SAP* — This guide on how to view CA Introscope data in dashboards, Investigator, and Transaction Tracer, is for WebView for SAP users.
- *CA APM Workstation User Guide* — Using CA Introscope dashboards, Investigator tree and application triage map, Transaction Tracer, and reporting. CA Introscope metrics overview and descriptions.

The following list shows the documentations specific to CA CEM. CA APM documentation is also pertinent for CA CEM.

- *CA APM Transaction Generator Implementation Guide* — Using the CA APM Transaction Generator (CA APM TG) 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 CA APM TG Agent to generate synthetic transactions that you can monitor using CA APM.

The following list shows the documentations specific to extensions and integrations. Some CA APM, CA Introscope, and CA CEM documentations are also pertinent for these extensions and integrations.

- *CA APM for CA SiteMinder Web Access Manager Guide* — Installing, configuring, and using the CA APM for CA SiteMinder Web Access Manager.
- *CA APM for IBM WebSphere Application Server for Distributed Environments Guide* — Installing, configuring and using the CA APM for IBM WebSphere Distributed Environments.

CA Technologies Proprietary Documentation
Authorized Use Only

- *CA APM for Oracle Databases Guide* — Installing, configuring, and using the CA APM for Oracle Databases.
- *CA APM for Microsoft SharePoint Guide* — Installing and configuring the CA APM for Microsoft SharePoint to monitor your SharePoint components during development, QA, staging, and production.
- *CA APM for CA SiteMinder SNMP Collector Guide* — Installing and configuring the CA APM for CA SiteMinder SNMP Collector. Understanding the associated metrics.
- *CA APM for Web Servers Guide* — Installing, configuring, and using the CA APM for Web Servers.
- *CA APM for Oracle WebLogic Portal Guide* — Installing, configuring, and using the CA APM for Oracle WebLogic Portal.
- *CA APM for Oracle WebLogic Server Guide* — Installing, configuring, and using the CA APM for Oracle WebLogic Server.
- *CA APM for IBM WebSphere MQ Guide* — Using Introscope to view metrics from IBM WebSphere MQ.
- *CA APM for IBM WebSphere Portal Guide* — Installing, configuring, and using the CA APM for IBM WebSphere Portal.

CA APM documentation

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

- *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 CA 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 CEM, 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 CEM. 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, CA Introscope, and CA CEM 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 CA Introscope. CA APM documentation is also pertinent for CA 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.
- *CA Cross-Enterprise Application Performance Management 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.
- *CA APM WebView User Guide* — Using WebView to view Introscope data in dashboards and Investigator.
- *<wyisc> WebView User Guide for SAP* — This guide on how to view CA Introscope data in dashboards, Investigator, and Transaction Tracer, is for WebView for SAP users.
- *CA APM Workstation User Guide* — Using CA Introscope dashboards, Investigator tree and application triage map, Transaction Tracer, and reporting. CA Introscope metrics overview and descriptions.

CA Technologies Proprietary & Confidential Information. Authorized Use Only. Do not copy any part of this documentation.

CA CEM Documentation

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

- *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 CEM.
- *CA APM TG Implementation Guide* — Using the CA APM TG (CA APM TG) 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 CA APM TG 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 APM PowerPack for IBM WebSphere Application Server Distributed.
- *PowerPack for Oracle Database Guide* — Installing, configuring, and using the CA APM 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 APM SiteMinder SNMP Collector. Understanding the associated metrics.
- *PowerPack for Web Servers Guide* — Installing, configuring, and using the CA APM PowerPack for Web Servers.
- *PowerPack for WebLogic Portal* — Installing, configuring, and using the CA APM 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 APM PowerPack for WebSphere Portal.

CA Technologies Proprietary Documentation
Authorized Use Only

Chapter 2: Starting and Stopping the Enterprise Manager

This chapter provides instructions for starting and stopping the Enterprise Manager and for automating startup in supported environments.

Note: See Properties Files Reference (see page 403) for details of properties discussed in this guide.

This section contains the following topics:

Starting the Enterprise Manager on Windows (see page 33)
Starting the Enterprise Manager on UNIX (see page 35)
Starting the Enterprise Manager on OS/400 (see page 37)
Starting and configuring the Enterprise Manager on z/OS (see page 39)
Configuring the Enterprise Manager to use a different JVM (see page 47)
About the Enterprise Manager log file (see page 47)
Rolling Enterprise Manager logs by date or size (see page 48)
Configuring the Enterprise Manager config directory (see page 49)
Stopping the Enterprise Manager (see page 50)
Restarting after an unplanned shutdown (see page 52)
Enterprise Manager shell and terminal window commands (see page 53)

Starting the Enterprise Manager on Windows

When running the Enterprise Manager in a Windows environment, you can either start the Enterprise Manager manually or run it as a Windows Service.

Starting the Enterprise Manager Manually on Windows

You can start the Enterprise Manager manually in one of these ways:

- Navigate to Start > Programs > CA APM > Administration > CA APM Enterprise Manager.
- Run the *IntroscopeEnterpriseManager.exe* located in the *<EM_Home>* directory.

The Enterprise Manager starts in accordance with the property settings in *<EM_Home>/Introscope Enterprise Manager.lax*, described in Introscope Enterprise Manager.lax.

CA Technologies Pre-GA Documentation. Authorized Use Only

Running the Enterprise Manager as a Windows Service

Installing the Enterprise Manager as a Windows Service increases its availability. If you configure the Enterprise Manager as a Windows Service, when the machine on which it runs starts, the Enterprise Manager starts automatically. It shuts down when the machine is shut down.

Note: Introscope uses a public domain Java Service Wrapper from Tanuki Software (http://wrapper.tanukisoftware.org) to enable the Enterprise Manager to run as a Windows Service. See the Configuration Property Overview (http://wrapper.tanukisoftware.org/doc/english/properties.html) for information about the Tanuki wrapper.

These instructions assume you are familiar with configuring applications as Windows Services, and are familiar with using the Windows Services console.

Prerequisites for running the Enterprise Manager as Windows Service

To run Enterprise Manager as a Windows Service, verify that Enterprise Manager is already installed and that it runs properly in Console mode.

Configuring the Enterprise Manager as a Windows Service

This section provides instructions on how to register a single Enterprise Manager on a Windows machine that is running with Introscope-default JVM system parameters.

1. Stop the Enterprise Manager—see Stopping the Enterprise Manager (see page 50).
2. In the *IntroscopeEnterpriseManager.lax* file, in your *<EM_Home>/* directory, set the *lax.stdin.redirect* property to blank, and save the file.
3. In the *IntroscopeEnterpriseManager.properties* file, in the *<EM_Home>/config* directory, set the *introscope.enterprisemanager.disableInteractiveMode* property to *true* and save the file.
4. Assign a unique name to the Enterprise Manager service. By default, a service name is assigned in the *EMService.conf*, in the *<EM_Home>/bin* directory. If no other services are configured on the machine, you can use the default service name. Otherwise, modify *EMService.conf* to specify a unique name and a display name for the service using these properties:
   - *wrapper.ntservice.name=<unique Enterprise Manager name>*
   - *wrapper.ntservice.displayname=<Enterprise Manager display name>*

CA Technologies Proprietary/Confidential Information
Authorized Use Only

5. Save changes to the file.
6. Run the *<EM_Home>/bin/EMCtrl32.bat register* script to register the Enterprise Manager instance as a service.

The Windows Services console shows that the Enterprise Manager is configured as a Windows Service.

Repeat these steps for each Enterprise Manager you want to run as a service.

Changing Windows Service properties

If you have to change configuration properties in the *EMService.conf* file, first unregister the service, as described in the following section. Make the changes, then re-register the service for the changes to take effect.

Unregistering an Enterprise Manager Windows Service

Unregister the Enterprise Manager Service:

- before you make changes to the *EMService.conf* properties file.
- to return the Enterprise Manager instance to Console control.
- before you uninstall the Enterprise Manager instance.

To unregister an Enterprise Manager Service:

1. Stop the Enterprise Manager Service and shut down the Windows Services GUI.
2. Run *<EM_Home>/bin/EMCtrl32.bat unregister.*

Note: If either the Enterprise Manager Service or the Windows Services GUI is active, the service is not deleted, but is *marked* for deletion. This is standard Windows Service behavior, and occurs because the Windows Registry lock prevents the service from being removed. When you restart the machine, the service is deleted.

Stopping and Restarting the Enterprise Manager Windows Service

Use the standard controls in the Windows Services GUI to start, stop, and restart the Enterprise Manager running as a Windows Service.

Starting the Enterprise Manager on UNIX

When running the Enterprise Manager in a UNIX environment, you can either start the Enterprise Manager manually, or run it in *nohup* mode.

CA Technologies Prerelease Documentation. Authorized Use Only.

Starting the Enterprise Manager Manually on UNIX

To manually start the Enterprise Manager on UNIX:

- To start the Enterprise Manager on UNIX, run Introscope_Enterprise_Manager by typing this at the command prompt:

'./Introscope Enterprise Manager'

The Enterprise Manager starts in accordance with the property settings in *<EM_Home>*, described in Properties Files Reference.

Enterprise Manager output appears in the UNIX shell. The shell does not accept user input, unless you start up the Enterprise Manager in *nohup* mode, as described in the following section.

Running Enterprise Manager control scripts on UNIX

You can manage the Enterprise Manager by running control scripts that allow you to start, stop, and check the server application status. You can use the *EMCtrl.sh startup* script to start and stop the Enterprise Manager instead of directly using the *'./Introscope Enterprise Manager'* command.

To run the Enterprise Manager control scripts on UNIX:

Use the following commands at the *bin* directory:

- ./EMCtrl.sh start—This starts the Enterprise Manager.
- ./EMCtrl.sh stop—This stops the Enterprise Manager.
- ./EMCtrl.sh status—This gives the status of the Enterprise Manager.
- ./EMCtrl.sh help—This shows the supported options for the script.

CA Technologies Prerelease Documentation. Authorized Use Only.

Running Enterprise Manager in nohup Mode on UNIX

The UNIX *nohup* command runs a command, but suppresses the action of the *HUP* (hangup) signal, enabling the command to continue running after the user who issued it logs out. In this case, output that normally goes to the terminal goes to a file named *nohup.out*.

To run the Enterprise Manager in nohup mode on UNIX:

1. In the *IntroscopeEnterpriseManager.lax* file, in the *<EM_Home>/* directory, set the *lax.stdin.redirect* property to blank, and save the file.
2. In the *IntroscopeEnterpriseManager.properties* file, in the *<EM_Home>/config* directory, set the property introscope.enterprisemanager.disableInteractiveMode to true and save the file.
3. Start the Enterprise Manager with this command:

   nohup Introscope_Enterprise_Manager&

Note: Do not run the Enterprise Manager in *nohup* mode without performing the configuration described above. Otherwise, the Enterprise Manager might not start, or might start and consume excessive system resources.

Starting the Enterprise Manager on OS/400

This section provides instructions for starting the Enterprise Manager on OS/400 using the OS/400 start script, and invoking from the command line.

Running the OS/400 start script

When the Enterprise Manager is installed on OS/400, the Enterprise Manager installer automatically generates two start scripts. One for Enterprise Manager and one for WebView. To invoke the start script from a QShell command prompt from the *<EM_Home>/* directory, use one of the following commands:

- Enterprise Manager start script:

  nohup ./runem.sh &

- WebView start script:

  nohup ./runwv.sh &

CA Technologies Pre-release Documentation. Authorized Use Only

If you run the script from a directory other than *<EM_Home>*, identify the home directory as shown below:

```
nohup ./runem.sh /<EM_Home> &
```

where *<EM_Home>* is the root of your Enterprise Manager installation directory.

The output will default to the *nohup.out* in the directory from which the script is executed.

Note: Use the *runem.sh* script to start the Enterprise Manager on Red Hat Linux 5.

Using control scripts for OS/400 and Z/OS

You can use control scripts in OS/400 and Z/OS to start the Enterprise Manager.

Use these control scripts for OS/400 and Z/OS:

1. Go to the *<EM_Home>\bin* directory.
2. Use these commands:
   - *./EMCtrl.sh start*—Starts the Enterprise Manager.
   - *./EMCtrl.sh status*—Verifies whether the Enterprise Manager is running or not.
   - *./EMCtrl.sh* —Stops the Enterprise Manager.

Running the Enterprise Manager as a Java program

To start Enterprise Manager from the command line, enter this command from the *<EM_Home>* directory:

```
lib/EnterpriseManager.jar:lib/WebServices.jar:lib/IntroscopeClient.jar:lib/SNMPAd
apter.jar:lib/Workstation.jar:lib/IntroscopeServices.jar
```

Start command syntax can vary, depending on your version of Java and other environmental factors.

Note: If you are running the Enterprise Manager as a Java Program on SAP, use this command:
lib/EnterpriseManager.jar:lib/WebServices.jar:lib/IntroscopeClient.jar:lib/SNMPAdapter.jar:lib/Workstation.jar:lib/IntroscopeServices.jar:lib/WebViewAPI.jar:lib/WebViewServer.jar

CA Technologies Prerelease Documentation. Authorized Use Only.

Starting and configuring the Enterprise Manager on z/OS

To start the Enterprise Manager on z/OS, use the *runem.sh* file.

For instructions on running the Enterprise Manager as a Java program, see Running the Enterprise Manager as a Java program (see page 38).

To facilitate use of Introscope in the z/OS environment, run the product in a manner familiar to z/OS operations and development staff—as a z/OS batch job, or as a started task. You can also run Enterprise Manager as a shell process. This section describes running the Enterprise Manager as a batch job or started task:

- Prerequisites for running Enterprise Manager as a z/OS batch job (see page 39) briefly describes the z/OS environment and the BPXBATCH utility. This section describes software prerequisites, security, authorization and address space size requirements.
- Configuring Enterprise Manager to run as a z/OS batch job (see page 42) describes the tasks necessary to customize the sample JCL and other files in preparation for running the Enterprise Manager as a batch job on the z/OS platform.
- Running the Enterprise Manager on z/OS (see page 44) describes the activities required to run the Enterprise Manager as a batch job or started task on the z/OS platform.

Prerequisites for running Enterprise Manager as a z/OS batch job

There are a variety of prerequisites for running an Enterprise Manager as a z/OS batch job. The following sections provide more detail about these requirements.

z/OS base infrastructure for Java

The MVS utility program, BPXBATCH, enables users to invoke UNIX System Services to run shell commands, shell scripts and executable files in MVS batch. You can therefore use the BPXBATCH utility to start the Enterprise Manager.

Software prerequisites

The specific software prerequisites to run Java (and therefore the Enterprise Manager) on z/OS can be found in the Java for z/OS Program Directory GI10-0614. The minimum software requirements are:

- The Enterprise Manager requires JVM 1.5 (either 32 or 64 bits).

  Note: WebView needs JVM 1.5 to run.
- UNIX System Services must be enabled

CA Technologies - Prerelease - Unauthorized Use Documentation

- Language Environment 1.5 or higher
- Java for z/OS

Security and authorization

Authorization is determined by the RACF profile associated with a user's USERID. To successfully start the Enterprise Manager, a user must have authority to access the following resources:

- UNIX System Services as specified by the OMVS parameter in the RACF profile
- EXEC PGM= BPXBATCH
- EXEC WILYPROC

It is important to understand how the user is identified in these contexts:

- Batch—When a user submits a batch job through the TSO/E or ISPF interface, the user's TSO UserID and password are propagated to the batch job. Provided the submitted Batch JOB CARD does not explicitly contain a User and Password, the batch job acquires the same access authority as the TSO user. If User and Password are coded on the Batch JOB CARD, they override the UserID and Password of the TSO/E user.
- Started Task—If the Enterprise Manager is initiated as a started task, it runs under the authorization of the STC user ID. A Site-specific algorithm might be used to relate the name of the started PROC to a particular UserID. You must make sure that the algorithm produces a UserID with the correct authorization to run the Enterprise Manager.

Virtual memory

If the OMVS address space is not large enough, an out-of-memory condition occurs when an attempt is made to run the Enterprise Manager. The OMVS ASSIZEMAX(address-space-size) parameter in the RACF profile for the owner of the Enterprise Manager job or started task must be large enough to execute the Enterprise Manager. The maximum value for this parameter is 2G.

In addition, in order to run the Enterprise Manager with Java 1.5 and avoid encountering Java 1.5 GC Heap Initialization errors, it may be necessary to raise the *MEMLIMIT* in your *SMFPRMxy PARMLIB* member. Update your *SMFPRMxy* to add the following line if it does not already exist:

```
MEMLIMIT(256M)          /* SET MINIMUM MEMLIMIT FOR JAVA 1.5 */
```

Note: The 256M value is minimum requirement. Recommended value can be 1024M or more.

CA Technologies Proprietary Information. Authorized Use Only. Documentation.

DB2

1. To access JDBC and SQLJ from UNIX Systems Services, you must set the following environment variables:
   - *STEPLIB* must include the *SDSNEXIT*, *SDSNLOAD* and *SDSNLOD2* data sets
   - *LIBPATH* and *LD_LIBRARY_PATH* must include the DLL libraries for the JDBC and SQLJ drivers
2. In order to run the Enterprise Manager as a z/OS job, the *runem.sh* shell script must contain the following statements where *DSNXXX* denotes the High Level Qualifier of your DB2 libraries:
   - *export STEPLIB=DSNXXX.SDSNEXIT:DSNXXX.SDSNLOAD:DSNXXX.SDSNLOD2*
   - *export LIBPATH=/usr/lpp/db2/db2710/lib:$LIBPATH*
   - *export LD_LIBRARY_PATH=/usr/lpp/db2/db2710/lib:$LD_LIBRARY_PATH*

   Note: While *DB2SQLJPROPERTIES*—the directory containing the *db2sqljjdbc.properties* file—is used to specify the name of a *SQLJ* properties file, it is not always necessary to set this environment variable.
3. The CLASSPATH must include the JDBC and SQLJ driver classes. Include this file in the *classpath* in the runem.sh shell script:

   /usr/lpp/dsn710/db2/db2710/classes/db2sqljclasses.zip

   Note: The character string "*710*" in *dsn710* and *db2710* refers to the version and release of DB2 being used. If another version of DB2 is being used, these directories must be changed.
4. Specify the DB2 subsystem identifier:

   SYS_DB2_SUB_SYSTEM_NAME=xxx, where xxx is the subsystem identifier.

CA Technologies Proprietary Documentation
Authorized Use Only

JDBC

JDBC for z/OS is based on the DB2 Call Level Interface (CLI). You must ensure that the DB2 CLI has been set up as follows:

1. The DB2 CLI application plan (the default name is DSNACLI) must be bound before you can use the DB2 CLI. A sample bind job can be found in DSNXXX.SDSNSAMP(DSNTIJCL) where DSNXXX denotes the high level qualifier of your DB2 installation.
2. The DB2 CLI application plan must be public. Execute the following command from SPUFI or from a batch job:

   GRANT EXECUTE ON PLAN DSNACLI TO PUBLIC
3. Set DSNAOINI to point to the CLI initialization file. The CLI initialization file provides information about DB2 subsystem names and additional configuration parameters. You can place the file either in an HFS file or in an MVS data set using one of the following statements:
   - *export DSNAOINI=/usr/lpp/db2/db2510/dsnaoini*
   - *export DSNAOINI=DSNxxx.CLIINI*
   - *export DSNAOINI=DSNxxx.CLIINI(CONF1)*
4. Make sure that the DB2 subsystem referenced in the CLI initialization file references the DB2 subsystem that you want to connect to. DB2 subsystem names usually take the form DB2n where n is a unique qualifier.

Configuring Enterprise Manager to run as a z/OS batch job

Before the Enterprise Manager can run as a z/OS batch job, the following files (extracted from the z/OS Enterprise Manager installer file) must be copied into the z/OS environment and customized to ensure compatibility with site-specific requirements.

- *IntroscopeEnterpriseManager.properties* – the CA Introscope Enterprise Manager properties file, which specifies information pertinent to the operation of the Enterprise Manager
- *runem.sh* – the shell script used to launch the Enterprise Manager process
- *WILYPROC* – the z/OS PROC used to execute the z/OS BPXBATCH program that launches the shell script
- *readme.txt* - a description of the runem.sh and WILYPROC files and how to use and customize them

CA Technologies Prerelease Documentation. Authorized Use Only.

Copying and customizing Enterprise Manager files

IntroscopeEnterpriseManager.properties and *runem.sh* are UNIX files that must be copied into the Hierarchical File System (HFS) on the z/OS platform. These files are used by UNIX System Services and the Enterprise Manager respectively. Use a UNIX-based editor such as vi to make required changes to these files. Do not use the ISPF editor because it does not properly handle insertion of characters that increase line length.

WILYPROC is a z/OS file that is used by the JES2 component of z/OS and by DB2 for z/OS. Make changes to this file using the ISPF editor.

To copy the Enterprise Manager files:

- Copy the *IntroscopeEnterpriseManager.properties* and *runem.sh* files into the Introscope installation directory, as shown in the following examples, where *<UID>* is the UID assigned to CA Technologies:

  /u/<UID>/introscope9.0/runem.sh

  /u/<UID>/introscope9.0/config/IntroscopeEnterpriseManager.properties

To customize Enterprise Manager properties file:

1. Open the *u/xxxx/introscope9.0/config/IntroscopeEnterpriseManager.properties* file.
2. Find the Interactive Mode heading and make sure the value for the following property is set to true:

   introscope.enterprisemanager.disableInteractiveMode=true

To customize the UNIX shell script runem.sh:

1. Verify the version and release of DB2. The DB2 files in the *runem.sh* shell script reference DB2 V7.1.

   Note: If you are connecting to a DB2 subsystem other than version 7.1, you must change all instances of *710* to the appropriate character string for the version being used.
2. Make the following changes to the *runem.sh* script:
   - In the CLASSPATH, */usr/lpp/dsn710/db2/db2710/classes/db2sqljclasses.zip*\
   - *export STEPLIB=DSN710.SDSNEXIT:DSN710.SDSNLOAD,DSNXXX.SDSNLOD2*
   - *#export LIBPATH=/usr/lpp/dsn710/db2/db2710/lib:$LIBPATH*
   - *#export LD_LIBRARY_PATH=/usr/lpp/dsn710/db2/db2710/lib:$LD_LIBRARY_PATH*

     Note: If *LIBPATH* and *LD_LIBRARY_PATH* have been commented out, remove the #.

CA Technologies Proprietary Documentation
Limited Use Only

To copying Enterprise Manager files to a z/OS partitioned dataset (PDS)

1. Using ISPF, create a partition dataset named *userid.JCL.CNTL*, (where userid is the userid assigned to CA Technologies) with fixed-block (FBA) of 80 byte records.
2. Copy WILYPROC into the PDS *userid.JCL.CNTL*.
3. Do not try to copy the shell script *runem.sh*, or the properties file *IntroscopeEnterpriseManager.properties*, into this PDS. These files have lines that are longer than 80 bytes. An attempt to copy them either fails, or lines are truncated.

To customize the z/OS PROC that invokes BPXBATCH

1. Edit WILYPROC using the ISPF editor. Make the following changes:
   - On the *ISPF* command line, type *CAPS OFF* and press Enter.
   - On the *ISPF* command line, type *NUM OFF* and press Enter.
   - Set the *IHOME* parameter equal to the fully qualified path to the directory that contains *runem.sh*.
2. This statement is coded to end the batch job and the UNIX shell script within 20 seconds:

   // PARM='SH nohup &IHOME/runem.sh &IHOME & sleep 20'

   If z/OS is too busy to initiate the Enterprise Manager within 20 seconds, the Enterprise Manager does not start. If you encounter this problem, try increasing the sleep time.
3. Modify the STDERR and STDOUT file names contained in the following statements to satisfy installation naming conventions:

   //STDOUT DD PATH='/<path>/stdout',

   //STDERR DD PATH='/<path>/stderr',

To copy WILYPROC into an active PROCLIB

- Using ISPF option 3, copy the PROC to *xxxx*.PROCLIB where *xxxx* represents an site-dependent high-level qualifier. If WILYPROC is not an acceptable name in your installation, rename the PROC accordingly.

CA Technologies Prerelease Documentation. Authorized Use Only.

Running the Enterprise Manager on z/OS

You can run the Enterprise Manager on z/OS in these ways:

- Running as a started task from the z/OS console or SDSF
- Running as a batch job

Running as a started task from the z/OS console or SDSF

You can run the Enterprise Manager as an MVS started task (STC). A started task is a familiar, administrator-friendly environment that allows the Enterprise Manager to be started, monitored, and cancelled from the MVS console. In this case, the Enterprise Manager runs under the authorization of the STC user ID, allowing assignment of specific authorities to the started task.

To run the Enterprise Manager as a started task:

- Type this command on the MVS console or on the command line of SDSF:

  /start WILYPROC

  You must type the member name of the PROC stored in an active PROCLIB. If you were required to change the name of WILYPROC, you must use the new name in the */start* command.

Running as a batch job

You can run the Enterprise Manager as a batch job, using JCL (job-control language) to describe the job's requirements.

To run the Enterprise Manager as a batch job:

- Submit the following JCL:

```
//jobname   JOB   installation jobcard parameters,NOTIFY=&SYSUID
ISCOPE     EXEC  WILYPROC
/*
```

  You must execute the member name of the PROC stored in an active PROCLIB. If you were required to change the name WILYPROC, you must use the new name in the EXEC statement.

  The z/OS batch job and the OMVS shell end shortly after the Enterprise Manager processes start. If you submitted the batch JCL from your ISPF session and your JOB card contains NOTIFY=&SYSUID, your ISPF session receives notification when your batch job ends.

CA Technologies Pre-release Documentation. Authorized Use Only.

Verifying the Enterprise Manager is running

The most accurate way to determine that the Enterprise Manager has started correctly is to view the STDOUT file. To view STDOUT, telnet into z/OS and view the STDOUT file defined in the WILYPROC. The last line of the file indicates that the Enterprise Manager has started.

You can also determine that the Enterprise Manager has started from SDSF using one of the following procedures:

- If you started the Enterprise Manager as a started task, on the DA screen you should see at least two active started tasks whose owner is the owner associated with the Enterprise Manager PROC.
- If you started the Enterprise Manager by submitting a batch job from your TSO session, look for two active tasks with your TSO USERID as the owner.

You might see the job that invoked the WILYPROC, as well as the OMVS shell script. These two entries disappear from the list after the Enterprise Manager starts.

The owner of your TSO session is also your TSO/E USERID. Do not mistake your TSO/E session for one of the Enterprise Manager tasks.

Starting the Enterprise Manager automatically

To start the Enterprise Manager automatically when JES2 starts:

1. Place the following statement in JES2PARM:

   $VS,'start WILYPROC

2. Type the member name of the PROC stored in an active PROCLIB. If you were required to change the name of WILYPROC, use the new name in the */start* command.

   You can also run the Enterprise Manager as a batch job using the job scheduling package you use to manage your batch workload.

CA Technologies Prerelease Documentation. Authorized Use Only.

Configuring the Enterprise Manager to use a different JVM

By default, an Introscope component uses the JVM it is bundled with. This section has instructions on how to configure the Enterprise Manager to use a different JVM.

- If you start the Enterprise Manager with *Introscope Enterprise Manager.exe*, edit the *lax.nl.current.vm* property in *Introscope Enterprise Manager.lax* to point to the desired JVM executable.
  - You can use an absolute or relative path to the executable.
  - On Windows, escape backslashes, and include the JVM executable file extension: *java.exe* instead of *java*.
- If you run the Enterprise Manager as a Windows service, edit the *wrapper.java.command* property in *EMService.conf* to point to the desired JVM executable.
  - You can use absolute paths; backslashes must be escaped.
  - You can also use relative paths; CA Technologies recommends that you use UNIX forward slashes and start with dot-slash (*./*).
  - The executable is *java* not *java.exe*.
- If you start the Enterprise Manager with the *runem.sh* script on z/OS or AS/400, edit the *JAVA_HOME* variable in *runem.sh* to point to the desired JVM root directory.
  - Use an absolute to the executable, not a relative path.

About the Enterprise Manager log file

The Enterprise Manager log file, *IntroscopeEnterpriseManager.log*, which is located in the *<EM_Home>/logs* directory, records information about how Introscope, CA CEM, and the APM database are running and error conditions: for example, messages about start up and shut down, certain types of exceeded overhead, third-party libraries, CA CEM communications with the Enterprise Manager, and data about specific Introscope and CA CEM properties.

You can configure the Enterprise Manager logging properties. For more information, see Logging (see page 411).

CA Technologies Pre-release Documentation. Authorized Use Only

Rolling Enterprise Manager logs by date or size

You can roll Enterprise Manager logs based on date or size, retaining a specified number of days of information and purging the rest.

To roll Enterprise Manager log files:

1. Open the *IntroscopeEnterpriseManager.properties* file in the *<EM_Home>\config* directory.
2. Locate the Logging Configuration section.
3. Modify properties to accomplish desired log rolling. See detailed instructions below.
4. Save and close the *IntroscopeEnterpriseManager.properties* file.

   Depending on the desired log rolling behavior, set the following properties.

To roll log files by date, or on a daily basis:

- Modify the properties:
  - *log4j.appender.logfile=com.wily.org.apache.log4j.DailyRollingFileAppender*
  - *log4j.appender.logfile.DatePattern*

  These properties work together to roll Enterprise Manager log files based on date. *log4j.appender.logfile.DatePattern* defines the log rolling time period that the *log4j.appender.logfile=com.wily.org.apache.log4j.DailyRollingFileAppender* property carries out.

  For example, this configuration rolls the log files every day:
  log4j.appender.logfile=com.wily.org.apache.log4j.DailyRollingFileAppender
  log4j.appender.logfile.DatePattern='.'yyyy-MM-dd

To roll log files by time interval:

- Modify the property *log4j.appender.logfile.DatePattern.* This Enterprise Manager property determines when to roll Enterprise Manager log files based on time of day.

  Here are two examples.

  - This configuration rolls the log files every hour:
    `log4j.appender.logfile.DatePattern='.'yyyy-MM-dd-hh`
  - This configuration rolls the log file rolls every day at midnight:
    `log4j.appender.logfile=com.wily.org.apache.log4j.DailyRollingFileAppender`
    `log4j.appender.logfile.DatePattern='.'yyyy-MM-dd`

CA Technologies Proprietary Documentation. Authorized Use Only.

To roll log files by log size and quantity:

- Modify the properties:
  - *log4j.appender.logfile.MaxBackupIndex*

    This property determines how many backup log files the Enterprise Manager keeps.
  - *log4j.appender.logfile.MaxFileSize.*

    This property determines the maximum size of the log file.

For example, this configuration keeps up to three backup/rolled Enterprise Manager logs, and each log is up to 2 kilobytes in size:

```
log4j.appender.logfile.MaxBackupIndex=3
log4j.appender.logfile.MaxFileSize=2KB
```

Configuring the Enterprise Manager config directory

If you have multiple Enterprise Managers and want to run them using a single instance of Introscope, you can create a separate configuration directory for each Enterprise Manager. For each Enterprise Manager, make a copy of the config, logs, data, and traces directories. In each copy of the config directory, set the appropriate properties in the IntroscopeEnterpriseManager.properties file to *point to the copies of your config, lo*gs, data, and traces directories.

CA Technologies Proprietary Use Documentation. Authorized Use Only.

By creating separate configuration directories, you can keep all the directories within one instance of Introscope while maintaining separate config, logs, data, and traces directories for each Enterprise Manager.

To specify a different configuration directory for your Enterprise Manager:

1. Open the *IntroscopeEnterpriseManager.properties* file and modify the following property:
   `introscope.enterprisemanager.directory.config=config`

   The default directory is *config*, which is the configuration directory created when Introscope was first installed. Change the location of this directory using either an absolute path (for example, *C:\\Introscope\em1\\config*) or a path relative to the Introscope installation directory (e.g. config2). For example:
   introscope.enterprisemanager.directory.config=
   C:\\Introscope\\em1\\config

   Note: In a Windows environment, backslashes must be escaped (doubled) for the path to function.

2. Save the changed *IntroscopeEnterpriseManager.properties* file.

3. Open the *Introscope Enterprise Manager.lax* file and modify the following property:
   `lax.nl.java.option.additional`

   This property modifies the Java command-line and can contain multiple properties. Use the location of the *config* file you selected in step one. For example:
   `lax.nl.java.option.additional=-Xms512m -Xmx1024m -Djava.awt.headless=false -Dcom.wily.introscope.em.properties=C:\Introscope\em1\config`

   Note: This property can point either to a directory containing a Enterprise Manager properties file, or to an *IntroscopeEnterpriseManager.properties* file directly. For example:
   `-Dcom.wily.introscope.em.properties=C:\Introscope\em1\config`

   or

   `-Dcom.wily.introscope.em.properties=C:\Introscope\em1\config\IntroscopeEnterpriseManager.properties`

4. Save the changes to the *Introscope Enterprise Manager.lax* file.

5. Restart the Enterprise Manager.

Stopping the Enterprise Manager

When you stop the Enterprise Manager:

- agents stop reporting data.
- SmartStor does not save data.
- you are logged out of the Workstation.
- all users connected to the Enterprise Manager are disconnected.

CA Technologies Pre-release Documentation. Author Edition Only

Stopping the Enterprise Manager from the terminal window

From the Enterprise Manager terminal window, enter:
`iscopeshutdown`

Note: Interactive mode must be enabled to use this command. Set *introscope.enterprisemanager.disableInteractiveMode* to *false* in the Enterprise Manager properties file. For more information, see Properties Files Reference (see page 403).

Stopping the Enterprise Manager from the Workstation

If you have shutdown privileges for the Enterprise Manager, you can stop the Enterprise Manager from the Workstation in a Management Module Editor window by choosing Manager > Shut Down Enterprise Manager.

Automatic shutdown of Enterprise Manager

If the Enterprise Manager detects errors writing to disk, indicating a shortage of disk space, it automatically shuts down. If failover is configured, agents can failover to another Enterprise Manager. For more information, see the *CA APM Java Agent Implementation Guide* or the *CA APM .NET Agent Implementation Guide*.

CA Technologies Pre-release Documentation. Authorized Use Only.

Restarting after an unplanned shutdown

Important! This topic only applies to CA APM. If you are using standalone CA Introscope (without TIMs), then skip this topic.

In most circumstances, restarting the Enterprise Manager is the same as starting it. However, if the Enterprise Manager has stopped in an error state instead of a planned shutdown, you may need to delete some files before restarting.

If an Enterprise Manager shut itself down (crashes) because of very heavy load or a file processing error, then there might be unprocessed files left in the *<EM_Home>/cem/tess/downloadfiles* and *<EM_Home>/cem/tess/startupfiles* directories. If the Enterprise Manager is restarted with these files present, the files are processed on restart. If the shutdown occurred because of an error in these files or the load caused by these files, then another crash might occur immediately on restart when the Enterprise Manager begins again to process these files.

The solution is to delete the files from the two directories before restarting the Enterprise Manager.

To restart the Enterprise Manager after a crash:

Important! Deleting files as described in this procedure does result in some data loss. You can lose recent defect details, customer experience metrics, and statistics.

1. On the Enterprise Manager that is running the TIM Collection Service, delete the files in the *<EM_Home>/cem/tess/downloadfiles* and *<EM_Home>/cem/tess/startupfiles* directories.

   The files are XML files and have names such as *btstats-2010-01-05-23_26_13-172.16.252.20.xml* and *defect-1262733968000057712-172.16.252.20.xml*

   For information about Enterprise Manager services, including the TIM Collection Service, see Distributing Enterprise Manager services (see page 252).

2. Start the Enterprise Manager as described in one of these topics:

   - Starting the Enterprise Manager on Windows (see page 33)
   - Starting the Enterprise Manager on UNIX (see page 35)
   - Starting the Enterprise Manager on OS/400 (see page 37)
   - Starting and configuring the Enterprise Manager on z/OS (see page 39)

Troubleshooting multiple crashes on an Enterprise Manager:

Symptom:

The Enterprise Manager that is running the TIM Collection Service crashed and crashes again on restart.

CA Technologies Pre-release Documentation. Authorized Use Only.

A possible cause is that TIM-related data files that are putting a heavy load on the Enterprise Manager or that contain errors.

Solution:

Delete these files as described in Restarting after an unplanned shutdown (see page 52).

Enterprise Manager shell and terminal window commands

When the Enterprise Manager is running, if *lax.stdin.redirect* (defined in Properties Files Reference) is set to *console*, you can enter the following commands in the Enterprise Manager terminal window or shell:

Command	Result
iscopedebugpolling [true | false] [delay in seconds]	Periodically reports debugging information to the log
iscopehelp	Display help screen
iscopeshutdown	Shuts down the Enterprise Manager

Note: Interactive mode must be enabled to use these commands. Set *introscope.enterprisemanager.disableInteractiveMode* to *false* in the Enterprise Manager properties file. For more information, see Properties Files Reference (see page 403).

CA Technologies Pre-release Documentation
Authorized Use Only.

Chapter 3: Configuring Enterprise Manager Communications

This section includes information about how the Enterprise Manager communicates with other components in an Introscope environment, default communications behaviors, and options for configuring Enterprise Manager communications.

This section contains the following topics:

The default Enterprise Manager communications channel (see page 55)
Connecting to the Enterprise Manager across a firewall (see page 58)
Enabling SSL for Enterprise Managers on OS/400 (see page 59)
Configuring the Enterprise Manager web server for HTTPS (see page 59)

The default Enterprise Manager communications channel

By default the Enterprise Manager communicates with agents, Workstations, WebView, and in a cluster of Enterprise Managers on port 5001. See Working with Enterprise Manager Clustering (see page 101), for more information.

All socket communication to the Enterprise Manager is obfuscated, and passwords are encrypted.

The properties in the Enterprise Manager properties file that define communication channels and their use are:

- *introscope.enterprisemanager.serversocketfactory.<ChannelName>*
- *introscope.enterprisemanager.port.<ChannelName>*
- *introscope.enterprisemanager.serversockets.reuseaddr*
- *introscope.enterprisemanager.enabled.channels*

Each agent that communicates with the Enterprise Manager has analogous properties that must specify matching values. See the *CA APM Java Agent Implementation Guide* or the *CA APM .NET Agent Implementation Guide*, as appropriate for your environment, for information about agent properties.

Note: The .NET agent does not support SSL connections to the Enterprise Manager.

Similarly, the WebView properties file specifies the Enterprise Manager port in the *introscope.webview.enterprisemanager.tcp.port* property.

CA Technologies Proprietary Documentation
Authorized Use Only

Rebinding a communication channel to a port in use

There are socket states where nothing is actually using the socket, but memory is still in use, thus rendering the socket unusable. Examples of these states are FIN_WAIT and TIME_WAIT.

The property *introscope.enterprisemanager.serversockets.reuseaddr* (see page 407) allows an Enterprise Manager communication channel to rebind to a local port that is in use. Use this feature if an Enterprise Manager socket is stuck in an unusable state, but you still want to start the Enterprise Manager.

Note: Do not enable this property by default and use it only when a socket is in an unusable state. You should disable it after the Enterprise Manager starts, as there may be a valid reason that the port is in use (for example, a Collector on the same machine is already bound to the port).

Configuring SSL

Configuring the Enterprise Manager for SSL is done in *IntroscopeEnterpriseManager.properties* by enabling a secure communications channel. This can be in addition to the default non-secure channel.

The following is an example of enabling both a secure and non-secure communications channel. It sets up the default non-secure listener on port 5001, and a secure listener on port *5443*:

```
# The enabled communication channels.
introscope.enterprisemanager.enabled.channels=channel1, channel2

# The default communication channel.
introscope.enterprisemanager.serversocketfactory.channel1=com.wily.isengard.postofficehub.link.net.server.DefaultServerSocketFactory
introscope.enterprisemanager.port.channel1=5001

# The secure communication channel.
introscope.enterprisemanager.serversocketfactory.channel2=com.wily.isengard.postofficehub.link.net.server.SSLServerSocketFactory
introscope.enterprisemanager.port.channel2=5443
# Location of a keystore containing certificates for authenticating the Enterprise Manager to clients.
# Either an absolute path or a path relative to the config directory.
# On Windows, backslashes must be escaped.  For example:
#
introscope.enterprisemanager.keystore.channel1=C:\\Introscope\\config\\internal\\server\\keystore
introscope.enterprisemanager.keystore.channel2=internal/server/keystore
# The password for the keystore.
```

CA Technologies Proprietary - Documentation Authorized Use Only.

```
introscope.enterprisemanager.keypassword.channel2=password
# Location of a truststore containing trusted client certificates.
# Either an absolute path or a path relative to the config directory.
# On Windows, backslashes must be escaped.  For example:
#
introscope.enterprisemanager.truststore.channel1=C:\\Introscope\\config\\inte
rnal\\server\\keystore
# A truststore is only needed if client authentication is required.
# If no truststore is specified, the Enterprise Manager trusts all client
certificates.
```

SSL property rules

The Enterprise Manager must be configured with an SSL server socket in *IntroscopeEnterpriseManager.properties* by enabling a secure communications channel in addition to the default channel.

When configuring SSL properties, the following rules:

- The keystore property is the location of a keystore. It should contain a certificate used by the Enterprise Manager to authenticate itself to the client. The Enterprise Manager's default keystore contains a self-signed certificate that is inherently untrustworthy. This doesn't matter because by default the Workstation and agent don't validate server certificates.
- The keystore property is either an absolute path or a path relative to the *config* directory.
- If the *needclientauth* property is set to *true*, the Enterprise Manager requires the client to authenticate.
  - The Enterprise Manager must be configured with a trustore.
  - Agents and Workstations must be configured with a keystore containing a certificate trusted by the Enterprise Manager.
  - If *needclientauth = true*, and a truststore is not configured, the Enterprise Manager does not start.
  - This requires setting the trustore property.
- The *truststore* property is the location of a truststore. It should contain the certificates of trusted clients. It is only needed if clients are required to authenticate.
  - If no truststore is specified, all client certificates are trusted.

CA Technologies Proprietary & Documentation
Authorized Use Only.

- The *ciphersuites* property is used to set the enabled cipher suites.
  - This can be used to enable anonymous authentication or avoid data encryption. For example, for anonymous authentication, you might set the cipher suites to *SSL_DH_anon_WITH_RC4_128_MD5*. For example:
    `introscope.enterprisemanager.ciphersuites.channel2=SSL_DH_anon_WITH_RC4_128_MD5`
  - A particular JSSE implementation may have several cipher suites available that permit anonymous authentication (they all contain "*anon*"). However, the cipher suite will be common to both client and server, and enabled on both.
  - You might also specify a cipher suite that does not encrypt data. For example:
    `introscope.enterprisemanager.ciphersuites.channel2=SSL_RSA_WITH_NULL_SHA`

Connecting to the Enterprise Manager across a firewall

HTTP tunneling is enabled by default. Tunneling technology enables one network to send its data via another network's connections.

To use the tunneling transport on the Workstation, connect to the port for the Enterprise Manager's embedded web server. By default the port is 8081. You can configure it with the *introscope.enterprisemanager.webserver.port* property in the *Enterprise Manager properti*es file. For more information, see introscope.enterprisemanager.webserver.port (see page 408).

In addition to using tunneling for the Enterprise Manager, you can configure agents to send information using tunneling technology. This feature enables agents to connect to an Enterprise Manager remotely. See the *CA APM Java Agent Implementation Guide* for more information about agent HTTP tunneling.

CA Technologies Pre-Release Documentation
Authorized Use Only

Enabling SSL for Enterprise Managers on OS/400

By default the AS/400 JVM uses a native JSSE provider. You must configure use of the pure Java JSSE provider to allow an Enterprise Manager hosted on an AS/400 to communicate with other Introscope components such as the Workstation, or in a cluster, other Enterprise Managers.

To enable SSL for Enterprise Managers on OS/400:

1. Open your Java security file, usually in your JVM's *lib/security* directory.
2. In the list of security providers, look for *com.ibm.jsse.IBMJSSEProvider*. If it exists, make sure it is higher up the list and thus takes priority over other SSL providers, including *com.ibm.as400.ibmonly.net.ssl.Provider*. If it not present, add it to the list.
3. Set the following SSL properties, replacing any previous definitions:

```
ssl.KeyManagerFactory.algorithm=IbmX509
ssl.TrustManagerFactory.algorithm=IbmX509
ssl.SocketFactory.provider=com.ibm.jsse.JSSESocketFactory
ssl.ServerSocketFactory.provider=com.ibm.jsse.
JSSEServerSocketFactory
```

Configuring the Enterprise Manager web server for HTTPS

You can configure the Enterprise Manger web server for HTTPS by using the Jetty configuration file. The Jetty file is provided for advanced configuration and can be found in the config directory of the *<EM_Home>* installation.

This file is activated by uncommenting the *introscope.enterprisemanager.webserver.jetty.configurationFile* property in *IntroscopeEnterpriseManager.properties*.

For more information about configuration, see these topics:

- Default Jetty configuration file (see page 60) and Modifying the Jetty-config.xml file (see page 60).
- http://docs.codehaus.org/display/JETTY/How+to+configure+SSL to access the Jetty website for additional information.

CA Technologies Pre-release Documentation. Authorized Use Only.

Understanding how HTTPS works

The Enterprise Manager web server supports HTTP tunneling over SSL.

This is done by introducing an optional Jetty 6.0.2 XML configuration file for the embedded Jetty web server in the Enterprise Manager. A default configuration file is provided that sets up a secure listener on port 8444, which allows the WorkStation and agent to access the Enterprise Manager using HTTPS tunneling through the Enterprise Manager web Server.

The Jetty configuration file is specified by the *introscope.enterprisemanager.webserver.jetty.configurationFile* property in *IntroscopeEnterpriseManager.properties*. The property specifies a path to a Jetty 6.0.2 XML configuration file. It can be absolute or a path relative to the *config* directory. The configuration file is loaded at startup and does not support hot configuration. The default value is *em-jetty-config.xml*. For example, it would appear as follows:

```
introscope.enterprisemanager.webserver.jetty.configurationFile=em-jetty-config.xml
```

When you configure the Jetty configuration properties in the *EM Web Server Settings* section of *IntroscopeEnterpriseManager.properties* the following rules apply:

- If the Jetty config file is successfully loaded, the default non-secure listener specified by *introscope.enterprisemanager.webserver.port* is not created.
- If a non-secure HTTP listener is desired in addition to a secure HTTPS listener, the default config file has a block of XML that can be uncommented to do this.

Default Jetty configuration file

A default *em-jetty-config.xml* is provided in the config directory.

Modifying the Jetty-config.xml file

You can modify many components of SSL using *Jetty-config.xml*.

For more information about Jetty, see http://docs.codehaus.org/display/JETTY/How+to+configure+SSL.

- Change the HTTPS port from the default 8444, by changing the value of the *port* attribute, such as in the following example:

  ```
  <Set name="port">8444</Set>
  ```

- The location of the config directory is specified in the XML by the *introscope.config system property*. This is set by the Enterprise Manager web server at startup and is designed as such to support a configurable config directory.

CA Technologies Proprietary/Confidential Documentation. Authorized Use Only.

- By default, the SSL listener is configured to use the keystore located at */internal/server/keystore*, a path relative to the *config* directory. This contains a self-signed certificate that is untrustworthy. It is highly recommended that you substitute your own keystore, containing a certificate signed by a recognized Certificate Authority (CA). See http://docs.codehaus.org/display/JETTY/How+to+configure+SSL for information about this.
- The default keystore password is password. Passwords can be obfuscated in the Jetty XML file. See http://docs.codehaus.org/display/JETTY/How+to+configure+SSL for information about this.
- By default, SSL is configured to accept self-signed certificates and not verify that the host name in the client's URL matches the host name in the digital certificate. This is to ensure that SSL works out of the box with the untrusted certificate in the default keystore. Customers needing strict authentication should create a keystore containing a trusted certificate, and set these attributes to in the XML, replacing the default keystore with their own:

```
<Set name="validateCertificates">true</Set>
<Set name="verifyHostnames">true</Set>
<Set name="keystore"><SystemProperty name="introscope.config"
default="./config" />/internal/server/keystore</Set>
<Set name="password">password</Set>
```

- Additionally, to require client authentication, create a truststore containing client certificates and set these attributes in the XML, replacing the default truststore with their own:

```
<Set name="needClientAuth">true</Set>
<Set name="truststore"><SystemProperty name="introscope.config"
default="./config" />/internal/server/keystore</Set>
<Set name="trustPassword">password</Set>
```

- If client authentication is required, agents and Workstations must be configured with a keystore containing a certificate trusted by the Enterprise Manager web server.

  To set the enabled cipher suites, set the *cipherSuites* attribute to a list of cipher suites:

```
<Set name="cipherSuites">
  <Array type="java.lang.String">
    <Item>SSL_DH_anon_WITH_RC4_128_MD5</Item>
  </Array>
</Set>
```

- The agents and Workstations must have an enabled cipher suite in common with the Enterprise Manager in order to use SSL. You might restrict the enabled cipher suites to prevent vulnerability to weak cipher suites, to allow anonymous connections, or to not encrypt data.

CA Technologies Internal/Restricted Distribution. Authorized Use Only.

Chapter 4: Configuring Enterprise Manager databases

This section describes configuration options for CA Introscope's internal databases.

For information about the APM database, see Administering the APM database (see page 63).

This section contains the following topics:

Administering the APM database (see page 63)
Configuring SmartStor data storage (see page 77)
Configuring Transaction Event database location and aging (see page 80)
Configuring location of baseline storage (see page 81)
Backing up SmartStor, baselines, and Transaction Events data (see page 81)
Using SmartStor tools to tune SmartStor data (see page 83)
Pruning SmartStor metadata (see page 90)
Trimming SmartStor data and metadata (see page 92)

Administering the APM database

This chapter contains information related to exporting and importing data to and from the APM database. You can import and export either the entire database, or just the configuration data.

For information about the other CA Introscope databases (SmartStor and the Transaction Event database), see Configuring Enterprise Manager databases (see page 63).

Backing up or restoring the database

Backup and restore capabilities include the ability to back up the APM database and configuration information.

For example, you can keep a copy of your database on a separate computer to protect your data. To do this, first backup the database, then copy the database backup file to another computer.

CA Technologies Prerelease Documentation. Authorized Use Only

Process for APM database backup and restore

This list shows the main tasks involved in protecting CA CEM data:

1. Understand the prerequisites (see page 64).
2. Back up the database on a regular basis (see page 65).
3. Restore the database if necessary (see page 66).

The backup and restore utilities for the PostgreSQL database come in the form of batch files (or shell scripts) that can be executed from the command line.

Note: The APM database includes CA CEM configuration settings, users and groups, business services and transaction definitions, as well as the application triage map data.

Before you backup or restore

WARNING! The restore function destroys the specified database if it already exists before it creates a new database with the name specified.

Both backup and restore require:

- The machine on which you are running a database script for backing up or restoring the PostgreSQL database must have PostgreSQL installed. If desired, you can run the script on a remote machine, as long as that remote machine also has PostgreSQL installed.
- Restoring a version 4.5 or earlier database is supported on Linux. Restoring a pre-4.5 database on Windows is not supported.

Before restoring the APM database on any of these platforms:

- OS/400 (IBMi)
- z/OS
- Solaris x86

Verify that that the JAVA_HOME environment variable has been set on the machine on which you will run the dbrestore script.

Set the JAVA_HOME environment variable and point it to the JRE root folder. JRE 1.6 or later is required, and the path cannot contain spaces. For example, if your JRE is installed in c:\program files\java\jre1.6.0, you set JAVA_HOME to the shortened version of the path which might be c:\progra~1\java\jre1.6.0.

CA Technologies Internal Use Documentation. Authorized Use Only

Backing up the database

The backup utility creates a *.backup* file of the APM database. You can customize the backup filename. If you do not specify a filename, the script creates a backup file with the name of the database and filename suffix of *.backup*.

To backup the APM database:

1. Navigate to the database-scripts directory for your operating system.
   - Linux: *<EM_Home>/install/database-scripts/unix/*
   - Windows: *<EM_Home>\install\database-scripts\windows*
2. Open a command line, and run the following script for your operating system:
   - Windows: *dbbackup-postgres.bat*
   - Linux: *dbbackup-postgres.sh*

   Use the following syntax and arguments:

   ```
   [dbbackup-postgres.bat | sh dbbackup-postgres.sh][dbserverhostip]
   [dbinstalldir] [dbname] [dbuser] [dbpassword] [dbport] [outputfile]
   ```

   Include the following arguments:

   dbserverhostIP

   The IP address of the machine hosting the database. If you are running it on the machine directly, you can specify *localhost*.

   dbinstalldir

   The directory location of the APM database.

   On Linux, this must be the Escaped or quoted version of the path if there are spaces in it. For Linux, an example is the directory */root/Introscope Enterprise Manager* can be represented as */root/Introscope\ Enterprise\ Manager*.

   If you didn't change the defaults, then the location is as follows:

   Linux: *opt/database*

   Windows: *<EM_Home>\database*

   database_name

   Specifies the database instance name. If you are using PostgreSQL and did not change the default, you should set this argument to *cemdb*. If you are using Oracle, you should set this argument to the Oracle system identifier (SID).

   dbuser

   Specifies the user name of the database user associated with the APM database. If you are using PostgreSQL and did not change the default, you should set this argument to *admin*. If you are using Oracle, you should set this argument to the APM database schema owner.

   Note: This is NOT the PostgreSQL database administrator user name.

CA Technologies Pre-Release Documentation. Authorized Use Only.

password

Specifies the password for the database user associated with the APM database.

port_number

Specifies the port number on which the database communicates to the Enterprise Manager and other components. If you are using PostgreSQL and did not change the default, you should set this argument to 5432. If you are using Oracle and the default Listener port, you should set this argument to 1521.

dbbackupdir

The existing directory location where you want to save the database backup file. For Linux, this must be the Escaped version of the path if there are spaces in it. For Linux, an example is the directory */root/Introscope Enterprise Manager* can be represented as */root/Introscope\ Enterprise\ Manager*.

There cannot be quotes around the directory and no ending slash.

outputfile

(Optional) The name of the database backup file. If you don't specify this parameter, the script generates a file with named *dbname.backup*, with the database name supplied. The backup filename suffix is *.backup* by standard convention, but it is not a requirement.

The following is an example that generates a cemdb.backup file in the /opt/database/backups directory:

```
sh dbbackup-postgres.sh 127.0.0.1 /opt/database cemdb admin quality 5432
./backups
```

3. When the backup process completes, a confirmation message displays.

   If the script encounters errors, here are common explanations:

   - The database name is incorrect.
   - PostgreSQL is not installed.
   - The location of the APM database is incorrect.

Restoring the database

The restore utility allows you to restore a previously backed up APM database backup file to the current system. For example, you can restore your database to recover after a system crash or a hard disk failure.

WARNING! The restore function destroys the specified database if it already exists before it creates a new database with the name specified.

Note: This script is not supported on Windows if you are restoring a database version 4.5 or earlier.

CA Technologies Proprietary Documentation. Authorized Use Only.

Note: If the database is on a UNIX machine, the restore process also upgrades the database.

Note: The amount of time it takes to restore a database is directly related to the database size. Large databases take longer to restore, and restore times also depend on machine capacity and memory. For example, a 1 GB database could take 1 to 2 hours to restore, depending on machine resources.

To restore the APM database:

1. Ensure that no users are connected to the database. You cannot restore the database if users are connected to it.
2. Navigate to the database-scripts directory for your operating system.
   - Linux: *<EM_Home>/install/database-scripts/unix/*
   - Windows: *<EM_Home>\install\database-scripts\windows*
3. Open a command line, and run the following script for your operating system:
   - Windows: *dbrestore-postgres.bat*
   - Linux: *dbrestore-postgres.sh*

   Use the following syntax and arguments:

   ```
   [dbrestore-postgres.bat | sh dbrestore-postgres.sh][dbserverhostip]
   [dbinstalldir] [dbserviceuser] [dbservicepassword] [dbname] [dbuser]
   [dbpassword] [dbport] [backupfile]
   ```

   Include the following arguments:

   dbserverhostIP

   The IP address of the machine hosting the database. If you are running it on the machine directly, you can specify *localhost*.

   dbinstalldir

   The directory location of the APM database.

   On Linux, this must be the Escaped or quoted version of the path if there are spaces in it. For Linux, an example is the directory */root/Introscope Enterprise Manager* can be represented as */root/Introscope\ Enterprise\ Manager*.

   If you didn't change the defaults, then the location is as follows:

   Linux: *opt/database*

   Windows: *<EM_Home>\database*

   dbserviceuser

   The PostgreSQL administrator username. By default, this is *postgres*.

CA Technologies Pre-release Documentation. Authorized Use Only.

dbservicepassword

The password associated with *dbserviceuser*. (The PostgreSQL administrator password.)

database_name

Specifies the database instance name. If you are using PostgreSQL and did not change the default, you should set this argument to *cemdb*. If you are using Oracle, you should set this argument to the Oracle system identifier (SID).

dbuser

Specifies the user name of the database user associated with the APM database. If you are using PostgreSQL and did not change the default, you should set this argument to *admin*. If you are using Oracle, you should set this argument to the APM database schema owner.

Note: This is NOT the PostgreSQL database administrator user name.

password

Specifies the password for the database user associated with the APM database.

port_number

Specifies the port number on which the database communicates to the Enterprise Manager and other components. If you are using PostgreSQL and did not change the default, you should set this argument to 5432. If you are using Oracle and the default Listener port, you should set this argument to 1521.

backupfile

The directory location and name of the database backup file.

The following is an example that restores a *cemdb.backup* file from the */opt/database/backups* directory:

```
sh dbrestore-postgres.sh 127.0.0.1 /opt/database postgres C@wilyapm90 cemdb
admin quality 5432 ./backups/cemdb.backup
```

4. When the restore process completes, a confirmation message displays.

   If you encounter errors:

   - The most common error is that users are still connected to the database. You can restart the PostgreSQL server to bump off errant connections.
   - If you get an error that says that database does not exist, this is okay. This means that the script is being used to create a new database in the process.

CA Technologies Pre-release Documentation: Authorized Use Only

Exporting or importing database configuration

In some cases, it is not possible to place the CA CEM system in a production environment for the first installation. For example, perhaps it is company policy to install all new hardware and software in a test lab prior to entering the production network.

Once you have established and tested your CA CEM configuration in a QA or test environment, you can easily move that proven configuration into production via the CA CEM export and import.

Configuration information includes users and user groups and business services and transactions.

Process for database configuration export and import

This list shows the main tasks involved in exporting the CA CEM configuration:

1. Understand the prerequisites (see page 69).
2. Export the (see page 70)CA CEM configuration (see page 70).
3. Import the (see page 71)CA CEM configuration (see page 71).

The export and import utilities come in the form of batch files (or shell scripts) that can be executed from the command line.

For information about importing and exporting business transaction definitions, see the *CA APM Transaction Definition Guide*.

For information about backing up and restoring all information (not only configuration information) in the APM database, see Backing up or restoring the database (see page 63).

Before you export or import database configuration data

WARNING! It is highly recommended that you back up the APM database before attempting to run the two utilities (especially import, because it destroys the existing configuration before rebuilding it). The import script destroys the database, and the result is a database that contains only the configuration data for the old database. See Backing up or restoring the database (see page 63) for more information.

Both export and import require:

- You run the scripts on the machine where the APM database is installed.
- You run the scripts in the location where the Enterprise Manager installer put them because they refer to files that are relative to these locations.

CA Technologies Pre-release Documentation
Authorized Use Only.

- Before running these scripts, turn off all connections to the APM database. Shut down the Enterprise Manager, close any database administration tools (for example, pgAdmin), verify that all users are logged off or have closed the CEM console. Although you do not have to expressly disable the TIMs, be aware you will need to re-enable them after import.
- Set the JAVA_HOME environment variable and point it to the JRE root folder. JRE 1.6 or later is required, and the path cannot contain spaces. For example, if your JRE is installed in *c:\program files\java\jre1.6.0*, you set JAVA_HOME to the shortened version of the path which might be *c:\progra~1\java\jre1.6.0*.
- If you are importing the configuration from a previous version of CA CEM, be aware that you can import any previous version on Linux, and only a 5.0 version (or higher) on Windows. The import utility automatically upgrades the database schema from the older release to the current release.

Exporting the APM database configuration

The *configexport.bat* or *configexport.sh* script creates an XML file with all the CA CEM configuration information. Configuration information includes users and user groups and business services and transactions. Exporting configuration information enables you to move data from test or lab systems as a base for a production system.

To export the APM database configuration:

1. Set the JAVA_HOME environment variable.
2. Navigate to the *<EM_Home>/install/database-scripts* directory for your operating system. For example:
   `cd /home/Introscope9.0.5.0/install/database-scripts/unix`
3. Open a command or terminal window and run the *configexport.bat* on Windows or *configexport.sh* on other platforms with the following arguments:
   `ConfigExport <dbhostIP> <database_name> <dbuser> <password> <port_number> <databasetype> <export-file>`

   dbhostIP

   Specifies the IP address of the computer hosting the database. You cannot use localhost for this command.

CA Technologies Pre-release Documentation. Authorized Use Only

database_name

Specifies the database instance name. If you are using PostgreSQL and did not change the default, you should set this argument to cemdb. If you are using Oracle, you should set this argument to the Oracle system identifier (SID).dbuser

Specifies the user name of the database user associated with the APM database. If you are using PostgreSQL and did not change the default, you should set this argument to admin. If you are using Oracle, you should set this argument to the APM database schema owner.

Note: This is NOT the PostgreSQL database administrator user name.

password

Specifies the password for the database user associated with the APM database.

port_number

Specifies the port number on which the database communicates to the Enterprise Manager and other components. If you are using PostgreSQL and did not change the default, you should set this argument to 5432. If you are using Oracle and the default Listener port, you should set this argument to 1521.

databasetype

Specifies the type of database you are using. This argument should be set to postgres or oracle.

export-file

Specifies the absolute path and file name for the XML file that contains the exported data.

For example, to run the *configexport* program on Linux, you would use a command line similar to the following:

```
./configexport.sh 197.168.144.10 arcade apmdba mypwd 1521 oracle
/home/apmdba/export_oracle.xml
```

4. When the configuration export process completes, a confirmation message displays.

Importing the APM database configuration

You can import the configuration settings from one APM database to another. For example, you can use the *configimport.bat* or *configimport.sh* program to import the database configuration from a QA system onto a production system. The configuration settings to be imported include users and user groups and business services and transactions.

CA Technologies Prerelease Documentation
Authorized Use Only

The APM database configuration import process does the following:

- deletes the previous database configuration information (about operators, users, groups, business service process and transaction definitions) and replaces it with the imported configuration data.
- deletes defect, report, and statistical data and does not replace it.

Note: You can import configuration from an earlier version of CA CEM. For example, you can download and export a configuration from CA CEM 4.1 and then upload and import it into CA CEM 4.2 (Linux only).

To import the APM database configuration:

1. Navigate to the *<EM_Home>/install/database-scripts* directory for your operating system. For example:
   `cd /home/Introscope9.0.5.0/install/database-scripts/unix`
2. Open a command or terminal window and run the *configimport.bat* on Windows or *configimport.sh* on other platforms with the following arguments:
   `ConfigImport -dbhost <hostname> -dbname <database_name> -dbport <port_number> -databasetype <database_type> -dbuser <dbuser> -dbpassword <password> -dbscriptsdir <dir> -importfile <import-filename> -dbserviceuser <dbserviceuser> -dbservicepwd <dbservicepwd> [-postgresinstalldir <postgresinstalldir>] [-nolog] [-jdbcbatchsize <jdbcbatchsize>] [-promptbeforeimport true|false] [-is64Bit true|false]`

   dbhost

   Specifies the host name or IP address of the computer hosting the database. You cannot use *localhost* for this command.

   dbname

   Specifies the database instance name. If you are using PostgreSQL and didn't change the default, you should set this argument to *cemdb*. If you are using Oracle, you should set this argument to the Oracle system identifier (SID).

   dbport

   Specifies the port on which the database communicates to the Enterprise Manager and other components. If you are using PostgreSQL and didn't change the default, you should set this argument to 5432. If you are using Oracle and the default Listener port, you should set this argument to 1521.

   databasetype

   Specifies the type of database you are using. This argument should be set to *postgres* or *oracle*.

CA Technologies Proprietary and Confidential Documentation
Authorized Use Only

dbuser

Specifies the user name of the database user associated with the APM database. If you are using PostgreSQL and didn't change the default, you should set this argument to *admin*. If you are using Oracle, you should set this argument to the APM database schema owner.

dbpassword

Specifies the password for the database user associated with the APM database.

dbscriptsdir

Specifies the directory for database specific SQL scripts. For example, you should use *<EM_Home>/install/oracle/database-scripts* if you are using an Oracle database schema or *<EM_Home>/install/database-scripts* if you are using a PostgreSQL database.

importfile

Specifies the absolute path and file name for the XML file that contains the data to be imported.

dbserviceuser

Specifies the PostgreSQL administrative service user account. The default administrative service account is *postgres*. The argument is not applicable for Oracle.

dbservicepwd

Specifies the password for the PostgreSQL administrative service user account.

postgresinstalldir

(Optional) Specifies the directory location of the APM database when using PostgreSQL. On Linux or Solaris, the path must be enclosed in quotes or include escape sequences if the path has spaces in it. For example, if the directory is */root/Introscope Enterprise Manager*, you can specify the location as *"/root/Introscope Enterprise Manager"* or /root/Introscope\ Enterprise\ Manager.

If you didn't change the defaults, then the location is *opt/database* on 32-bit Linux or Solaris, *opt/database/postgres/8.3-community* on 64-bit Linux or Solaris, or *<EM_Home>\database* on Windows.

The argument is not applicable for Oracle.

nolog

(Optional) Specifies that you want to turn off logging to standard output (StdOut).

CA Technologies Prerelease Documentation
Authorized Use Only

jdbcbatchsize

(Optional) Specifies the maximum size of data packets transferred at a time through Java database connections (JDBC). The default value is 2500.

promptbeforeimport

(Optional) Specifies whether you want the *configimport* program to display a prompt before it begins importing data. You can set this parameter to *true* if you want a prompt displayed. The default value is *false*.

is64Bit

(Optional) Specifies whether the PostgreSQL database is 32-bit or 64-bit. You can set this parameter to *true* if you are importing data into a 64-bit PostgreSQL database. This argument is not applicable for Oracle databases.

For example, to run the *configimport* program on Linux and a PostgreSQL database, you would use a command line similar to the following:

```
./configimport.sh -dbhost sfqa07 -dbname cemdb -dbport 5432 -databasetype
postgres -dbuser admin -dbpassword mypassword -dbscriptsdir
/apps/Introscope/install/database-scripts -importfile /tmp/apm_data.xml
-dbserviceuser postgres -dbservicepwd %pas5w0rd -postgresinstalldir
/opt/database -promptbeforeimport true
```

The import script performs an initial validation check to ensure the XML is properly formatted.

The import script offers you the option to validate the XML more thoroughly to ensure the XML content can be processed properly. Enter Y to perform the validation, or enter N to continue with the import without validating the XML content.

Note: This second validation can take awhile, depending on the size of the XML file. It's important to ensure the XML data can be properly imported; otherwise, errors occur.

The script displays progress messages, and displays a message when validation completes, import begins, and when import completes.

If either validation check fails, an error message displays and the configuration import does not proceed. If both validation checks succeed, the script imports the specified database configuration data.

CA Technologies Prerelease Documentation. Authorized Use Only

Moving the APM database to another operating system

In some cases, you may want to move the APM database from one computer to another or from one platform to another as part of your upgrade process. For example, you may want to move the APM database from PostgreSQL on Windows or Linux to PostgreSQL on Solaris. To move the database from one platform to another, you must export the business transaction and configuration data from a source database then import the business transaction and configuration data into a the target database on the new platform. Alternatively, you can use the *dbbackup-postgres* and *dbrestore-postgres* utilities to create a complete backup of an existing PostgreSQL database and restore the database on a new computer.

For more information about exporting and importing data from a PostgreSQL database or using the *dbbackup-postgres* and *dbrestore-postgres* utilities, see the *CA APM Configuration and Administration Guide*. For information about migrating data from a PostgreSQL database to an Oracle database (see page 75), see the *CA APM Installation and Upgrade Guide*.

Migrating data from a PostgreSQL database to an Oracle database

If you are upgrading from CA APM 9.0, the APM database schema is a PostgreSQL database. After upgrading, you can choose to migrate the APM database to Oracle, if desired. For example, if your organization is standardized on Oracle as a database platform, you may need to migrate the APM database to Oracle to enable production-level monitoring. If you want to move your existing data from a source database in PostgreSQL to a target database in Oracle, you can use the CA APM *migration* utility to specify the source and target databases.

To migrate from a PostgreSQL database to an Oracle database:

1. Plan and prepare for the upgrade by reviewing the upgrade scenarios and upgrade paths.
2. Create a backup copy of all components.
3. Upgrade the Enterprise Manager and APM database using the interactive or silent Enterprise Manager installer.
4. Navigate to the *<EM_Home>/install/migration* directory.

CA Technologies Internal Use Only. Authorized Use Only. Documentation

5. Run the *migration.bat* or *migration.sh* script as appropriate for your operating environment using the following arguments:

   `migration -srcDatabaseName <database_name> -srcDatabaseType <database_type> -srcHost <hostname> -srcPort <port_number> -srcUser <username> -srcPassword <password> -tgtDatabaseName <datatargetname> -tgtDatabaseType <database_type> -tgtHost <hostname> -tgtPort <port_number> -tgtUser <username> -tgtPassword <password>`

   srcDatabaseName

   Specifies the name of the source database. If you used the default settings, the database name is *cemdb*.

   srcDatabaseType

   Specifies the type of the source database. In this release, the only valid value is *postgres*.

   srcHost

   Specifies the host name or IP address of the database server for the source database.

   srcPort

   Specifies the port number for connecting to the source database. If you used the default settings, the port is 5432.

   srcUser

   Specifies the database user name for the source database. If you used the default settings, the database user name is *admin*.

   srcPassword

   Specifies the password for the database user in the source database.

   tgtDatabaseName

   Specifies the name or service account identifier of the target database. For an Oracle database, the name should be the unique system identifier (SID) for the target database.

   tgtDatabaseType

   Specifies the type of the target database. In this release, the only valid value is *oracle*.

   tgtHost

   Specifies the host name or IP address of the database server for the target database.

   tgtPort

   Specifies the port number for connecting to the target database. If you used the default settings, the port is 1521.

CA Technologies Prerelease Documentation. Authorized User Only.

tgtUser

Specifies the database user name for the target database.

tgtPassword

Specifies the password for the database user in the target database.

Note: All arguments are case-sensitive and required to run the migration tool.

For example, to migrate data from the PostgreSQL database *cemdb* to an Oracle database with the SID *orcl* with the schema owner of *apmdb* on a Linux computer, the command line might be similar to the following:

```
./migration.sh -srcDatabaseName cemdb -srcDatabaseType postgres -srcHost nyc16
-srcPort 5432  -srcUser admin -srcPassword quality -tgtDatabaseName orcl
-tgtDatabaseType oracle -tgtHost nyc23 -tgtPort 1521 -tgtUser apmdb -tgtPassword
cent3R
```

6. Check the *migration.log* file in the *logs* directory to verify the schema has been migrated successfully or to troubleshoot problems with the migration. For example, navigate to the *<EM_Home>/install/migration/logs* directory.

   Note: You can modify the location of the *migration.log* file by modifying the *log4j.appender.logfile.File* property in the *<EM_Home>/install/migration/log4j-migration.properties* file. For more information about customizing migration and log file properties, see Customizing migration properties.

Configuring SmartStor data storage

SmartStor records all application performance data (Introscope metrics) at all times while it is in production, without the need for an external database. This enables you to analyze historical data, to identify root causes of application downtime, or perform capacity analysis-without the need for an external database. SmartStor is enabled with default settings when installing Introscope. However, you can disable SmartStor. For more information about disabling SmartStor, see introscope.enterprisemanager.smartstor.enable (see page 416). SmartStor requires about 10 GB of disk storage. SmartStor data is set to age-out over time so that the data store does not get excessively large. By default, SmartStor continuously records all Introscope data in the *<EM_Home>/data* directory.

SmartStor has three different maintenance sub-processes that transform metric data into persistent storage flat files. In addition, asynchronous threads perform related data processing.

The first sub-process takes raw files with names that end in *.spool and converts them to *.data type files. This happens hourly, on the hour.

The second sub-process is the transformation of hour-long data files into day long data files. This happens at midnight every day.

CA Technologies Prerelease Documentation. Unauthorized Use Only.

The third sub-process transforms, reperiodizes, and archives data that is more than one day old according to the tiers defined in the *IntroscopeEnterpriseManager.properties* file. This is also part of the SmartStor maintenance that runs on the Enterprise Manager at midnight.

Here are the default tier settings in the *IntroscopeEnterpriseManager.properties* file:

introscope.enterprisemanager.smartstor.tier1.frequency=15

introscope.enterprisemanager.smartstor.tier1.age=7

introscope.enterprisemanager.smartstor.tier2.frequency=60

introscope.enterprisemanager.smartstor.tier2.age=30

introscope.enterprisemanager.smartstor.tier3.frequency=900

introscope.enterprisemanager.smartstor.tier3.age=365

SmartStor uses a storage hierarchy that assumes that recent data is accessed the most and needs to have the highest resolution. Here are the data resolutions based on the default property settings listed above:

Store 1 week of data in 15 second chunks (uncompressed)

Store 30 days of data in minute (60 seconds) chunks (compressed)

Store 1 year of data in 15 minute (900 seconds) chunks (compressed)

Important! You must configure SmartStor to write to a dedicated disk. This avoids severely affecting Enterprise Manager performance and responsiveness, which could be caused by disk contention with other Enterprise Manager activities or with other processes running on the host.

When you set up the dedicated disk for SmartStor data, set the property to *true* (the Enterprise Manager installer sets this property to *false* by default). See introscope.enterprisemanager.smartstor.dedicatedcontroller (see page 416).

These sections provide SmartStor planning and configuration instructions:

- Determining SmartStor system requirements (see page 79)
- Specifying SmartStor data tier values (see page 79)

CA Technologies Proprietary Documentation. Unauthorized use only.

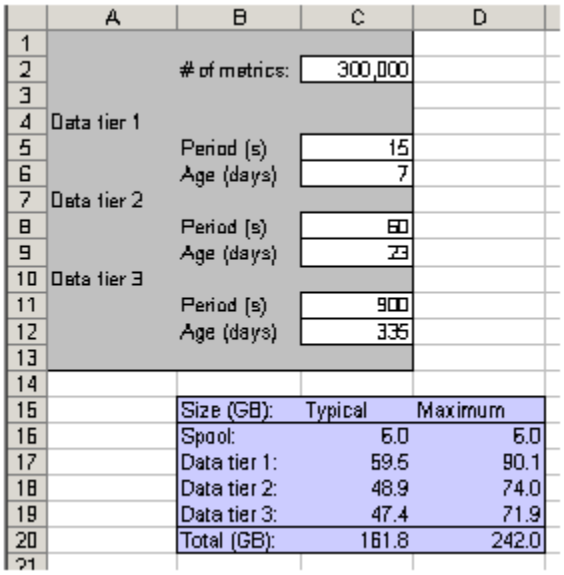
Determining SmartStor system requirements

Your installation includes a spreadsheet to help determine SmartStor system requirements:
<EM_Home>/docs/SmartStorSizing.xls

You can experiment with different values for number of metrics, and different periods and aging ranges for your data tiers. The resulting GB size usage appears in the table below.

SmartStor's performance is sensitive to the operating system, JVM, and file system choices.

This is the SmartStor Sizing Tool formula window:



	A	B	C	D
1				
2		# of metrics:	300,000	
3				
4	Data tier 1			
5		Period (s)	15	
6		Age (days)	7	
7	Data tier 2			
8		Period (s)	60	
9		Age (days)	23	
10	Data tier 3			
11		Period (s)	900	
12		Age (days)	335	
13				
14				
15		Size (GB):	Typical	Maximum
16		Spool:	6.0	6.0
17		Data tier 1:	59.5	90.1
18		Data tier 2:	48.9	74.0
19		Data tier 3:	47.4	71.9
20		Total (GB):	161.8	242.0
21				

After you determine the appropriate data tier frequency and aging periods, you enter your specifications into the appropriate SmartStor properties.

Specifying SmartStor data tier values

Three data tiers configure SmartStor data storage. Each tier consists of a property pair:

- the storage frequency (in seconds)
- the data aging period (in days)

The rules for the frequency and aging values are:

- The storage frequency must be a multiple of 15 seconds.
- The frequency must not be greater than 1800 seconds (30 minutes).
- Each tier's frequency must be an even multiple of the previous tier's frequency (this multiple can be one; each tier can have the same frequency).
- The age must not be zero.
- All three tiers must be defined, and you cannot define more than three tiers.

To configure the SmartStor data tier values:

1. Open the *IntroscopeEnterpriseManager.properties* file.
2. For the *introscope.enterprisemanager.smartstor.tier1.frequency* property, enter the appropriate value.
3. For the *introscope.enterprisemanager.smartstor.tier1.age* property, enter the appropriate value.
4. Repeat for the remaining two data tier property pairs.
5. Restart the Enterprise Manager.

Note: If you change your data tiers and reboot the Enterprise Manager, SmartStor attempts to adjust data to your new configuration, but it might lose a day's worth of data.

Configuring Transaction Event database location and aging

The Enterprise Manager stores transaction traces and other event data in the Transaction Event database.

The location of the Transaction Event database is configured during the Enterprise Manager installation. Unless you identify a different location during installation, the database is created in the *<EM_Home>/traces* directory. After installation, you can choose a different location with the *introscope.enterprisemanager.transactionevents.storage.dir* property in *IntroscopeEnterpriseManager.properties*.

The number of days that Introscope retains transaction event data is configured during the Enterprise Manager installation. The default value of 14, can be changed using the *introscope.enterprisemanager.transactionevents.storage.max.data.age* property.

Note: Increasing length of time that Introscope saves event data increases system overhead and disk storage requirements.

CA Technologies Pre-release Documentation
Authorized Use Only.

Configuring location of baseline storage

By default, Introscope stores metric baseline information on the Enterprise Manager machine, in *<EM_Home>/data/baselines.db*.

The Enterprise Manager installer does not offer the option of selecting a different location for baseline storage. However, you can define a different location for *baselines.db* in the *introscope.enterprisemanager.dbfile* property in the Enterprise Manager properties file.

For information about how Introscope uses metric baselines to enable high-level indicators of application health, see the *CA APM Workstation User Guide*.

Backing up SmartStor, baselines, and Transaction Events data

You can backup and work with Enterprise Manager SmartStor, baselines, and Transaction Events data files after stopping the Enterprise Manager.

By default the Enterprise Manager stores all SmartStor and baselines data in the *<EM_Home>/data* directory.

The Enterprise Manager data files for SmartStor are:

- the *metrics.metadata* file, which tracks information about the various metrics that the agents have reported. Most importantly, *metrics.metadata* tracks what the metrics are named).
- files ending with *.spool* (for example, *1251226800000.spool*) that hold the 15-second metric values sent to SmartStor since the Enterprise Manager performed the last spool-to-query conversion. The *.spool* file never holds more than one hour's worth of data.
- files ending with *.data* (for example, *1251183600000.data*) that hold metric values in SmartStor's read-only query optimized format. These files hold data point values for 15 seconds, 1 minute, and 1 hour.

The Enterprise Manager data file for baselines data is:

- *baselines.db*.

CA Technologies Prerelease Documentation. Authorized Use Only

By default the Enterprise Manager stores all Transaction Trace and events data. in the *<EM_Home>/traces* directory. The Enterprise Manager data file for Transaction Events data is:

- *traces_YYYYMMDD.db*. For example: *traces_20090923.db*.

While the Enterprise Manager is running, it expects to be able to create and write *metrics.metadata, .spool*, and *.data* files without error. Therefore CA Technologies strongly recommends that you stop the Enterprise Manager before moving, copying, and working with Enterprise Manager data.

Warning! Copying or editing Enterprise Manager data files while the Enterprise Manager is running may result in data loss.

When an Enterprise Manager is shut down, you can do all of the following to backup and work with Enterprise Manager data:

- save data to backup media
- edit SmartStor data using Smartstor tools
  For more information about SmartStor tools, see Using SmartStor tools to tune SmartStor data (see page 83).
- compress data
- restore data
- copy data for offline viewing by another Enterprise Manager.

CA Technologies Prerelease Documentation. Authorized Use Only.

Creating a query-only duplicate cluster

You can copy SmartStor data files to duplicate cluster machines to serve as a query-only source historical (non-live) data.

Note: CA Technologies recommends creating duplicate clusters only in production environments where the service demand considerably exceeds the capacity of the hardware on which the cluster is running.

To create a query-only duplicate cluster:

1. Create an exact duplicate of your production or test cluster.

   Use one machine for the MOM and one for each Collector in your cluster.
2. Shut down your production cluster.
3. Copy the SmartStor, baselines, and Transaction Events data files from the production or test MOM to the duplicate MOM and from each of the production or test Collectors to the duplicate Collectors.
4. Start up both the production cluster and the query-only duplicate cluster.
5. Run any desired queries on the query-only duplicate cluster.
6. Repeat necessary steps once a day to keep current the data on the query-only duplicate cluster.

Using SmartStor tools to tune SmartStor data

Introscope provides SmartStor command-line tools that enable you to:

- clean up from metric leaks and explosions

  For information about metric leaks and explosions, see the *CA APM Sizing and Performance Guide.*
- split a single Enterprise Manager into multiple Enterprise Managers
- merge multiple Enterprise Managers into a single Enterprise Manager.

CA Technologies Prerelease Documentation. Authorized Use Only.

The SmartStor tools start scripts are located in the *<EM_Home>/tools/* directory. Run SmartStor tools in a Windows environment using the *SmartStorTools.bat* start script. Run SmartStor tools in a UNIX environment using the *SmartStorTools.sh* start scrip*.*

When you run the SmartStor tools commands from the *<EM_Home>* command line using valid arguments, the tools perform the requested commands. If you run the SmartStor tools commands using invalid arguments, the script displays a help message.

WARNING! The source Enterprise Manager *must* be shut down when using SmartStor tools. CA Technologies recommends shutting down the source and destination Enterprise Managers when running SmartStor tools. If an Enterprise Manager is running when SmartStor tools are used, there is a risk of SmartStor corruption. This can occur if the Enterprise Manager is writing data to SmartStor when SmartStor tools are also making changes to SmartStor data.

Run these SmartStor tools commands using this format:

<SmartStorTools start script> <command> <arguments>

- merge
- remove_metrics
- remove_agents
- keep_agents
- prune
- list-agents

WARNING! Use the prune tool with caution. Since this tool deletes data, take care that you do not accidentally delete important data.

- test_regex

SmartStor tools commands

SmartStor tools commands enable you to perform several different maintenance and administration tasks on the SmartStor database.

Important! When specifying the names of agents and metrics in these commands, special characters must be escaped using a backslash. Special characters are *([{\^-$|]})?*+.* For example, the metric name *SuperDomain|machineName|WebLogic|WebLogic* would be specified *SuperDomain\|machineName\|WebLogic\|WebLogic*.

CA Technologies Prerelease Documentation Authorized Use Only.

merge

Merges a specified SmartStor source directory into a different destination directory.

Arguments:

- <destination SmartStor directory>
- -help Prints a help message
- <source SmartStor directory>

Examples:

Windows:

```
SmartStorTools.bat merge <source directory> <destination directory>
```

For example:

```
SmartStorTools.bat merge <EM_Home>\data  C:\data
```

When the command executes successfully, all the SmartStor data present in the source directory is merged with the data in the destination directory.

UNIX:

```
SmartStorTools.sh merge <source directory> <destination directory>
```

CA Technologies Prerelease Documentation. Authorized Use Only.

remove_metrics

Removes metrics matching the given regular expression. Use this tool to clean up metrics leaks and explosions.

WARNING! Use the remove_metrics tool with caution. Since this tool deletes data, take care that you do not accidentally delete important data.

Arguments:

- *-dest* —Destination SmartStor directory
- *-help* —Prints a help message
- *-metrics* —Regular expression for metrics to remove
- *-src* —Source SmartStor directory

Examples:

Windows

```
SmartStorTools.bat remove_metrics –dest <destination directory> -metrics <metrics to be removed> -src <source directory>
```

For example:

```
<EM_Home>\tools\SmartStorTools.bat remove_metrics –dest C:\SWDump\destination -metrics ".*Socket.*" -src <EM_Home>\data
```

When the command executes successfully, all the metrics except the socket metrics are present in the destination directory and the source directory has all the metrics intact.

UNIX

```
SmartStorTools.sh remove_metrics –dest < destination directory> -metrics <metrics to be removed> -src <source directory>
```

CA Technologies Pre-release Documentation. Authorized Use Only.

remove_agents

Removes agents matching the given regular expression.

Warning! Use the remove_agent tool with caution. Since this tool deletes data, take care that you do not accidentally delete important data.

Arguments:

- *-agents* —Regular expression for agents to remove
- *-dest* —Destination SmartStor directory
- *-help* —Prints a help message
- *-src* —Source SmartStor directory

Examples:

Windows

```
SmartStorTools.bat remove_agents –agents <Agent name> –dest <destination directory>
-src <source directory>
```

For example:

```
<EM_Home>\tools>SmartStorTools.bat remove_agents –agents  ".*WebSphere.*" -dest
C:\SWDump\destination -src <EM_Home>\data
```

When the command executes successfully, all the metrics related to WebSphere agents are removed from the destination directory and all the metrics in the source directory remain intact.

UNIX

```
SmartStorTools.sh remove_agents –agents  ".*WebSphere.*" -dest <destination
directory> -src <source directory>
```

CA Technologies Prerelease Documentation. Authorized Use Only

keep_agents

Keeps agents matching the given regular expression

Arguments:

- *-agents* —Regular expression for agents to keep
- *-dest* —Destination SmartStor directory
- *-help* —Prints a help message
- *-src* —Source SmartStor directory

Examples:

Windows

```
SmartStorTools.bat keep_agents –agents  <agent(s) name>  –dest <destination
directory>  -src <source directory>
```

For example:

```
<EM_Home>\tools>SmartStorTools.bat keep_agents -agents ".*Websphere.*" -dest
C:\SWDump\destination\keepagents -src data
```

When the command executes successfully, all the metrics related to Websphere agents are present in the destination directory and the contents of source directory are intact.

UNIX

```
SmartStorTools.sh keep_agents -agents <Agent(s) name> -dest <destination directory>
-src <source directory>
```

prune

Finds agents that have no historical data in SmartStor, and removes all associated metric metadata. Use this tool to clean up metrics leaks and explosions.

WARNING! Use the prune tool with caution. Since this tool deletes data, take care that you do not accidentally delete important data.

Using the prune tool requires several steps. See Pruning SmartStor metadata (see page 90).

Arguments:

- *-dest* — directory containing pruned metadata file along with unchanged *.data files of src.
- *-help* — prints a help message

CA Technologies Prerelease Documentation. Authorized Use Only.

- *-silent* — silent mode
- *-src* — source SmartStor directory

Examples

Windows

```
SmartStorTools.bat Prune  –dest <destination directory>  -src <source directory>
```

For example:

```
<EM_Home>\tools>SmartStorTools.bat Prune -dest C:\SWDump\destination\backup -src
data
```

When the command executes successfully, metadata files are backed up to the *dest* directory with the name *metrics.metadata.*

UNIX

```
SmartStorTools.sh prune -dest <destination directory> -src <source directory>
```

test_regex

Prints to the Windows or UNIX console the metrics or agents matching the given regular expression.

Arguments:

- *-help* —Prints a help message
- *-src* —Source SmartStor directory
- -agents —Regular expression for agents to match
- *-metrics* —Regular expression for metrics to match

Examples:

Windows

```
SmartStorTools.bat test_regex -agents <agent(s) name> –src <sourcedirectory>
```

For example:

```
<EM_Home>\tools>SmartStorTools.bat test_regex -agents ".*Websphere.*" -src
<EM_Home>\data
```

When the command executes successfully, all the metrics for WebSphere agent are printed.

UNIX

```
SmartStorTools.sh test_regex -agents <agent(s) name> -src <source directory>
```

CA Technologies Prerelease Documentation. Authorized Use Only.

list_agents

Lists the agents matching a given regular expression.

The SmartStor tool list_agents tools lists the agents matching the given regular expression.

Arguments:

- *-agents* —Regular expression for agents to match
- *-src* —Source SmartStor directory
- *-help* —Prints a help message

Examples:

Windows

```
SmartStorTools.bat list_agents –agents  <agent(s) name>  -src <source directory>
```

For example:

```
<EM_Home>\tools>SmartStorTools.bat list_agents -agents ".*Websphere.*" -src
<EM_Home>\data
```

When the command executes successfully, all the agent names matching the regular expression (in this case ".*WebSphere.*") are listed with their last reported date.

UNIX

```
SmartStorTools.sh list_agents -agents <Agent(s) name> -src <source directory>
```

Pruning SmartStor metadata

Use the *prune* SmartStor tool to finds agents that have no historical data in SmartStor, then remove all associated metric metadata. Use this tool to clean up metrics leaks and explosions.

Important! Because Enterprise Managers create SmartStor data when running, CA Technologies recommends shutting down the Enterprise Manager when running SmartStor tools. Using SmartStor tools while the Enterprise Manager is running produces unpredictable results.

To prune SmartStor metadata:

1. Shut down the Enterprise Manager.
2. From a command prompt, run the *prune* command.

CA Technologies Pre-release Documentation. Authorized Use Only

For example:

On Windows: *<EM_Home>\tools>SmartStorTools.bat prune -src <path to data directory> -dest <path to dest backup directory>*

On UNIX: *<EM_Home>/tools>SmartStorTools.sh prune -src <path to data directory> -dest <path to dest backup directory>*

When the command executes successfully:

- metadata files are backed up to the *dest* directory with the name *metrics.metadata.*
- the Windows or UNIX console displays the pruned metrics during the prune operation
- *.data files will be copied to *dest,* without any change, from the *src* directory.

To redirect the Windows or UNIX console pruned metrics list to a text file, include the *>>* argument in the command.
For example: *<EM_Home>\tools>SmartStorTools.bat prune -src <path to data directory> -dest D:\backup >>D:\console.txt*

1. Rename the *metrics.metadata* file in the *<EM_Home>/data* directory. For example: *noprunemetrics.metadata*.
2. Run the *test_regex* command on the *<EM_Home>/data* directory.

   The Windows or UNIX console displays the command output showing all the metrics based on the regular expression you included in the command. Notice that the pruned metrics are no longer present.

   For example, if you were pruning a large number of socket metrics, you run this *test_regex* command:

   On Windows:
   ```
   <EM_Home>\tools>SmartStorTools.bat test_regex -metrics ".*Socket.*" -src <path to data directory>
   On UNIX:
   <EM_Home>/tools>SmartStorTools.sh test_regex -metrics ".*Socket.*" -src <path to data directory>
   ```
3. Restart the Enterprise Manager.

CA Technologies Pre-release Documentation. Authorized Use Only.

Trimming SmartStor data and metadata

Use the SmartStor tools to trim excess SmartStor data.

Important! CA Technologies recommends shutting down the Enterprise Manager when running SmartStor tools. Because Enterprise Managers create SmartStor data when running, if you use SmartStor tools, the results can be unpredictable.

For an example about how to find a SmartStor data problem, see Example: Using the test_regex command to find a SmartStor data problem (see page 94). Also see the topics about metrics leaks and explosions in the *CA APM Sizing and Performance Guide*.

To trim SmartStor data and metadata:

1. Shut down the Enterprise Manager.
2. From a command prompt, run the *test_regex* command to print the current metric list.

   For example, if Introscope is installed in the *<EM_Home>* directory and SmartStor is configured to save data in the *<EM_Home>\data* directory, run this command:

   <EM_Home>\tools>SmartStorTools.bat test_regex -agents -src <EM_Home>\data

   The result is a list of fully-qualified metric names, including the host name, agent name, metric name, and metric IDs. The total line count gives an indication of total number of metrics, although it is not as accurate as the metric count you get from the supportability metrics in the Investigator.

   By doing line counts against specific regular expressions matching an agent or metric name, you can get an idea of what metrics are occupying the majority of the space. This also helps identify what metrics you are not interested in and would like to remove.
3. From the list of metrics, pick metrics from specific agents to be removed.
4. Run the SmartStor tools *remove_metrics* command to remove a set of excess metrics.

   For example, if you have high SQL metric counts, use this command to remove all the SQL metrics:

   ```
   <EM_Home>\tools>SmartStorTools.bat remove_metrics
   –dest c:\sql_metrics_removed -metrics ".*SQL.*" -src <EM_Home>\data
   ```

   When the command executes successfully, all the metrics except the SQL metrics are present in the destination directory and the source directory has all the metrics intact.

   Note: This command needs approximately as much target space (although it will shrink by a few GB) to run as the space occupied by the original data (specified by the *–src* option). In the example above, the new data without the SQL metrics (specified by the *–dest* option) will now be in the *c:\sql_metrics_removed* directory.

   This command is a long-running operation and can run for a few hours.

CA Technologies Proprietary Documentation. Authorized Use Only.

5. If necessary, repeat the *remove_metrics* command on a different set of excess metrics.

   Sockets are another common source of leaked metrics. After running the *remove_metrics* command, use this command to remove socket metrics:

   *<EM_Home>\tools>SmartStorTools.bat remove_metrics –dest c:\sql_and_sockets__metrics_removed -metrics ".*Sockets.*" -src* c:\sql_metrics_removed

   This command makes use of the data derived from the first command that removed SQL metrics as the source data (the data in the *c:\sql_metrics_removed directory*). The data that has both SQL and socket metrics removed will then be in this location specified by the *–dest* option: *c:\sql_and_sockets__metrics_removed*. Again, this command needs almost as much extra space to run as the original data.

6. Check the metrics data value for the cleaned metrics data and metadata.

   a. Open the *IntroscopeEnterpriseManager.properties* file in the *<EM_Home>\config* directory.

   b. Make a note of the current property setting, then set the property *introscope.enterprisemanager.metrics.historical.limit=500000*.

   c. Make a note of the current setting, then set the property *introscope.enterprisemanager.smartstor.directory*= to point to the destination location for the metrics removal.

      For example, the example destination location is *c:\sql_and_sockets__metrics_removed*. Therefore, this property setting would be *introscope.enterprisemanager.smartstor.directory=c:\sql_and_sockets__metrics_removed*

   d. Restart the Enterprise Manager.

   e. Check the value of this metric:
      Enterprise Manager | Data Store | SmartStor | Metadata | Metrics with Data

      The value should be far below the original value. If it is not, repeat the trimming steps using suitable SmartStor tools commands against agent or metric names that are not needed. Once the historical metric count is below 300K, the Enterprise Manager will run more efficiently.

CA Technologies Proprietary Documentation
Unauthorized Use Prohibited

7. If needed, reset the changed *IntroscopeEnterpriseManager.properties* values back to their original settings.
   a. Shut down the Enterprise Manager.
   b. Open the *IntroscopeEnterpriseManager.properties* file in the *<EM_Home>\config* directory.
   c. Reset these properties back to the original settings:
      - *introscope.enterprisemanager.metrics.historical.limit=*
      - *introscope.enterprisemanager.metrics.historical.limitintroscope.enterprisemanager.smartstor.directory=*
   d. Restart the Enterprise Manager.

Example: Using the test_regex command to find a SmartStor data problem

In this example, you use the *text_regex* command to discover that JMX metrics data have exploded.

1. Run a *test_regex* command to produce a readable version of metrics.metadata.

   For example,

   <EM_Home>\tools>SmartStorTools.bat test_regex -metrics -src <EM_Home>\data

   The result is a list of the metadata metrics.

2. Review the metadata metrics list to determine the probable source of the metrics leak or explosion.

   For example, you may find a lot of metrics metadata entries for JMX metrics.

3. (Optional) Run a script to sort and group the metrics.

   For example, to verify that JMX metrics are the source of the metrics explosion you can run this script:

   grep '|'. | awk -F'|' '{print $1}' | sort | uniq -c*

   If you are using a UNIX system, this would be the full *test_regex* command with sort script:

   SmartStorTools.sh test_regex -metrics -src ../data | grep '|'. | awk -F'|' '{print $1}' | sort | uniq -c*

CA Technologies Pre-release Documentation. Authorized Use Only.

4. Review the sort results.

   For example, you might see a list of metric counts and types that looks like this:

   1054327 JMX

   1564 JSP

   1051 OlamWebApp

   889 EJB

   569 CEF

   496 Struts

   173 Edocs

   167 Servlets

   36 CPU

   29 Sockets

   In this case, SmartStor is storing over a million metrics, and most are JMX metrics. This indicates there's been a JMX metrics explosion that can be cleaned up using SmartStor tools. For more information on metrics explosions, see the *CA APM Sizing and Performance Guide*.

CA Technologies Prerelease Documentation. Authorized Use Only

Chapter 5: Using Virtual Agents to Aggregate Metrics

This section has information about configuring and using Virtual agents.

This section contains the following topics:

Understanding Virtual Agents (see page 97)
Virtual Agent requirements (see page 97)
Configuring Virtual Agents (see page 98)

Understanding Virtual Agents

You can configure multiple physical agents into a single *Virtual agent*. A Virtual agent enables an aggregated, logical view of the metrics reported by multiple Java Agents.

A Virtual agent is useful if you manage clustered applications with Introscope—a Virtual agent is comprised of the agents that monitor different instances of the same clustered application, and appears in the Introscope Workstation as a single agent. This allows metrics from multiple instances of a clustered application to be presented at a logical, application level, as opposed to separately for each application instance.

You can view performance and availability data for a specific application instance, by scoping your views and interactions in terms of a single agent.

Virtual Agent requirements

If you have multiple stand-alone Enterprise Managers, a Virtual agent can contain only agents that report to the same Enterprise Manager.

In an Enterprise Manager cluster, agents that report to Enterprise Managers within a single cluster can belong to the same Virtual agent, regardless of the Collector Enterprise Manager to which they report.

Consider these conditions when configuring Virtual agents:

- An agent can be assigned to multiple Virtual agents.
- Virtual agents cannot include other Virtual agents.
- If you define multiple Virtual agents, they must have unique names, including custom metric agents—all agents in a cluster must have unique names.

CA Technologies Proprietary Documentation. Authorized Use Only

Configuring Virtual Agents

You configure virtual agents using the *agentclusters.xml* file, located in the *<EM_Home>/config* directory of the Enterprise Manager to which the agents report. If you run clustered Enterprise Managers, configure virtual agents using the *agentclusters.xml* file in the MOM *<EM_Home>/config* directory. A sample agentclusters.xml is available in the *<EM_Home>/*config directory.

This sample *agentclusters.xml* defines a virtual agent named *BuyNowAppCluster*, in the Introscope SuperDomain. The virtual agent includes all agents, on any host, whose agent name starts with *BuyNow*.

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<agent-clusters xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="agentclusters0.1.xsd" version="0.1" >
   <agent-cluster name="BuyNowAppCluster" domain="SuperDomain" >
      <agent-specifier>.*\|.*\|BuyNow.*</agent-specifier>
      <metric-specifier>Frontends\|.*</metric-specifier>
   </agent-cluster>
</agent-clusters>
```

The root element, *<agent-clusters>*, is required. The *<agent-cluster>* element defines a virtual agent, and has two required attributes:

- *name*—If you define multiple virtual agents, each must have a unique name.
- *domain*—Assigns the virtual agent to an Introscope domain.

  If no domain is defined (as *domain=""*) in the agent-cluster definition, the virtual agent defaults to the SuperDomain.

If you define multiple virtual agents, you define an *<agent-cluster>* element for each. The *<agent-cluster>* element requires two child elements:

- *<agent-specifier>*—Contains a regular expression that specifies the agents in the virtual agent, using the standard fully qualified agent name:

  <host> | *<process>* | *<agentName>*

- *<metric-specifier>*—Contains a prefix that specifies the metrics to collect from the agents in the virtual agent, in terms of resource type, or subsets of the instances of a resource type. Recommended prefixes are:
  - Frontends
  - CPU
  - JMX
  - WebSpherePMI

Note: While these are recommended prefixes, any resource can be used as a metric specifier.

CA Technologies Proprietary Document. Authorized use only.

The *<agent-cluster>* element can contain multiple *<metric-specifier>* stanzas.

Note: A higher volume of matching metrics imposes high overhead on the Enterprise Manager and can ultimately affect Enterprise Manager capacity.

Important! Regular expressions and wildcard metric specifiers such as ".*" and "(.*)" are allowed, but used with caution, if at all. Use of wildcards can result in a high volume of metrics and a performance impact

To display a virtual agent under a specific Browse tree domain node, define the *tgtDomain* attribute.

- *tgtDomain*—Defines a target domain where virtual agent metrics defined in the domain attribute are copied for display in the Browse tree.
  - If *tgtDomain* is not defined, then the domain is named tgtDomain.
  - If the *tgtDomain* value is not a valid domain defined in domains.xml, Introscope cannot create the target domain Browse tree node.
  - To view the domain defined using the *tgtDomain* attribute, users must have permission defined in domains.xml. For information about domains.xml permissions, see the *CA APM Security Guide.*

This sample *agentclusters.xml* defines a virtual agent named *BuyNowAppCluster*, which is defined in the Introscope SuperDomain node and is displayed in the MyDomain node.

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<agent-clusters xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="agentclusters0.1.xsd" version="0.1" >
   <agent-cluster name="BuyNowAppCluster" domain="SuperDomain"
tgtDomain="MyDomain"
      <agent-specifier>.*\|.*\|BuyNow.*</agent-specifier>
      <metric-specifier>Frontends\|.*</metric-specifier>
   </agent-cluster>
</agent-clusters>
```

In this case, CA APM users with permission to view only the MyDomain domain see this Browse tree:

```
Custom Metric Host(Virtual)
    Customer Metric Process(Virtual)
       BuyNowAppCluster
```

CA Technologies Internal / Partner Use Only. Not for External Distribution.

Chapter 6: Working with Enterprise Manager Clustering

This section describes Introscope's Enterprise Manager clustering functionality and provides instructions for configuring and starting an Enterprise Manager cluster.

This section contains the following topics:

Understanding Enterprise Manager clustering (see page 101)
Using Enterprise Manager hot failover (see page 104)
Configuring a cluster of Enterprise Managers (see page 109)

Understanding Enterprise Manager clustering

In a clustered environment, an Enterprise Manager serves as a *Manager of Managers* (MOM), managing the other Enterprise Managers in the cluster. The managed Enterprise Managers are named *Collectors*.

Workstation users logged in to a MOM can

- See and access the data reported by the agents that report to each of the Collectors that the MOM manages.
- Perform all usual Workstation functions on those Collectors.

This section describes Enterprise Manager clustering. For instructions on configuring a cluster, see Configuring a cluster of Enterprise Managers (see page 109). For information about clustering and agent load balancing, see Configure agent load balancing.

When to cluster Enterprise Managers

Clustering is useful in large Introscope environments that consist of multiple Enterprise Managers and many agents—a cluster configuration reduces the complexity and cost of managing a multiple-Enterprise Manager environment.

CA Technologies Pre-release Documentation
Authorized Use Only.

Understanding how MOMs manage Collectors

When you start a MOM, configured as described in Configuring a cluster of Enterprise Managers (see page 109), it attempts to establish a connection with each of the Collectors it manages.

Every agent in the Investigator automatically has a metric called Enterprise Manager host. This metric is helpful when running a cluster, so you can tell to which Collector Enterprise Manager the agent is actually connected.

The MOM uses a built-in, pre-configured user account for connecting to Collectors. The password is protected using public/private key encryption. For more information, see the *CA Wily APM Security Guide*.

If the MOM's initial attempt to connect to one of its Collectors results in an unrecoverable error, such as an authentication error, an unknown host exception, or a *java.rmi.ConnectIOException*, then that Collector is removed from the cluster—the MOM stops trying to connect to that Collector until the next restart of the MOM.

For recoverable errors, such as a lost connection, or a *java.rmi.ConnectionException*, the MOM repeats the connection attempt every 10 seconds.

Logging of connection status occurs in this fashion:

- If the logging level is DEBUG, a stack trace of the connection failure is written to the log.
- If the logging level is VERBOSE, a message is written to the log for the first 10 consecutive failed attempts; after 10 consecutive failed attempts, a message is written to the log that no further logging of failed connections attempts will be made.
- A message is always written to the log when the MOM connects or re-connects to a Collector.

As long as it remains running, the MOM polls the Collectors it manages.

Clustering and Enterprise Manager clock synchronization

In an Enterprise Manager cluster, Collector system clocks must be within three seconds of the MOM clock setting. If at startup, a Collector's clock is not synchronized with the MOM's clock, the Collector will fail to join the cluster. During operation, if a previously synchronized Collector's clock become more than three seconds off, the MOM will no longer accept data from that Collector.

Clock synchronization among the MOM and the Collector Enterprise Managers is not automatic.

CA Technologies Prerelease Documentation. Authorized Use Only.

To keep the MOM and Collector Enterprise Manager clocks synchronized, CA Technologies recommends that all Enterprise Managers synchronize their clocks at regular intervals to an NTP server. If a Collector clock becomes more than three seconds off from the MOM's clock, error messages in the MOM's log indicate that the MOM is no longer accepting data from the Collector. If this occurs, increase the frequency of synchronization with the NTP server. No restart of Enterprise Managers is necessary after clock synchronization. The MOM will automatically re-connect to the Collector when it detects the clocks are synchronized.

Note: It is necessary to run a time server software which synchronizes the clocks of all the machines in the cluster at regular intervals.

Clustering and agent naming

In an Enterprise Manager cluster, the names of the agents reporting to the Collector Enterprise Managers must be unique.

The MOM requires unique agent names across the cluster. Although each Enterprise Manager ensures that the agents that connect to it have a unique name, as described in the discussion about resolving agent naming conflict in the *CA APM Java Agent Implementation Guide* or the *CA APM .NET Agent Implementation Guide*, you must ensure that the agents across the Collectors in the cluster have unique names. If more than one agent in the cluster have the same name, only the last encountered by the MOM appears in Workstation views. Configuring agents to obtain their names from the application server is a technique for assigning unique agent names. For more information, see the discussion about obtaining an agent name from the application server in the *CA APM Java Agent Implementation Guide* or the *CA APM .NET Agent Implementation Guide*, as appropriate to your environment.

A Collector automatically renames custom metric agents, adding a unique identifier to the agent name to identify the agent's Collector. The default identifier has the form (*host@port*), where *host* is the hostname of the Collector and *port* is the Collector's connection port, as defined in its property introscope.enterprisemanager.port.<ChannelName>. Alternatively, the identifier can be set using the *introscope.enterprisemanager.clustering.collector.identifier* property. The identifier should be unique.

For example, the name of a custom metric agents in the Super Domain in a Collector on *myhost* that listens for connections on port 5001 would be:

```
Custom Metric Agent (myhost@5001) (*SuperDomain*)
```

Clustering and Management Modules

In an Enterprise Manager cluster, only the Management Modules installed on the MOM are used. For performance reasons, it is recommended that no Management Modules reside on the Collector Enterprise Managers.

CA Technologies Product Documentation. Author Copy

Deploying PowerPack extension files in a cluster environment

If you have activated extensions in your CA APM deployment, extension files must be administered in a certain way when deployed in a clustered environment.

Follow these steps:

1. Locate the extension folder where it has been installed in *<EM_Home>/examples/<Extension_Home>*, where *<Extension_Home>* is the name of a extension.
2. Copy the entire contents of this folder to the *<EM_Home>/* location on the MOM.
3. Copy the entire contents of this folder, except the directory *<EM_Home>/examples/<Extension_Home>/config*, to the *<EM_Home>/* location on each Collector Enterprise Manager.

   The *<EM_Home>/examples/<Extension_Home>/config* directory is not copied because Management Module *.jar* files must only reside on the MOM.

Using Enterprise Manager hot failover

Enterprise Manager hot failover works by designating one Enterprise Manager as a primary and the other as a secondary or backup, or both Enterprise Managers as primaries. The responsibility of the secondary is to take control if the primary fails and to relinquish control when the primary recovers. The responsibility of the primary is to take control from the secondary.

You can configure Enterprise Manager hot failover if a pair of Enterprise Managers shares a single Introscope installation. Typically, the two Enterprise Managers are on different machines. The Introscope installation must reside on a network file system since both Enterprise Managers need read access to the Introscope installation, although only one Enterprise Manager at a time needs write access. For example, the Introscope installation might be shared using a Network Attached Storage (NAS) protocol such as Network File System (NFS) or Server Message Block (SMB).

In order to provide cluster stability, CA Technologies recommends that you configure hot failover for your MOM Enterprise Manager. In addition, you can configure hot failover for one or more Collector Enterprise Managers within a cluster or for a standalone Enterprise Manager. If you are configuring hot failover for the MOM and Collectors in a cluster, you configure each Enterprise Manager individually.

By default, Enterprise Manager hot failover is not enabled.

CA Technologies Pre-Released Documentation. Authorized Users Only.

To enable Enterprise Manager hot failover:

1. Navigate to the *<EM_Home>\config* directory and open the *IntroscopeEnterpriseManager.properties* file.
2. Locate the Hot Failover section.
3. Set the hot failover properties. See Enterprise Manager hot failover (see page 431).

   When agents, TIMs, and Workstations try to connect to an Enterprise Manager, they try all the IP addresses for a host name. Thus, if you have defined a logical host in DNS with the IP addresses of the primary and secondary Enterprise Managers, then the agents, TIMs, and Workstations can use this for the Enterprise Manager host name and connect to whichever Enterprise Manager is running.

   Note: In a failover environment, you can use the *EMCtrl* script to start a primary Enterprise Manager. The secondary Enterprise Manager should not use the control script. For example, on a UNIX platform, start the primary using the control script *EMCtrl.sh* start and use the *'./Introscope Enterprise Manager'* command to start the secondary.

Enterprise Manager hot failover rules

There are a series of rules that apply to hot failover:

- When you start the Enterprise Manager with hot failover, you must start each Enterprise Manager one at a time. In case of a primary/secondary configuration, the primary must be started first.
- If both Enterprise Managers are primary, then the second primary waits for the first primary to fail and does not relinquish control when the first primary recovers.
- For hot failover, the primary, secondary, and installation directories are typically on different hosts. However, the Introscope installation (where Introscope is originally installed) and Enterprise Managers may reside on the same host. For more information, see Configuring Enterprise Manager failover to work on a single host (see page 107).
- If the Enterprise Manager's host matches one of the primary hosts, then the Enterprise Manager assumes the role of primary.
- If the Enterprise Manager's host matches one of the secondary hosts, then the Enterprise Manager assumes the role of secondary.
- If the Enterprise Manager's host matches neither the primary nor secondary host, then the Enterprise Manager does a normal startup. However a warning message is logged.

CA Technologies Pre-release Documentation
Authorized Use Only.

- A secondary Enterprise Manager checks every two minutes (configurable) to see if a primary Enterprise Manager is waiting to start. If so, the secondary yields to the primary and shuts down.
- A primary Enterprise Manager does not yield to a secondary Enterprise Manager or another primary Enterprise Manager.
- The startup and shutdown sequence is as follows: When starting the hot failover pair, start the primary Enterprise Manager first then the secondary, and when shutting down the hot failover pair, stop the secondary first then the primary.

Restarting the secondary Enterprise Manager after yielding to the primary

Once the secondary Enterprise Manager has yielded to the primary, it exits the JVM. It is up to you to restart the secondary Enterprise Manger. You can perform the restart using a script or by running the secondary Enterprise Manager as a Windows service. The secondary exits with a status code of 23 (*org.eclipse.equinox.app.IApplication.EXIT_RESTART*) if it shuts down because it needs to yield to a primary. The secondary exits with a 0 status code if it shuts down normally because of a command-line interrupt or by the Workstation.

If you are using Windows, you can use a command like this:

```
:RESTART
"Introscope Enterprise Manager"
IF ERRORLEVEL 23 GOTO RESTART
```

The *introscope.enterprisemanager.failover.interval* property controls how often a running secondary Enterprise Manager checks to see if a primary Enterprise Manager is waiting to start. For more information, see Enterprise Manager hot failover (see page 431).

CA Technologies Pre-release Documentation
Authorized Use Only

Configuring Enterprise Manager failover to work on a single host

You can configure Enterprise Manager failover to work on a single host rather than a shared file system. You set up Enterprise Manager failover on a single host to permit a second Enterprise Manager to take over if the first crashes when you do not want to:

- run an Enterprise Manager as a Windows service
- restart an Enterprise Manager using a shell script.

To configure Enterprise Manager failover for a single host:

1. Set the following Enterprise Manager properties in the *IntroscopeEnterpriseManager.properties* file as follows:
   - *introscope.enterprisemanager.failover.enable=true*
   - *introscope.enterprisemanager.failover.primary=localhost*
2. Start the first primary Enterprise Manager by clicking Introscope Enterprise Manager.exe.
3. Start the second primary Enterprise Manager by clicking Introscope Enterprise Manager.exe.

   When the second primary Enterprise Manager starts properly, in the command prompt (the console) a message displays stating: *acquiring primary lock*. If this message does not appear, then either Enterprise hot failover is not configured correctly or the primary Enterprise Manager is not running.

About MOM failover and CA CEM TIMs

If you are deploying CA CEM in your CA APM environment, you configure the MOM to connect with each TIM. When you enable a TIM, the MOM gives the TIM two IP addresses: the MOM IP address and the IP address of the Collector running the TIM Collection Service. This Collector connects with the TIM to collect the TIM data.

If you have configured MOM failover, when a TIM is running and enabled, the primary MOM provides the TIM with:

- its primary MOM IP address
- if set, the secondary MOM IP addresses that you defined in the *IntroscopeEnterpriseManager.properties* file.

The TIM stores these IP addresses, which you can view in the Wily System Setup > TIM > TIM Settings page.

CA Technologies Proprietary Use Documentation. Unauthorized Use Only.

In the case of MOM failover, in order for the failover MOM to automatically connect to the TIM, you must make sure that these conditions are met:

- the TIM is running and enabled
- in the *IntroscopeEnterpriseManager.properties* file,
  - the *IntroscopeEnterpriseManager.failover.enable* property is set to true
  - one or more primary MOM IP addresses are defined in the *introscope.enterprisemanager.failover.primary* property

  Note: If you define more than one primary MOM IP address, it is not necessary to define a secondary MOM IP address.
- If you define only one primary MOM IP address, you must define one or more secondary MOM IP addresses in the *introscope.enterprisemanager.failover.secondary* property in the *IntroscopeEnterpriseManager.properties* file.

  Note: You can define more than one secondary MOM if desired.
- If the address of the failover MOM changes, then you must disable and enable all the TIMs so that they can obtain the IP address of the new failover MOM.
  You can see the failover MOM IP addresses in the CEM console in the TIM Settings page in the *TessIpAddr* field.

If these conditions are met, upon primary MOM failure and failover to either another other primary or the secondary MOM, the failover MOM automatically connects to the TIM.

If these conditions are not met, here's what happens upon a MOM failure:

If...	**Then upon MOM failure...**
TIM is not running and enabled	No connection to the TIM before or after MOM failure. The TIM must be running and enabled to be recognized by the MOM and receive the MOM connection.
IntroscopeEnterpriseManager.failover.enable property is set to *false*	The Collector continues to pull data from the TIM, but no alternative primary or secondary MOM tries to connect to the TIM.
No MOM additional primary or secondary IP addresses defined	Collector continues to pull data from the TIM, but no alternative primary or secondary MOM tries to connect to the TIM. Any MOM configuration changes are not sent to the TIM.

CA Technologies Pre-release Documentation. Authorized Use Only.

Configuring a cluster of Enterprise Managers

You configure a cluster of Enterprise Managers by defining cluster-related properties in the MOM's properties file.

Important! CA Technologies strongly recommends that you do not have differing domains.xml files for the MOM and the Collectors. The MOM can handle differing domains within the cluster for live agents (agents presently sending data to the cluster). However, having differing domains for historical agents on the MOM and Collectors can cause inconsistency in the MOM view of historical data. In such mixed-domain clusters, historical agents on Collectors are not tracked, so their data does not display in Workstation graphs made through the MOM unless you explicitly mount the historical agents. To avoid this situation, CA Technologies strongly recommends that you place the identical domains.xml file on the MOM and all the Collectors in a cluster. This allows live and historical agent data to be visible at all times from the MOM Workstation without mounting historical agents to view historical data on a Workstation associated with a specific Collector. For information about working with domains.xml, see the *CA APM Sizing and Performance Guide*.

Configuring the MOM

To configure the MOM for a set of Collector Enterprise Managers:

1. Set *introscope.enterprisemanager.clustering.manager.enable* in the MOM's properties file to *true.*

   Define the connection properties for each of the Collector Enterprise Managers that the MOM will manage. For each Collector that the MOM will manage, define these properties in the MOM's properties file where <*emName*> in each of the properties is the same arbitrary identifier of the Collector Enterprise Manager:

   - introscope.enterprisemanager.clustering.login.<emName>.host—specify the Collector hostname or IP address.
   - introscope.enterprisemanager.clustering.login.<emName>.port—specify the listener port on the Collector.
   - introscope.enterprisemanager.clustering.login.<emName>. publickey—uncomment this property to set it to the default value of

   config/internal/server/EM.public

2. Restart the MOM.

CA Technologies Prerelease Documentation Authorized Use Only.

Configuring a Collector

To configure a Collector in a cluster of Enterprise Managers:

1. Set *introscope.enterprisemanager.clustering.collector.enable* Collector's properties file to *true*. (For more information, see introscope.enterprisemanager.clustering.collector.enable (see page 425).)
2. If you want this Collector to have a special identifier, you can set the optional property introscope.enterprisemanager.clustering.collector.identifier in the Collector's properties file.

   If you don't specify a special identifier for the Collector, the default of *host@port* is used, where *host* is the Enterprise Manager's host name, and *port* is the Enterprise Manager's connection port.

Important! Enterprise Manager MOMs and Collectors must use the same *domains.xml* file, or you may lose some historical data when querying agent metric data from the Workstation.

CA Technologies Prerelease Documentation
Authorized Use Only.

Chapter 7: Configure MOM Agent Load Balancing

CA APM clustered environments support agent load balancing. To use this capability, you configure agents to connect to the MOM instead of a Collector. The MOM provides agents with Collectors available for connection to based on the metric load and your CA Introscope configuration. The agents connect to an available Collector.

The MOM equalizes the metric count among the Collectors by ejecting participating agents from over-burdened Collectors. The ejected agents reconnect to the MOM, and then are reallocated to under-burdened Collectors. Load balancing is only for agent metric load; it is not applicable to the CA CEM data load.

Pre-8.0 agents cannot connect to the MOM; they must instead directly connect to one of the Collectors. If a pre-8.0 agent tries to connect to the MOM

- MOM denies the connection.
- CA Introscope writes a warning message to both the MOM log and the agent log.

Note: If your cluster is handling a full load that includes agents and TIMs, use weighting factors on Collectors running the Enterprise Manager services. Enterprise Manager services can require increased resources, so using weighting factors can compensate for possible reduced metric capacity. For more information about weighting factors, see Set the agent load balance metric weighting factor (see page 113).

This section contains the following topics:

How the MOM Assigns Agents to Collectors (see page 112)
How the MOM Rebalances the Cluster (see page 112)
Set the Agent Load Balance Metric Weighting Factor (see page 113)
Set the Agent Load Balance Metric Threshold (see page 114)
Set the Agent Load Balance Interval Property (see page 115)
Agent Load Balancing Usage Scenarios (see page 115)
Configure the Assignment of Agents to Collectors (see page 118)

CA Technologies Pre-release Documentation
Distribution Private Only

How the MOM Assigns Agents to Collectors

Several factors determine how the MOM assigns an agent to a Collector.

- Connection type.

  An agent is only assigned to a Collector that supports the same connection type that the agent uses to connect to the MOM. For example, if the agent connects to the MOM using HTTP, then the Collector must have enabled HTTP connections. For information about setting connection types, see the *CA APM Java Agent Implementation Guide* or *CA APM .NET Agent Implementation Guide*.

- Configuration done in the *loadbalancing.xml* file.

  You configure the *loadbalancing.xml* file to use these statements:

  - <include> statements to allow agents to connect to a specific set of Collectors.
  - <exclude> statements to disallow agents from connecting to specific set of Collectors.

- Agent connection history with a specific Collector.

  CA Introscope includes the capability to prevent an explosion of SmartStor data as agents are transferred from one Collector to another in the cluster. If an agent has connected to a Collector previously, the MOM favors that Collector for future connections unless these factors are in effect:

  - An alternative under loaded Collector is available.
  - A favored Collector is overloaded.

- Agent connection when no history (upon startup or restart)

  - If an agent has no history with any Collectors, the MOM assigns the agent to an available under-loaded Collector.

How the MOM Rebalances the Cluster

The MOM can disconnect agents from a Collector and redirect them to other Collectors when it detects that the cluster metrics load has become unbalanced. This capability is named the *MOM load balancer;* the MOM performs *agent load balancing*. A cluster is unbalanced if any Collector metrics load is either under loaded or overloaded.

A cluster can become unbalanced for a number of reasons, as shown in these examples:

- A Collector can crash and be restarted.
- A new Collector is added to the cluster.
- Many agents connect directly to one Collector.

CA Technologies Product Documentation. Authorized Use Only.

The MOM checks for cluster unbalance at regular intervals. If the cluster is unbalanced, the MOM seeks to restore the balance by disconnecting agents from Collectors that are over a weight-adjusted average. Disconnected agents reconnect to the MOM and are redirected to Collectors that are under the weight-adjusted average. The goal is an assignment of agents that equalizes the weight-adjusted metric load across the cluster.

When rebalancing, the MOM performs these actions on agents that are connected directly to Collectors:

- MOM disconnects pre-9.1 agents. The agents use agent profile file configuration information to find a Collector for reconnection.
- MOM redirects 9.1 agents to an allowed available under-loaded Collector.

Set the Agent Load Balance Metric Weighting Factor

When you configure agent load balancing, the MOM allots agents to Collectors based on weight-adjusted load. You can adjust the weighting factor of individual Collectors in a cluster. The number of metrics that a Collector can handle determines the relative power of a Collector in a cluster. A number of factors affect the number of metrics that a Collector can handle including these factors:

- CPU power.
- Memory.
- Any assigned Enterprise Manager services handling the CA CEM data load.

When setting up agent load balancing, you can configure the weighting factor to require the MOM to assign fewer metrics to Collectors with less capacity. You can help avoid cluster performance problems by setting up the metric weight load so that the more powerful Collectors handle the bigger metric loads.

The value of the introscope.enterprisemanager.clustering.login.<emID>.weight property is a positive number that controls the relative load of the Collector. The weight load on a specific Collector divided by the total weight load of the cluster is the percentage of the metric load assigned to that Collector. Factors determining how the MOM assigns agents to Collectors (see page 112) modify the weight load.

Note: In the introscope.enterprisemanager.clustering.login.<emID>.weight property name, *<emID>* is an arbitrary identifier. Each Collector has a unique identifier. Provide an appropriate identifier for your environment.

The MOM then uses weight-adjusted metric counts when assigning agents to Collectors and when rebalancing the agent metric load. For example, a MOM connects to three Collectors that all have zero metrics currently being reported.

CA Technologies Prerelease Documentation. Authorized Use Only.

Example

Three Collectors have these weight settings:

- Collector A has a weight of 150
- Collector B has a weight of 100
- Collector C has a weight of 50

The MOM assigns metrics to Collectors A, B, and C approximately in the ratio of 3:2:1.

Follow these steps:

1. Open the IntroscopeEnterpriseManager.properties file located in the <*EM_Home*>/config directory.
2. Set this property to assign a weighting load factor to the Collector.
   `introscope.enterprisemanager.clustering.login.<emID>.weight`
3. Save and close IntroscopeEnterpriseManager.properties.

Set the Agent Load Balance Metric Threshold

The introscope.enterprisemanager.loadbalancing.threshold property sets the cluster tolerance, or threshold, for metric imbalance.

The default introscope.enterprisemanager.loadbalancing.threshold property setting is 20,000 metrics, meaning a Collector has to be 20,000 metrics out of balance before the MOM rebalances the agents. A Collector is out of balance when it is either under or over the weight-adjusted cluster average by more than the threshold. See the agent load balancing usage scenarios (see page 115) for some examples.

Configure a load balancing threshold value that avoids frequent rebalancing. Agent metric load tends to be fairly stable. Set the threshold to be large enough to avoid triggering rebalancing when transient metrics are generated for problem investigation.

Follow these steps:

1. Open the IntroscopeEnterpriseManager.properties file located in the <*EM_Home*>/config directory.
2. Set this property to configure the cluster tolerance for metric imbalance.
   `introscope.enterprisemanager.loadbalancing.threshold`
3. Save and close IntroscopeEnterpriseManager.properties.

CA Technologies Prerelease Documentation. Authorized Use Only.

Set the Agent Load Balance Interval Property

The introscope.enterprisemanager.loadbalancing.interval property tells the MOM how often to check the cluster for possible rebalancing. The default interval is 600 seconds (10 minutes) and the minimum interval is 120 seconds (2 minutes).

The cluster requires a readjustment period of several minutes following any agent reassignment. If the load balancing interval is too short, there is a danger that rebalancing starts again before the cluster has adjusted to the previous occurrence. The risk a new rebalancing starting before the current rebalancing is finished is greater if you use the application triage map. In this case, when agents are reassigned, they must resend application triage map data to the new Collector. If you have performed these actions, watch for signs of instability that can result from too-frequent rebalancing:

- Deployed CA Introscope 8.x load balancing.
- Set the load balancing interval below the default.

Follow these steps:

1. Open the IntroscopeEnterpriseManager.properties file located in the <*EM_Home*>/config directory.
2. Set this property to configure how often the MOM checks the cluster for load balancing.
   `introscope.enterprisemanager.loadbalancing.interval`
3. Save and close IntroscopeEnterpriseManager.properties.

Agent Load Balancing Usage Scenarios

Here are some examples to help you understand how agent load balancing operates in a clustered environment.

Agent Load Balancing When a New Collector is Added to the Cluster

You determine that your cluster of three Collectors A, B, and C is overloaded.

You dynamically add Collector D by adding a set of connection properties to the IntroscopeEnterpriseManager.properties file on the MOM. In addition, you have set the metric threshold for all the Collectors in this cluster to 10,000 metrics. You did so by configuring the introscope.enterprisemanager.loadbalancing.threshold property in the IntroscopeEnterpriseManager.properties file.

CA Technologies Prerelease Documentation. Authorized Use Only.

The average metric load is 22,500 metrics reporting each to Collectors A, B, and C, while Collector D has zero metrics being reported to it. The MOM rebalances the cluster because:

- The difference between the average load of 22,500 metrics and zero metrics

is greater than

- The threshold of 10,000 metrics.

The MOM rebalances 7,000 metrics each from Collectors A, B, and C and redirects all 21,000 metrics to report to Collector D. After rebalancing, 23,000 metrics report to Collector A, 23,000 metrics to Collector B, 23,000 metrics to Collector C, and 21,000 metrics to Collector D.

Agent Load Balancing When a Collector with Agents is Added

A MOM connects to Collectors A and B. 36,000 metrics report to Collector A and 30,000 metrics to Collector B. You set the metric threshold to 10,000 metrics using the introscope.enterprisemanager.loadbalancing.threshold property in the IntroscopeEnterpriseManager.properties file. Collector C, which has 24,000 metrics reporting to it, is added to the cluster.

The MOM does not rebalance the metric load because:

- The average is 30,000 metrics per Collector. (36,000 + 30,000 + 24,000)=90,000 metrics/3 Collectors).
- None of the Collectors differs from 30,000 by more than the 10,000 metric threshold.

Agent Load Balancing When a Collector Fails

A MOM connects to Collectors A, B, and C. 36,000 metrics report to Collector A, 30,000 metrics to Collector B, and 24,000 metrics to Collector C. Collector A fails and the agents that reported to Collector A reconnect to the MOM. The MOM redirects approximately 15,000 metrics to Collector B and 21,000 to Collector C. Now Collectors B and C both have 45,000 metrics being reported to them.

CA Technologies Product Documentation: Authorized Use Only

Agent Load Balancing When a Collector Restarts

When Collector A recovers after failure, the cluster is unbalanced. Ideally there would be an average of 30,000 metrics reporting to each Collector. However, Collectors B and C each have 45,000 metrics reporting to them, which:

- Is 15,000 metrics over the average.
- Exceeds the threshold of 10,000 metrics.

Therefore the MOM disconnects agents with a total of 15,000 metrics each from Collectors B and C and redirects all 30,000 metrics to Collector A. This results in 30,000 metrics being reported each to Collectors A, B, and C.

Agent Load Balancing with Weights

A MOM connects to Collectors A, B, and C. The threshold is set to 10,000. 24,000 metrics report to Collector A, 30,000 to Collector B, and 36,000 to Collector C, making a total of 90,000 metrics. The cluster has an average of 30,000 metrics per Collector. You have set the weight of Collector A to 150, Collector B to 100, and Collector C to 50. The average weight is 100 (the sum of the weights divided by the number of Collectors or (150 + 100 + 50)/3).

Ideally each Collector has a metric load proportional to its relative weight. Because Collector A has a weight of 150, it would ideally have 45,000 metrics. Its weight is 50 percent higher than average, so its metric load would ideally be 50 percent above the 30,000 metric average. Collector B has an average weight and therefore would ideally have the average metric load, or 30,000 metrics. Collector C has a weight 50 percent of the average and therefore would ideally have 50 percent of the average metric load, or 15,000 metrics.

Based on these relative weights and metric averages, the cluster is unbalanced. Collector A is under loaded because it is under the weight-adjusted average by more than the threshold (24,000 - 45,000= -21,000). Collector B is perfectly balanced because at 30,000 metrics, its metric load is equal to its weight-adjusted average. Collector C is overloaded because it is over the weight-adjusted average by more than the 10,000 metric threshold (36,000 - 15,000= 21,000). The MOM rebalances the cluster by ejecting agents with 21,000 metrics from Collector C and redirecting them to Collector A.

Load Balancing When Agents Prefer a Specific Collector

If you have an agent that prefers a particular Collector, configure the latched property in the loadbalancing.xml file (see page 118). This configuration allows specific agents to connect a specific Collector whenever that Collector is available.

If you do not configure the latched property, then the MOM can redirect the agent from an overloaded Collector to an under-loaded Collector.

CA Technologies Product Release Documentation. Authorized Use Only.

Configure the Assignment of Agents to Collectors

Use the MOM loadbalancing.xml to configure agent elements that control the assignment of connecting agents to Collectors by the MOM load balancer. You can configure include and exclude elements to allow and disallow agent connections to specific Enterprise Managers.

Follow these steps:

1. Open loadbalancing.xml in the *<EM_Home>/config* directory on the MOM.
2. Define the properties to meet the needs of your environment.

   Use the following guidelines and the agent load balancing configuration examples (see page 120) when defining these properties:

   - The MOM load balancer evaluates the *<agent-collector>* elements in sequential order.

     If a connecting agent matches any of the *<agent-specifiers>*, it is assigned to one of the Collectors in the *<include>* list. If there is no *<include>* list, it is assigned to a Collector that is not in the *<exclude>* list. If there are no Collectors available that meet these restrictions, the agent request for a Collector connection is denied.
   - For each *<agent-collector>* element, define one or more *<agent-specifier>* elements to match a set of agents.
   - Be sure that text content of *<agent-specifier>* is a Perl5 regular expression, such as .*\|.*\|EPAgent1</agent-specifier>.
   - To include all agents, do not use ".*" (dot star), because this code by itself is not used in agent load balancing semantics. Instead, use .*\|.*\|.*, which tells the MOM to look for all agents on all hosts and to include all processes and all agent names ([all hosts|all processes|all agents]).
   - For the *<agent-collector>* element, the include list can be empty. The include list can also contain Collectors to which the agents are permitted to connect.
     - If the include list is empty, agents are not allowed to connect to any Collectors.

     ```
     <include>
     </include>
     ```

     - Within <include> and </include> elements, you can add Collectors to which the agent is allowed to connect.

       In this example, the agent is allowed only on the Collector named localhost@5002.

     ```
     <include>
         <collector host="localhost" port="5002"/>
     </include>
     ```

CA Technologies Internal/Restricted Use Only. Not for distribution.

- For the *<agent-collector>* element, the exclude list can be empty. The exclude list can also contain Collectors to which the agents are not allowed to connect.
  - If the exclude list is empty, agents are allowed to connect to all Collectors.

```
<exclude>
</exclude>
```

  - Within <exclude> and </exclude> elements, you can add Collectors to which the agent is not allowed to connect.

    In this example, the agent is not allowed only on the Collector named localhost@5002.

```
<exclude>
   <collector host="localhost" port="5002"/>
</exclude>
```

- Set the latched property to allow specific agents to connect a specific Collector whenever that Collector is available. This Collector is named a *latched* Collector.

  If you add only one Collector to the include list, then no latched property setting is needed.

  If more than one Collector has latched=true, then an agent can connect to any of those latched Collectors.

  When the latched property is set and the Collector is overloaded, the MOM does not disconnect the agent from the Collector. If a latched Collector goes down, then the disconnected agent connects to another allowed available latched Collector. For example, if your environment has ten Collectors, you could configure agent to latch to two or three of the Collectors. If the latched Collectors all become unavailable, then the MOM can redirect the agents to the other seven unlatched available Collectors.

  When the latched property is not set on a Collector, the agent connects to that Collector only if all of the latched Collectors are unavailable.

  If the latched value is not set, then CA Introscope assumes the value to be false.

```
<include>
    <Collector host="newCollector-ip" port="6001" latched=”true” />
</include>
</agent-Collector>
```

- The *<name>* attribute is optional.

3. Save and close loadbalancing.xml.

   The new values are put into effect when the Enterprise Manager checks the property file during a hot deploy. Hot deploys occur about once every 60 seconds. The MOM load balances the cluster at the next load balance interval.

CA Technologies Internal Use Only. Authorized Release Documentation.

Agent Load Balancing Configuration Examples

You can configure the MOM agent load balancing properties in varied ways. These configuration examples provide code samples:

- Assign an agent to a particular Collector or set of Collectors (see page 121)
- Exclude an agent from a particular Collector or set of Collectors (see page 122)
- Assign multiple agents to multiple Collectors (see page 123)

CA Technologies Prerelease Documentation. Authorized Use Only.

Assign an Agent to a Particular Collector or Set of Collectors

The following example shows how to configure assigning an agent named MyAgent to the Collector named "myhost" on port 5001.

```
<agent-collector name="Assign MyAgent to One Collector Example">
    <agent-specifier>.*\|.*\|MyAgent</agent-specifier>
    <include>
        <collector host="myhost" port="5001"/>
    </include>
</agent-collector>
```

The following example shows how to configure the latched property where:

- All agents are named EPAgent.
- Agents are assigned to connect to the Collector named "newCollector-ip" on port 6001.

```
<agent-Collector name="Latch True Example">
    <agent-specifier>.*\|.*\|EPAgent.*</agent-specifier>
    <include>
        <Collector host="newCollector-ip" port="6001" latched=”true” />
    </include>
</agent-Collector>
```

The following example shows how to configure the latched property where:

- All agents are named EPAgent.
- Agents are assigned to connect to the Collector named "newCollector-ip" on port 6001.
- When newCollector-ip is unavailable, the MOM can assign all agents named EPAgent to either Collector1-ip or Collector2-ip.

```
<agent-Collector name="Latch when Collector Unavailable Example">
    <agent-specifier>.*\|.*\|EPAgent.*</agent-specifier>
    <include>
        <Collector host="newCollector-ip" port="6001" latched=”true” />
        <Collector host="Collector1-ip" port="6001" />
        <Collector host="Collector2-ip" port="6001" />
    </include>
</agent-Collector>
```

The following example shows how to assign All EP agents where:

- Collectors are named "localhost" on ports 5001 and 5002.
- Collector localhost@5001 is preferred over localhost@5002.

```
<agent-collector name="Allow EPAgents on Two Collectors and One is Preferred Example">
    <agent-specifier>.*\|.*\|EPAgent.*</agent-specifier>
    <include>
        <collector host="localhost" port="5001" latched="true"/>
        <collector host="localhost" port="5002"/>
    </include>
```

CA Technologies Prerelease Documentation. Authorized Use Only.

```
</agent-collector>
```

The following example shows how to assign All EP agents where:

- Collectors are named "localhost" on ports 5001, 5002, 5003, and 5004.
- Collectors localhost@5001 and localhost@5003 are preferred over localhost@5002 and localhost@5004.

```
<agent-collector name="Allow EPAgents on Four Collectors and Two are Preferred
Example">
    <agent-specifier>.*\|.*\|EPAgent.*</agent-specifier>
    <include>
        <collector host="localhost" port="5001" latched="true"/>
        <collector host="localhost" port="5002"/>
        <collector host="localhost" port="5003" latched="true"/>
        <collector host="localhost" port="5004"/>
    </include>
</agent-collector>
```

Exclude an Agent from a Particular Collector or Set of Collectors

The following example shows how to exclude an agent named MyAgent from a Collector named MyHost.

```
<agent-collector name="Exclude MyAgent from a Collector Example">
    <agent-specifier>.*\|.*\|MyAgent</agent-specifier>
    <exclude>
        <collector host="MyHost" port="5001"/>
    </exclude>
</agent-collector>
```

The following example shows how to exclude multiple agents from multiple Collectors.

```
<agent-collector name="Multiple Agents Assigned to Multiple Collectors Example>
    <agent-specifier>.*\|.*\|MyAgent.*</agent-specifier>
    <agent-specifier>.*\|.*\|MyOtherAgent.*</agent-specifier>
    <exclude>
        <collector host="MyHost" port="5001"/>
        <collector host="MyOtherHost" port="5001"/>
    </exclude>
</agent-collector>
```

CA Technologies Pre-release Documentation. Authorized Use Only.

Assign Multiple Agents to Multiple Collectors

The following example shows how to assign multiple agents to multiple Collectors.

```
<agent-collector name="Multiple Agents Assigned to Multiple Collectors Example>
    <agent-specifier>.*\|.*\|MyAgent.*</agent-specifier>
    <agent-specifier>.*\|.*\|MyOtherAgent.*</agent-specifier>
    <include>
        <collector host="MyHost" port="5001"/>
        <collector host="MyOtherHost" port="5001"/>
    </include>
</agent-collector>
```

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 8: Configure Agent - Enterprise Manager Network Topology

You can plan and configure which specific agents or groups of agents can connect to these computers:

- specific standalone Enterprise Managers.
- Collectors.
- groups of Collectors.

CA APM names the process as planning and configuring your CA APM agent - Enterprise Manager network topology (or structure).

The goal of planning and configuring your agent to Enterprise Manager network topology is to prevent or reduce CA APM performance issues. Such as preventing these example situations:

- New agents from randomly being started up and flooding a specific Collector with connections.
- Agents in a test environment from connecting to your production environment.

You can use these files and properties to control how and when agents connect to specific standalone Enterprise Managers or Collectors:

- loadbalancing.xml file.
- introscope.apm.agentcontrol.agent.allowed property in the EnterpriseManager.properties file.
- introscope.agent.enterprisemanager.failbackRetryIntervalInSeconds property in the IntroscopeAgent.profile file.

When you configure MOM agent load balancing in loadbalancing.xml (see page 118), you tell the MOM where to move agents based on these factors:

- Collector capacity.
- Number of agent metrics being produced.
- Collector capacity weighting factors.

You can also use loadbalancing.xml for fine-grained agent to Enterprise Manager connection setting, which the MOM then carries out across the cluster. This chapter describes how you can implement fine-grained agent to Enterprise Manager connections.

CA Technologies Prerelease Documentation. Authorized Use Only

Example

Your environment consists of a MOM and five Collectors. You want agent 1 and agent 2 to be load balanced on Collector 1 and Collector 2. You configure loadbalancing.xml for these fine-grained agent to Enterprise Manager connections:

- Allow agent 1 and agent 2 to connect to Collectors 1 and 2.
- Do not allow these agents to connect to Collectors 3, 4, and 5.

If there is an overload, the MOM handles the agents as follows:

- Allocates agent 1 and agent 2 on Collector 1 and Collector 2 depending on the runtime cluster load.
- Never allows these agents to connect to Collectors 3, 4, or 5.

The MOM sends to each Collector agent connection rules set in loadbalancing.xml upon startup and every 10 minutes. Collectors use these rules to know which agents to allow and disallow. Collectors then send the rules to each agent so each agent knows what to Collectors it can connect when these actions occur:

- MOM goes down.
- Agent is disallowed from a Collector.

If an agent is disallowed from an Enterprise Manager, then these actions occur:

- Agent is in a disallowed state.
- APM Status console Denied Agents list displays an entry.

When an agent restarts, it gets initial Collector connection information from the introscope.agent.enterprisemanager.connectionorder property in its IntroscopeAgent.profile file. The allowed Collectors are automatically supplied to this property in addition to the defined value. A list of allowed available Collectors or Standalone Enterprise Managers remains in the agent memory until agent restart. Upon restart, the list of allowed available Enterprise Managers is sent from the Collector to the agent.

CA Technologies Prerelease Documentation. Unauthorized use Only.

You can use the APM Status console Denied Agents list to monitor for repeated disallowed connections on specific agents or groups of agents. You can then configure loadbalancing.xml to allow the disallowed agents to connect to a different Collector.

This section contains the following topics:

What Happens When an Agent is Disconnected from an Enterprise Manager (see page 127)
What Happens When an Agent is in a Denied State (see page 128)
How Agent Connection Information Propagates Across CA Introscope (see page 128)
Process of Configuring Agent - Enterprise Manager Network Topology (see page 129)
Plan the Agent - Enterprise Manager Network Topology (see page 130)
Configure if New Agent Connections are Allowed When No loadbalancing.xml Rules (see page 132)
Configure Agent Reconnection Wait Time (see page 133)
Configure Disallowed Agent Connection Clamp Limit (see page 133)
Configure Agent Continual Reconnect to Any Allowed Enterprise Manager (see page 134)
Configure loadbalancing.xml for Allowed and Disallowed Agents by Enterprise Manager (see page 136)
Agent Connection Scenarios (see page 136)
Troubleshoot Agent - Enterprise Manager Network Topology (see page 137)

What Happens When an Agent is Disconnected from an Enterprise Manager

Agents are disconnected from Enterprise Managers when, for example, they are load balanced or the Enterprise Manager goes down.

Pre-9.1 agents are disconnected when either:

- Agents are disallowed from an Enterprise Manager.
- No Enterprise Manager is available for connection.

9.1 agents are disconnected when the agent is disallowed and the introscope.enterprisemanager.agent.disallowed.connection.limit clamp value is hit.

CA Technologies Pre-release Documentation. Authorized Use Only

What Happens When an Agent is in a Denied State

Based on the configurations set in loadbalancing.xml and introscope.apm.agentcontrol.agent.allowed property, agents are allowed or disallowed connection to specific Enterprise Managers.

When in a denied state, an agent:

- Disconnects from the Enterprise Manager.
- Tries to connect to other Enterprise Mangers to which it is allowed to connect.

The agent keeps trying Enterprise Managers until it can either connect or it runs out of Enterprise Managers:

- To which the agent is allowed to connect.
- Available for connection.

If an agent is denied on all Enterprise Managers, then how the agent acts differs based on its introscopeagent.profile file introscope.agent.enterprisemanager.failbackRetryIntervalInSecondsproperty setting. If the property is:

- Uncommented and a time duration set, then the agent keeps trying to connect to allowed Enterprise Managers periodically based on the duration.
- Not set (commented out), then based on the agent CA APM version the:
  - Pre-9.1 agent disconnects from the Enterprise Manager and keeps trying to reconnect.
  - 9.1 agent remains connected to the final allowed Enterprise Manager, and does not send any metric data to the Enterprise Manager. The agent is dormant as it waits for the Enterprise Manager to change the agent to an allowed state.

CA Technologies Prerelease Documentation. Authorized Use Only

How Agent Connection Information Propagates Across CA Introscope

Here is how CA Introscope sends agent to Enterprise Manager connection information across a cluster.

The MOM periodically sends loadbalancing.xml to all the Collectors, synchronizing loadbalancing.xml across the cluster. CA Introscope ignores any Collector or CDV loadbalancing.xml configurations.

If the MOM goes down after a Collector has received loadbalancing.xml, then the Collector uses its loadbalancing.xml file to allow or disallow agents. If a Collector has never received loadbalancing.xml from the MOM, it uses its own loadbalancing.xml file for agent allow and disallow instructions.

In addition, the Collector uses the introscope.apm.agentcontrol.agent.allowed property in its IntroscopeEnterprismanager.properties file to allow or disallow agent connections from agents not defined in loadbalancing.xml.

The MOM also provide agents with a list of Enterprise Managers to which the agents are allowed to connect and are available. This list is based on Enterprise Manager IP address, so can include Enterprise Managers in different clusters.

When the:

- MOM
- standalone Enterprise Manager
- Initial Collector

to which an agent is connected goes down, this list tells the agent to which Enterprise Manager to connect.

When:

- a standalone Enterprise Manager goes down
- Your environment contains both standalone Enterprise Managers and a cluster

loadbalancing.xml can direct agents to another standalone Enterprise Manager or a Collector.

Note: If an agent restarts, it uses its agent profile introscope.agent.enterprisemanager.connectionorder property for its initial Collector connection information.

CA Technologies Prerelease Documentation. Authorized Use Only.

Process of Configuring Agent - Enterprise Manager Network Topology

Configuring agent to Enterprise Manager network topology requires these high-level steps.

1. Plan your agent to Enterprise Manager or Collector topology.

2. Configure the introscopenenterprisemanager.properties file if appropriate for your deployment.
   a. (Optional) Configure whether agents not matched by any regular expressions in loadbalancing.xml are allowed or disallowed connection to the MOM or specific Enterprise Managers.

      Set the introscope.apm.agentcontrol.agent.allowed property.

   b. (Optional) Configure a randomized duration for agent reconnection attempts to the Enterprise Manager.

      Set the introscope.apm.agentcontrol.agent.reconnect.wait property

   c. (Optional) Configure the number of disallowed agents that can connect to an Enterprise Manager.

      Set the introscope.enterprisemanager.agent.disallowed.connection.limit.
3. Configure the IntroscopeAgent.profile file if appropriate for your deployment.
   - (Optional) Configure whether agents disconnected from an Enterprise Manager continually attempt to reconnect to any allowed Enterprise Manager.

     Set the introscope.agent.enterprisemanager.failbackRetryIntervalInSeconds property.
4. Configure loadbalancing.xml.
   a. (Optional) Configure the latch property so that specific agents always connect to specific Enterprise Managers.
   b. (Optional) Configure include and exclude elements to allow and disallow agent connections to specific Enterprise Managers
   c. Save loadbalancing.xml.

   The new values are put into effect when the Enterprise Manager checks the property file during a hot deploy, which occurs about every 60 seconds. The MOM sends the updated loadbalancing.xml to Collectors in the cluster at the next load balance interval.
5. Perform system monitoring and reconfigure if needed.
   a. Reconfigure any properties as needed.
   b. Update loadbalancing.xml.

CA Technologies Proprietary Documentation
Authorized Use Only

Plan the Agent - Enterprise Manager Network Topology

Plan your agent to Enterprise Manager or Collector topology before configuring agent connection assignments. In theses assignments, you configure which agents are allowed or not allowed to connect to specific individual or groups of Collectors.

If your business practices require it, set up agents into subclusters. You can assign agents to a specific set of Collectors, then configure loadbalancing.xml so the MOM load balances only that set of Collectors. For example, you could have the agents monitoring the applications for a specific business unit connect to three Collectors in your cluster of ten.

During your planning, ask and answer questions such as:

- Are classes or groups of agents or specific agents allowed or disallowed in the cluster or on specific Collectors?
- Should any Collectors be specially protected against overloading?

  For example, do some Collectors need protection against direct connections from mis-configured agents? If all agents have the identical agent profile and attempt to connect, they would flood a single IP address. If there are many agents and they are all allowed to connect, this situation can bring down a Collector.

Follow these steps:

1. For each group or class of agent, determine which set of Collectors you want to:

- Allow access.
- Disallow access.

1. Determine default connection activity of agents not matched in loadbalancing.xml.

   Do you want to allow or disallow all agents not matched in loadbalancing.xml on specific Enterprise Managers?

2. Will any agents be load balanced? If so, are they allowed on all the Collectors that you want them moved to?

   Carefully consider the effect of agent connection restrictions, as they restrict the load balancing options for the MOM. A MOM performing load balancing cannot remove an agent from an overloaded Collector when the agent is not allowed to connect to an under-loaded Collector. In this case, the MOM must load balance other agents from the overloaded Collector to an under-loaded Collector.

3. Determine how agents disconnected from an Enterprise Manager will act.

   Do you want agents to try to reconnect continuously to any allowed Enterprise Manager?

   Or based on the agent CA APM version is it better for:

   - Pre-9.1 agents to disconnect from the Enterprise Manager and keep trying to reconnect?
   - 9.1 agents to remain connected to the final allowed Enterprise Manager and not send any metric data to the Enterprise Manager?

4. Does your CA APM environment allow agents to connect across the following:

   - Clusters?

CA Technologies Prerelease Documentation. Authorized Use Only

- Collectors and Standalone Enterprise Managers?
- Any combination of clusters, Collectors, and Standalone Enterprise Managers?

Configure if New Agent Connections are Allowed When No loadbalancing.xml Rules

Determine you if want to allow or disallow all new agent connections when the agents do not match any regular expressions in loadbalancing.xml. Make this determination for:

- MOM.
- Each Collector.
- Any standalone Enterprise Managers.

You can set the introscope.apm.agentcontrol.agent.allowed property as a performance measure.

Note: The MOM property value always overrides the Collector value. The Collector value is used only until the MOM connects to the Collector and the MOM property values are used.

You can also use this property setting and configurations in loadbalancing.xml to fine-tune the agent to Enterprise Manager connections in your environments.

Follow these steps:

1. Go to the <*EM_Home*>\config directory and open the IntroscopeEnterpriseManager.properties file.
2. Set introscope.apm.agentcontrol.agent.allowed= true.
3. Save the IntroscopeEnterpriseManager.properties file.

   The setting is applied when the Enterprise Manager checks the property file during a hot deploy, which occurs about once every 60 seconds.

CA Technologies Prerelease Documentation. Authorized Use Only.

Configure Agent Reconnection Wait Time

You can vary the number of seconds that agents must wait before trying to reconnect to a standalone Enterprise Manager or Collector. A varied waiting period helps CA APM avoid performance problems that happen when many agents try to reconnect to an Enterprise Manager at the same time. You configure the introscope.apm.agentcontrol.agent.reconnect.wait property to vary the wait time.

Pre-9.1 agents are disconnected when either:

- They are disallowed from an Enterprise Manager.
- No Enterprise Manager is available for connection.

9.1 and later agents are disconnected when the:

- Agent is disallowed.
- Disallowed agents connection clamp value (see page 430) is hit.

Follow these steps:

1. Go to *<EM_Home>*\config directory and open the IntroscopeEnterpriseManager.properties file.
2. Set the introscope.apm.agentcontrol.agent.reconnect.wait property.
3. Save the IntroscopeEnterpriseManager.properties file.

   The setting is applied when the Enterprise Manager checks the property file during a hot deploy, which occurs about once every 60 seconds.

   When disconnected from an Enterprise Manager, an agent tries to reconnect based on a reconnection time generated by the Enterprise Manager. The connection time is based on a default value and a randomized additional positive or negative value.

Configure Disallowed Agent Connection Clamp Limit

You can limit the number of 9.1 agents that are disallowed but remain connected to a standalone Enterprise Manager or Collector. Clamping disallowed agents is performance measure because it limits the number of connections to an Enterprise Manager. Each connection requires some resources, so an overly high number of connections can decrease Enterprise Manager performance.

Follow these steps:

1. Go to the *<EM_Home>*\config directory and open the IntroscopeEnterpriseManager.properties file.

CA Technologies Prerelease Documentation. Authorized Use Only.

2. Set the introscope.enterprisemanager.agent.disallowed.connection.limit property.
3. Save the IntroscopeEnterpriseManager.properties file.

   The setting is applied when the Enterprise Manager checks the property file during a hot deploy, which occurs about once every 60 seconds.

   When the clamp is hit and a disallowed agent tries to connect to the Enterprise Manager, the Enterprise Manager:

   - Logs the event.
   - Does not accept the new connection.

Configure Agent Continual Reconnect to Any Allowed Enterprise Manager

You can configure whether agents that are denied connections on all Enterprise Managers can continually cycle through their allowed Enterprise Managers trying to reconnect. The following are the allowed Enterprise Managers:

- Enterprise Managers configured in the agent profile introscope.agent.enterprisemanager.connectionorder property value.
- Any Enterprise Manager allowed based on loadbalancing.xml configuration.

If a disallowed agent cannot connect to an Enterprise Manager, it:handles connections in these ways:

- Tries to connect to the next allowed Enterprise Manager on the Allowed Enterprise Manager List.
- Does not connect to any Enterprise Manager on which it is disallowed.

Configuring disallowed agents to keep trying to connect has these characteristics:

- Is useful for environments where agents are allowed to connect across these CA Introscope components:
  - Clusters.
  - Collectors and Standalone Enterprise Managers.
  - Any combination of clusters, Collectors, and Standalone Enterprise Managers.
- Forces disconnected agents to keep trying to connect to allowed Enterprise Managers under these conditions:
  - Until an Enterprise Manager is available for connection.
  - Even if the Enterprise Managers are in different clusters.

CA Technologies Proprietary Documentation. Authorized Use Only.

Otherwise, 9.1 and later agents handle connections in these ways:

- Connect to an Enterprise Manager in disallowed mode
- Do not try to connect to allowed Enterprise Managers that becomes available.

Follow these steps:

1. Go to <*Agent_Home*>/core/config directory and open the IntroscopeAgent.profile file.
2. Set the introscope.agent.enterprisemanager.failbackRetryIntervalInSeconds property.
3. Save the IntroscopeEnterpriseManager.properties file.
4. Restart the managed application.

   If these events occur:

   - Agent is disallowed on an Enterprise Manager
   - Enterprise Manager is down

   Then the agent immediately tries to connect to the next allowed Enterprise Manager on the Allowed Enterprise Manager List.

   If the agent cannot connect to any allowed Enterprise Manager in allowed mode, then the agent performs these actions:

   - Waits for the duration set in the introscope.agent.enterprisemanager.failbackRetryIntervalInSeconds property.
   - Retries connecting to allowed Enterprise Managers starting with the first Enterprise Manager on its Allowed Enterprise Manager List.

   This cycle repeats until the agent connects to an Enterprise Manager in allowed mode.

CA Technologies Pre-release Documentation. Authorized Use Only.

Configure loadbalancing.xml for Allowed and Disallowed Agents by Enterprise Manager

Configure loadbalancing.xml (see page 118) on the MOM or standalone Enterprise Manager to allow and disallow agent connections on specific Enterprise Managers.

Every 10 minutes the MOM sends the agent connection information in loadbalancing.xml to Collectors, which then send allowed Enterprise Manager information to agents. The MOM allows and disallows agent connections on specific Enterprise Managers based on the latest loadbalancing.xml file.

If you have an agent that you never want the MOM to load balance, you can also configure loadbalancing.xml. In this example, the MOM never load balances an agent named EPAgent from a Collector named Collector1-ip to another Collector.

```
<agent-Collector name="EPAgent connects to only Collector1-ip">
<agent-specifier>.*\|.*\|EPAgent</agent-specifier>
    <include>
        <Collector host="Collector1-ip" port="6001"/>
    </include>
</agent-Collector>
```

Agent Connection Scenarios

Situation where if there's a 200 agent clamp and 500 agents try to connect, the EM gets overloaded with attempted connects and isn't able to perform its standard job.

A disconnected 9.1 agent cycle through its Allowed Enterprise Manger list, and finds no allowed available Enterprise Managers. The agent connects in parks in a dormant/passive state to the last collector in its list. The agent displays in the APM Status console Denied Agents list.

An agent connects to a Collector and gets the latest Allowed Enterprise Manager list. The agent then disconnects from the Collector. The MOM sends the Collector an updated Allowed Enterprise Manager list, that changes the Collectors to which the agent is allowed to connect. When the agent is able to connect to a Collector, the agent gets the updated Allowed Enterprise Manager list. If the MOM goes down, the agent still has the most current Allowed Enterprise Manager list.

CA Technologies Prerelease Documentation
Authorized Use Only

Troubleshoot Agent - Enterprise Manager Network Topology

How Enterprise Manager Handles Disallowed Agents

The MOM sends loadbalancing.xml to all Collectors in the cluster every 10 minutes. The loadbalancing.xml file tells Enterprise Managers which agents are allowed and disallowed to connect to it.

Here's what happens when an Enterprise Manager disallows an agent connection.

Pre-9.1 agents

The Enterprise Manager disconnects pre-9.1 agents, which then try to connect to another allowed available Enterprise Manager. If the agent is not allowed to connect to any Enterprise Manager, it continues trying to connect to the last Enterprise Manager that is:

- Available.
- Allowed to connect to the agent.

9.1 agents

The Enterprise Manager disconnects 9.1 agents, which then try to connect to another allowed available Enterprise Manager. If the agent is not allowed to connect to any Enterprise Manager, it connects to the MOM. The agent remains connected to the MOM but, does not send any data.

When Collector Uses introscope.apm.agentcontrol.agent.allowed Property

When the MOM connects to a Collector, the MOM provides the introscope.apm.agentcontrol.agent.allowed property and loadbalancing.xml to the Collector. If the MOM goes down, the Collector uses the following information provided by the MOM to know which agent connections to allow and disallow:

- Current introscope.apm.agentcontrol.agent.allowed property setting in the IntroscopeEnterpriseManager.properties file.
- loadbalancing.xml.

If the MOM has never connected to the Collector, the MOM has not transferred its introscope.apm.agentcontrol.agent.allowed property and loadbalancing.xml file to the Collector. Therefore, the Collector uses its own introscope.apm.agentcontrol.agent.allowed property to decided whether to allow or disallow specific agent connections.

CA Technologies Pre-release Documentation. Authorized Use Only.

Agent Does Not Match loadbalancing.xml Regular Expression as Allowed or Disallowed

The value of introscope.apm.agentcontrol.agent.allowed property in the IntroscopeEnterpriseManager.properties file determines whether agents not matched by regular expressions in loadbalancing.xml can connect to a Collector.

If the property value is true, then agents that do not match any regular expression in loadbalancing.xml are allowed to connect to Collectors.

If the property value is false, then agents that do not match any regular expression in loadbalancing.xml are disallowed from connecting to Collectors.

All Collectors Assigned to an Agent are Unavailable and Other Collectors are Available

Pre-9.1 agents

The Collector disconnects the agents, which then try to connect to the MOM. With no Collectors available to which the agents are allowed to connect, the MOM denies every agent connection attempt. The agent keeps trying to connect to the MOM until an allowed Collector for the agent is available. The MOM redirects the agents to an allowed Collector when the Collector becomes available.

9.1 agents

The Collector disconnects the agents, which then connect to the MOM but, do not send any data. The MOM redirects the agents to an allowed Collector when the Collector becomes available.

How Collector Handles a Disallowed Agent with a Direct Connection

The MOM periodically sends loadbalancing.xml to all Collectors in the cluster. The loadbalancing.xml file makes each Collector aware of the agents that are allowed to connect to it. If loadbalancing.xml is updated so that an agent must connect to a Collector and the agent is disallowed, the following actions take place.

Pre-9.1 agents

The Collector denies the agent connection and the agent continues the connection attempt.

CA Technologies Prerelease Documentation. Authorized Use Only.

9.1 agents

The Collector gives latest allowed Collectors list to the agent. The agent is not allowed on the disallowed Collector. If MOM available then MOM redirects the Agent to valid allowed Collector. If MOM is not available, then the agent uses its allowed Collector list to iterate through until connects to allowed Collector.

Default Agent Connection Mode When introscope.apm.agentcontrol.agent.allowed Property is Missing

Allowed

How CA Introscope Determines Where Agents Connect When No MOM Available

The Collector to which the agent is trying to connect decides whether the connection is allowed or disallowed based on its current loadbalancing.xml file.

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 9: Using the Command-Line Workstation

This chapter describes the Introscope Command Line Workstation (CLW), a command-line interface to the Enterprise Manager.

This section contains the following topics:

[About the Command Line Workstation (CLW)](#) (see page 141)
[Running the CLW](#) (see page 142)
[CLW command reference](#) (see page 145)
[Sample scripts](#) (see page 186)

About the Command Line Workstation (CLW)

The Introscope Command Line Workstation (CLW) is available for performing some of the Introscope Workstation functions.

The CLW is useful for automating or scheduling Introscope operations tasks. For example, you can perform repetitive tasks such as disabling multiple alerts in several Management Modules more quickly with the CLW than with the Workstation interface. Similarly, the CLW is useful for scheduling tasks that must be performed automatically or on a periodic basis.

Using CLW commands you can:

- **Manage Management Modules and Elements**—You can list, copy, rename, and delete Management Modules and Elements, see if they are active, and change their current state.

  You can create Management Modules and Elements based on existing ones, using CLW copy and rename commands. You cannot, however, create Management Modules and Elements from scratch, or modify their details using the CLW.

- **Control Agents**—You can start, stop, and resume agent reporting remotely, and copy configuration files to remote agents.

CA Technologies Product Release Documentation. Authorized Use Only.

- **Start Transaction Trace Sessions**—You can write scripts to automatically start a Transaction Trace Session as a result of alert conditions.
- **Shut Down the Enterprise Manager**—You can write scripts to shut down the Enterprise Manager on scheduled basis. For example, you might want to stop the Enterprise Manager during off-hours, when using the Workstation might be inconvenient.
- **If the CLW exits**—CLW exits with a 0 value if nothing goes wrong. It will exit with a negative status code if one of the following happens:
  - Enterprise Manager authentication failure.
  - Connection problem with the Enterprise Manager such as a dropped network connection, or invalid host name.
  - There is an invalid WS and Enterprise Manager.

Running the CLW

This section is an overview of how to run the CLW. For information on specific commands, see CLW command reference (see page 145).

JVM requirements

The CLW requires JVM 1.6 or later.

Where to run CLW commands

You can include CLW commands in script or batch files, or enter them at the command line. The CLW *.jar* file is installed in the *<EM_Home>/lib* directory. Run CLW commands from that directory, or include the path to the directory in the commands.

CA Technologies Prerelease Documentation. Authorized Use Only.

Using CLW interactive mode

You can change from the default mode to interactive mode by using the *-i* flag, such as in the following example. If you plan on writing scripts for CLW, this mode cannot be used, as it will break them.

java -Xmx128M <EM_logon> <TT_properties> -jar CLWorkstation.jar > usage | -i | <command>

Interactive mode will read command line arguments for a command, execute, then exit. However, if no command line arguments are passed, other than JVM options like *-DHost=*, then it will read and execute one line of standard input and will continue to do so until it reaches the end of the file is reached or you exit the system. If no arguments are passed, and no input is provided, an error will be printed and it will exit the system. All lines that begin with the # symbol will be ignored.

Invoking CLW

CLW command syntax is:

```
java –Xmx128M <EM_logon> -jar CLWorkstation.jar <command>
```

Where:

- *–Xmx128M* sets the Java heap size, preventing *java.lang.OutOfMemoryError* errors that might otherwise occur.
- *<EM_logon>* is your Enterprise Manager logon information:

  ```
  -Duser=<user name>
  –Dpassword=<password>
  –Dhost=<Enterprise Manager host / ip address>
  –Dport=<Enterprise Manager port number>
  ```

- *<command>* is one of the commands described in the CLW command reference (see page 145).

For example:

```
java –Xmx128M -Duser=jdoe –Dpassword=mypassword –Dhost=jdoeDT
–Dport=5001 -jar CLWorkstation.jar <command>
```

Default logon values for accessing the Enterprise Manager are used if you invoke the CLW without supplying your logon information:

```
java –Xmx128M -jar CLWorkstation.jar <command>
```

Every CLW command you enter at the command line or in a script must start with the CLW invocation:

```
java –Xmx128M -jar CLWorkstation.jar
```

Note: To avoid getting a *java.lang.OutOfMemoryError*, set the Java heap size.

CA Technologies Prerelease Documentation. Authorized Use Only.

CLW is case-sensitive

CLW commands accept arguments that specify agents, Management Modules, and other named items in your Introscope deployment. CLW is case-sensitive—you must enter the names of Introscope objects exactly as they are named. For example, if an agent is named *WebLogic1*, a CLW command that refers to *weblogic1* does not execute for that agent.

CLW command output

CLW commands return their output to the command window. You can direct command output to a file, as appropriate.

CLW and regular expressions

You can use CLW commands to manipulate elements of your Introscope environment, such as agents, Management Modules, Dashboards, actions, and alerts. Depending on the type of target, you can perform operations such as copy, list, rename, activate, and deactivate.

You designate the target or targets of a CLW command using exact text strings or Perl regular expressions. For example, this CLW command lists all agents connected to the Enterprise Manager:

```
list agents matching (.*)
```

This command lists only agents whose name contains the string *weblogic*:

```
list agents matching .*weblogic.*
```

Backslashes and quotes in CLW commands

When you supply CLW command arguments that contain spaces, surround the argument with quotation marks ("), and include a backslash in front of each quotation mark, to prevent the shell from interpreting the quotation mark. For example:

```
\"2005/01/01 12:20:37\"
```

Using CLW in different environments

Regular expression syntax varies depending on your operating system and command shell, and in some cases, based on whether you enter the expression at the command line or in a script. The sample commands in this chapter are written for a Windows environment, and reflect the syntax required to use regular expression in a batch file.

CA Technologies Pre-release Documentation. Authorized Use Only

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.

    Note: The default value for this property is false.

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

    The default value for this property is INFO.

CLW command reference

This section contains details about CLW commands.

Note: The syntax shown for each command does not include the invocation portion of the command. As described in Invoking CLW (see page 143), each CLW command must begin with: *java -jar CLWorkstation.jar*

Enterprise Manager commands

The Enterprise Manager aggregates application performance data from agents, runs requested calculations, makes performance data available to authorized clients, and stores performance data.

Shut down Enterprise Manager

This command shuts down the Enterprise Manager. To run it, you must have shutdown or full permission. Any Workstation users connected to the Enterprise Manager are logged off when you shut down the Enterprise Manager.

Syntax

shutdown

CA Technologies Prerelease Documentation. Authorized Use Only.

Example

shutdown

Output

None.

Extract persistent data (time range)

This command extracts historical data for a specified time period from SmartStor. You supply a regular expression that specifies the agents, a regular expression that specifies the metric groups, and the start and end of the time period. The *hh:mm:ss* portion of the date-time string is optional—if you do not supply it, the default is 00:00:00 (midnight).

Syntax

get historical data from agents matching <REGULAR EXPRESSION> and metrics matching <REGULAR EXPRESSION> between <YYYY-MM-DD [hh:mm:ss]> and <YYYY-MM-DD [hh:mm:ss]>

Example

get historical data from agents matching [REGULAR EXPRESSION] and metrics matching [REGULAR EXPRESSION] between \"2011-10-12 00:00:00\" and \"2011-10-12 00:00:01\" with frequency of [EXACT MATCH] (second|seconds|sec|secs|s)

This command extracts the persisted data for the *Servlets:Average Response Time \(ms\)* Metric Grouping, for all agents whose name includes the string *weblogic*, for the time period from midnight on October 12, 2011 through midnight on October 12, 2011.

Output

The data is returned in comma-separated value (CSV) format.

Extract persistent data (previous N minutes with frequency of N seconds)

This command extracts historical data for a specified number of minutes prior to execution of the command, at a specific frequency, in seconds. You supply a regular expression that specifies the agents, a regular expression that specifies the Metric Groups, and the number of minutes of history to be extracted, and the frequency at which to extract the data.

CA Technologies Pre-release Documentation. Authorized Use Only.

Specify a pathname for authentication

Specifies the full pathname for a public key file to use instead of username/password authentication when connecting to the Enterprise Manager. The public key used is the same one used for Manager of Managers/Collector authentication.

For more information about configuring the public key, see the *CA APM Security Guide*.

Introscope agent commands

An agent runs as a part of the managed application on the Java Application machine. It collects and summarizes the probe-reported data and sends it to the Enterprise Manager.

List Agents

The List Agents command lists one or more agents that are connected to the Enterprise Manager. You supply a regular expression that determines which agents are listed.

Syntax

list agents matching <REGULAR EXPRESSION>

Example

list agents matching (.)*

This command lists all agents that are connected to the Enterprise Manager.

Output

Each agent is listed, prefixed by its Host and Process. Virtual agents are included in the list. For example:

jdoe-dt|Weblogic|WebLogicAgent

Custom Metric Host (Virtual)|Custom Metric Process (Virtual)|Custom Metric Agent (Virtual)

Turn on Agents

This command starts or resumes Metric reporting from one or more agents. You supply a regular expression that specifies the target agents.

CA Technologies Prerelease Documentation. Authorized Use Only.

Syntax

turn on agents matching <REGULAR EXPRESSION>

Example

turn on agents matching (.)*

This command turns on all agents.

Output

None.

Turn off Agents

This command stops Metric reporting from one or more agents. You supply a regular expression that specifies the target agents.

Send configuration file

This command sends a specified file to one or more agents. You supply the name of the file, a regular expression that specifies the target agents, and the source and destination directories.

Target agents must be configured to allow remote writes in their *introscope.agent.remoteagentconfiguration.enabled* property. The location and its default values vary. For more information, see the *CA APM Java Agen*t Implem*entation Guide* or the *CA APM .NET Agent Implementation Guide*.

You can only copy files explicitly specified in the agent's *introscope.agent.remoteagentconfiguration.allowedfiles property*.

The default value of this property is *domainconfig.xml*, which is the file that contains all the information that Introscope needs to recognize transactions for CA CEM.

APM Status Console Commands

This section provides CLW commands related to the CA APMStatus console. For information about using the APM Status console, see the *CA APM Workstation User Guide*.

You must have apm_status_console_control permission to run CA APM Status console CLW commands. For information about this permission, see the *CA APM Security Guide.*

CA Technologies Pre-release Documentation. Authorized Use Only.

This CLW category includes these commands:

Get Cluster Configuration (see page 149)
List Active Clamps (see page 149)
List Denied Agents (see page 151)
List Important Events (see page 152)

Get Cluster Configuration

This command displays the host name and port of all the Collectors in the cluster.

Syntax

get cluster configuration

Output

Lists the host names and ports of all the Collectors in your cluster. For example:

get cluster configuration
Triage Map Configurations.By Frontend|TradeService|Health:Average Response Time (ms)_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|loginAdmin.action|loginAdmin.action:Errors Per Interval_Summary_MetricAlert
Triage Map Configurations.By Frontend|ReportingService|Backend Calls|WebServices:Average Response Time (ms)_Location_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatients.action:Average Response Time (ms)_Location_MetricAlert

Example

get cluster configuration

This command list all Collectors in your cluster

List Active Clamps

This command lists the clamps properties in apm-events-thresholds-config.xml that currently meet or exceed their threshold values. These are called *active* clamps. This command provides a list of active clamps across all Collectors in a cluster or on a standalone Enterprise Manager.

When an apm-events-thresholds-config.xml clamp value meets or exceeds its threshold value, the Enterprise Manager notifies the APM Status console.

CA Technologies Confidential Release Documentation. Authorized Use Only.

Syntax

list active clamps

Output

Lists the name of all the active clamps in your cluster or standalone Enterprise Manager. For example:

ClampName,CollectorName,ThresholdValue,CurrentValue,LatestClampedTime

Triage Map Configurations.By Frontend|TradeService|Health:Average Response Time (ms)_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|loginAdmin.action|loginAdmin.action:Errors Per Interval_Summary_MetricAlert
Triage Map Configurations.By Frontend|ReportingService|Backend Calls|WebServices:Average Response Time (ms)_Location_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatients.action:Average Response Time (ms)_Location_MetricAlert

Example

list active clamps

This command will try to list all the active clamps in the cluster or on the standalone Enterprise Manager.

Command options

This section describes *list active clamps* command options.

CA Technologies Prerelease Documentation. Authorized Use Only

List active clamps for specific and multiple Collectors

You can add these two arguments to the *list active clamps* command:

- of COLLECTORNAME—Use this argument to return a list of the active clamps on a specific Collector.
  - Syntax

    list active clamps of <CollectorName>
  - Output

    ClampName,CollectorName,ThresholdValue,CurrentValue,LatestClampedTime
  - Example

    list active clamps of <CollectorName>
- of COLLECTORNAME1,COLLECTORNAME2 —Use this argument to return a list of the active clamps on multiple specific Collectors.
  - Syntax

    list active clamps of <CollectorName1>, <CollectorName2>
  - Output

    ClampName,CollectorName,ThresholdValue,CurrentValue,LatestClampedTime
  - Example

    list active clamps of <CollectorName1>, <CollectorName2>

List Denied Agents

This command lists the agents that are currently denied connections across all Collectors in a cluster or to a standalone Enterprise Manager.

Syntax

list denied agents

CA Technologies Prerelease Documentation. Authorized Use Only.

Output

Lists the names of all the agents in your cluster or on your standalone Enterprise Manager that are currently denied a Enterprise Manager connection. For example:

List of agent names

Triage Map Configurations.By Frontend|TradeService|Health:Average Response Time (ms)_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|loginAdmin.action|loginAdmin.action:Errors Per Interval_Summary_MetricAlert
Triage Map Configurations.By Frontend|ReportingService|Backend Calls|WebServices:Average Response Time (ms)_Location_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatients.action:Average Response Time (ms)_Location_MetricAlert

Example

list denied agents

This command list all agents denied Enterprise Manager connections across your cluster or on your standalone Enterprise Manager.

List Important Events

This command lists the Important Event properties in apm-events-thresholds-config.xml that currently meet or exceed their threshold values. This command provides a list across all Collectors in a cluster or for a standalone Enterprise Manager.

When an apm-events-thresholds-config.xml property value meets or exceeds its threshold value, the Enterprise Manager notifies the APM Status console.

Syntax

list important events

CA Technologies Prerelease Documentation. Authorized Use Only.

Output

Lists the data about all the important events in your cluster or standalone Enterprise Manager. For example:

DegreeOfImportance, eventdescription, eventtime

Triage Map Configurations.By Frontend|TradeService|Health:Average Response Time (ms)_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|loginAdmin.action|loginAdmin.action:Errors Per Interval_Summary_MetricAlert
Triage Map Configurations.By Frontend|ReportingService|Backend Calls|WebServices:Average Response Time (ms)_Location_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatients.action:Average Response Time (ms)_Location_MetricAlert

Example

list important events

This command list all the important events in the cluster or on the standalone Enterprise Manager.

Management Module commands

Management Modules organize Elements in the Workstation, making them easier to find and manipulate. CLW commands for Management Modules simplify the process of moving a Management Module and the Elements it contains from one Enterprise Manager to another.

List Management Modules

The List Management Module command lists one or more Management Modules. You supply a regular expression that specifies the Management Modules.

Syntax

list management modules matching <REGULAR EXPRESSION>

Example

list management modules matching (.)*

This command lists all Management Modules.

CA Technologies Prerelease Documentation. Authorized Use Only.

Output

The name of each Management Module is listed:

```
Sample
Sample13
Sample11
```

Query Management Modules

The Query Management Module command lists Management Modules and their current state: active or inactive. You supply a regular expression that specifies the target Management Modules.

Syntax

query state of management modules matching <REGULAR EXPRESSION>

Example

query state of management modules matching (.)*

This command returns the name and state of all Management Modules.

Output

The name and state of each Management Module is listed:

```
Sample1:Active
Sample2:Active
```

Activate Management Modules

The Activate Management Module command activates one or more Management Modules. You supply a regular expression that specifies the target Management Modules.

Deactivate Management Modules

The Deactivate Management Module command deactivates one or more Management Modules. You supply a regular expression that specifies the target Management Modules.

Note: The Deactivate Management Module command cannot deactivate Triage Map Configuration Management Modules.

CA Technologies Prerelease Documentation. Authorized Use Only.

Copy Management Module

The Copy Management Module command creates a copy of a Management Module in the *<EM_Home>/config/modules* directory.

This command is an efficient way to create one or more new Management Modules that are based on existing Management Modules. For example, to create a new Management Module that is identical to an existing one, except for its Metric Group, you copy it using the CLW, and then modify the new Management Module's Metric Group in the GUI Workstation.

In the Copy Management Module command, you identify the Management Module you want to copy, a name for new Management Module, and a name, excluding the file extension, for the *.jar* file for the new Management Module.

Note: The Deactivate Management Module command cannot deactivate Triage Map Configuration Management Modules.

Delete Management Module

The Delete Management Module command deletes a Management Module. You supply the exact name of the Management Module you want to delete.

Read-only Management Modules

This command sets Management Modules that match a specified regular expression to read-only.

WARNING! CA Technologies recommends that you save a version of your editable Management Module to a safe location for future use. Once a Management Module is set to read-only, you cannot return it to an editable state.

To set Management Modules to read-only, use this command:

set readonly management modules matching <regular expression>

Then log into the Workstation, open the Management Module Editor and note that management modules matching the regular expression have a lock icon next to them, have the Editable checkbox unchecked—changes to these Management Modules are no longer allowed.

Note: This feature is intended for use by CA Technologies partners. Please contact CA Support before using this feature in your environment.

Action Commands

An Action is a behavior, such as a notification, that can be triggered by an Alert.

CA Technologies Prerelease Documentation
Authorized Use Only.

List Actions

The List Actions command lists one or more Actions in one or more Management Modules. You supply two regular expressions—one that specifies the Actions, and one that specifies the Management Modules.

Syntax

list actions matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

list actions matching (.) in management modules matching (.*)*

This command lists all Actions in all Management Modules.

Output

Each Action is listed, prefixed by the Management Module that contains it:

Sample.SMTP Email Action

Sample10.Workstation Notification

Sample2.Workstation Notification

Activate Actions

The Activate Actions command activates one or more Actions in one or more Management Modules. You supply two regular expressions—one that specifies the Actions to activate, and one that specifies the Management Modules in which to activate them.

Syntax

activate actions matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

activate actions matching (SMTP.) in management modules matching (.*)*

This command activates Actions that begin with the string *SMTP* in all Management Modules.

Output

Each Action activated by the command is listed, prefixed by the Management Module that contains it, suffixed by its activation state:

CA Technologies Prerelease Documentation
Authorized Use Only.

Sample.SMTP Email Action:Active
Sample10.Workstation Notification:Active

Sample2.Workstation Notification:Active

Deactivate Actions

The Deactivate Actions command deactivates one or more Actions in one or more Management Modules. You supply two regular expressions—one that specifies the Actions to deactivate, and one that specifies the Management Modules in which to deactivate the Actions.

Query Actions

The Query Actions command lists the state—active or inactive—of one or more Actions in one or more Management Modules. You supply two regular expressions—one that specifies the Actions, and one that specifies the Management Modules.

Rename Action

The Rename Action command renames an Action. You specify the name of the Action and the Management Module that contains it, and a new name for the Action.

Delete Action

The Delete Action command deletes one or more Actions from one or more Management Modules. You supply two regular expressions—one that specifies the Actions, and one that specifies the Management Modules from which to delete them.

Note: You cannot delete an Action that is used or referenced by another Element—for example, one that is linked to an Alert.

Alert commands

An Alert indicates potential problems in your managed Java Application by comparing Metric values to user-defined threshold values.

List Alerts

The List Alerts command lists one or more Alerts in one or more Management Modules. You supply two regular expressions—one that specifies the Alerts, and one that specifies the Management Modules.

CA Technologies Prerelease Documentation
Authorized Use Only.

Syntax

list alerts matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

list alerts matching (.) in management modules matching Sample*

This command lists the Alerts in the Management Module named *Sample*.

Output

Each Alert is listed, prefixed by the Management Module that contains it:

Sample.Console Summary Alert

Sample.All Servlet Response Time

Sample.All EJB Response Time

Activate Alerts

The Activate Alerts command activates one or more Alerts in one or more Management Modules. You supply two regular expressions—one that specifies the Alerts to activate, and one that specifies the Management Modules in which to activate those Alerts.

Syntax

activate alerts matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Output

Each Alert activated by the command is listed, prefixed by the Management Module that contains it, and suffixed by its state:

Sample2.Application Summary Alert:Active

Sample2.SQL Average Processing Time #2 (ms):Active

Sample2.CPUApp1:Active

Example

activate alerts matching (.) in management modules matching Same2*

This command activates all of the Alerts in the Management Module named *Sample2*.

CA Technologies Pre-release Documentation. Authorized Use Only.

Deactivate Alerts

The Deactivate Alerts command deactivates one or more Alerts in one or more Management Modules. You supply two regular expressions—one that specifies the Alerts to deactivate, and one that specifies the Management Modules in which to deactivate those Alerts.

Query Alerts

The Query Alerts command lists the state—active or inactive—of one or more Alerts in one or more Management Modules. You supply two regular expressions—one that specifies the Alerts, and one that specifies the Management Modules.

Rename Alert

The Rename Alert command renames an Alert. You specify the name of the Alert and the Management Module that contains it, and a new name for the Alert.

Delete Alert

The Delete Alerts command deletes one or more Alerts from one or more Management Modules. You supply two regular expressions—one that specifies the Alerts, and one that specifies the Management Modules from which to delete them.

Note: You cannot delete a Alert that is used or referenced by another element—for example, one that is used on a dashboard.

Alert Downtime Schedule commands

Alert Downtime Schedules let you manage downtime periods from the Introscope Workstation Management Module Editor.

List Alert Downtime Schedules

The List Alert Downtime Schedules command lists one or more alert downtime schedules in one or more Management Modules. You supply two regular expressions—one that specifies the alert downtime schedule, and one that specifies the Management Modules.

Syntax

list alert downtime schedules matching [REGULAR EXPRESSION] in management modules matching [REGULAR EXPRESSION]

CA Technologies Pre-release Documentation. Authorized Use Only.

Example

List alert downtime schedules matching (.) in management modules matching Sample*

Output

Sample.Example Alert Downtime Schedule
Sample.Another Alert Downtime Schedule

Activate Alert Downtime Schedules

The Activate Alert Downtime Schedules command activates one or more alert downtime schedules in one or more Management Modules. You supply two regular expressions—one that specifies the alert downtime schedules to activate, and one that specifies the Management Modules in which to activate those alert downtime schedules.

Syntax

activate alert downtime schedules matching [REGULAR EXPRESSION] in management modules matching [REGULAR EXPRESSION]

Example

activate alert downtime schedules matching (.) in management modules matching Sample*

Output

Sample.Example Alert Downtime Schedule:Active

Sample.Another Alert Downtime Schedule:Active

Deactivate Alert Downtime Schedules

The Deactivate Alert Downtime Schedules command deactivates one or more alert downtime schedules in one or more Management Modules. You supply two regular expressions—one that specifies the alert downtime schedules to deactivate, and one that specifies the Management Modules in which to deactivate those alert downtime schedules.

Query Alert Downtime Schedules

The Query Alert Downtime Schedules command lists the state—active or inactive—of one or more alert downtime schedules in one or more Management Modules. You supply two regular expressions—one that specifies the alert downtime schedules, and one that specifies the Management Modules.

CA Technologies Prerelease Documentation. Authorized Use Only.

Rename Alert Downtime Schedules

The Rename Alert Downtime Schedules command renames an alert downtime schedule. You specify the name of the alert downtime schedule and the Management Module that contains it, and a new name for the alert downtime schedule.

Delete Alert Downtime Schedules

The Delete Alert Downtime Schedules command deletes one or more alert downtime schedules from one or more Management Modules. You supply two regular expressions—one that specifies the alert downtime schedule, and one that specifies the Management Modules from which to delete them.

Calculator commands

Calculators create custom Metrics by summing or averaging Metric data.

List Calculators

The List Calculators command lists one or more calculators in one or more Management Modules. You supply two regular expressions—one that specifies the calculators, and one that specifies the Management Modules.

Syntax

list calculators matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

list calculators matching (.) in management modules matching (.*)*

This command lists all calculators in all Management Modules.

Output

Each calculator is listed, prefixed by the Management Module that contains it:

Sample.Average Servlet Response Time

Activate Calculators

The Activate Calculators command activates one or more calculators in one or more Management Modules. You supply two regular expressions—one that specifies the calculators to activate, and one that specifies the Management Modules in which to activate those calculators.

CA Technologies Preferred Use Documentation
Authorized Use Only.

Syntax

activate calculators matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

activate calculators matching (.) in management modules matching Sample2*

This command activates all calculators in the Management Module named *Sample2*.

Output

Each calculator activated by the command is listed, prefixed by the Management Module that contains it, and suffixed by its state:

Sample2.Average Servlet Response Time:Active

Deactivate Calculators

The Deactivate Calculators command deactivates one or more calculators in one or more Management Modules. You supply two regular expressions—one that specifies the calculators to deactivate, and one that specifies the Management Modules in which to deactivate those calculators.

Query Calculators

The Query Calculators command lists the state—active or inactive—of one or more calculators in one or more Management Modules. You supply two regular expressions—one that specifies the calculators, and one that specifies the Management Modules.

Rename Calculator

The Rename Calculator command renames a calculator. You specify the name of the calculator and the Management Module that contains it, and a new name for the calculator.

Delete Calculators

The Delete Calculators command deletes one or more calculators from one or more Management Modules. You supply two regular expressions—one that specifies the calculators, and one that specifies the Management Modules from which to delete them.

Note: You cannot delete a calculator that is used or referenced by another Element, for instance, one that is used in a Dashboard.

CA Technologies Prerelease Documentation. Authorized Use Only.

Dashboard commands

A dashboard is a page in the Workstation that contains Data Viewer objects such as Graphs, Bar Charts, and alert status indicators, along with graphic and text objects.

List Dashboards

The List Dashboards command lists one or more Dashboards in one or more Management Modules. You supply two regular expressions—one that specifies the Dashboards, and one that specifies the Management Modules.

Syntax

list dashboards matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

*list dashboards matching (.*Triage) in management modules matching (.*)*

This command lists the Dashboards, in all Management Modules, whose name ends in the string *Triage.*

Output

Each dashboard is listed, prefixed by the Management Module that contains it:

Sample10.3. Triage

Sample2.3. Triage

Sample.3. Triage

Activate Dashboards

The Activate Dashboards command activates one or more dashboards in one or more Management Modules. You supply two regular expressions—one that specifies the Dashboards to activate, and one that specifies the Management Modules in which to activate those Dashboards.

Syntax

activate dashboards matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

CA Technologies Prerelease Documentation Authorized Use Only.

Example

activate dashboards matching (.) in management modules matching Sample2*

This command activates all dashboards in the Management Module named *Sample2*.

Output

Each dashboard activated by the command is listed, prefixed by the Management Module that contains it, and suffixed by its state:

Sample2.4b. Database Detail:Active

Sample2.1. Introscope Console:Active

Sample2.2. Detection:Active

Deactivate Dashboards

The Deactivate Dashboards command deactivates one or more dashboards in one or more Management Modules. You supply two regular expressions—one that specifies the Dashboards to deactivate, and one that specifies the Management Modules in which to deactivate those Dashboards.

Query Dashboards

The Query Dashboards command lists the state—active or inactive—of one or more dashboards in one or more Management Modules. You supply two regular expressions—one that specifies the dashboards, and one that specifies the Management Modules.

Rename Dashboard

The Rename Dashboard command renames a dashboard. Specify the name of the dashboard and the Management Module that contains it, and a new name for the dashboard.

Delete Dashboards

The Delete Dashboards command deletes one or more dashboards from one or more Management Modules. You supply two regular expressions—one that specifies the dashboards, and one that specifies the Management Modules from which to delete them.

Note: You cannot delete a dashboard that is used or referenced by another element. For example, you cannot delete a dashboard that is has a hyperlink to a data viewer on another dashboard.

CA Technologies Prerelease Documentation
Authorized Use Only.

Metric Grouping commands

A Metric Grouping specifies which Metrics to act upon; it is used as a building block for elements such as alerts and calculators.

List Metric Groups

The List Metric Groups command lists one or more Metric Groups in one or more Management Modules. You supply two regular expressions—one that specifies the Metric Groups, and one that specifies the Management Modules.

Syntax

list metric groups matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

*list metric groups matching (.*Response.*) in management modules matching Sample*

This command lists the Metric Groups in the Management Module named *Sample*, whose name includes the string *Response*.

Output

Metric Groups are listed in this format:

<Management Module>.<Metric Group>[Agent: "<Agent Expression>" Attribute: "<Metric Expression>"]

For example:

Sample.Servlet Average Response Time (ms) [Agent: "(.)\|(.*)\|(.*)" Attribute:\"Servlets:Average Response Time \(ms\)"]*

The sample above provides this information for the Metric Grouping:

Management Module—*Sample*

Metric Grouping Name—*Servlet Average Response Time (ms)*

Agent Expression—*(.*)\|(.*)\|(.*)*

Metric Expression—*Servlets:Average Response Time \(ms\)*

CA Technologies Prerelease Documentation Authorized Use Only.

Report Templates Commands

This section defines CLW commands for managing and generating reports. For more information about Introscope's reporting capabilities, see the Reporting chapter of the *CA APM Workstation User Guide*.

List Report Templates

The List Report Templates command lists one or more report templates in one or more Management Modules. You supply two regular expressions—one that specifies the report templates, and one that specifies the Management Modules.

Syntax

list report templates matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

list report templates matching (.) in management modules matching Sample*

This command lists all report templates in the Management Module named *Sample*.

Output

Each report template is prefixed by the Management Module that contains it.

Sample.QA/Test Application Performance
Sample.System Availability
Sample.Application Capacity Planning

Activate Report Templates

The Activate Report Templates command activates one or more report templates in one or more Management Modules. You supply two regular expressions—one that specifies the report templates to activate, and one that specifies the Management Modules in which to activate the report templates.

Syntax

activate report templates matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

activate report templates matching (.) in management modules matching Sample2*

This command activates all Report Templates in the Management Module named Sample2.

CA Technologies Prerelease Documentation
Authorized Use Only.

Output

Each report template activated by the command is listed, prefixed by the Management Module that contains it, and suffixed by its state:

Sample2.Application Capacity Planning:Active

Sample2.Production Application Health:Active

Sample2.System Availability:Active

Deactivate Report Templates

The Deactivate Report Templates command deactivates one or more report templates in one or more Management Modules. You supply two regular expressions—one that specifies the report templates to deactivate, and one that specifies the Management Modules in which to deactivate those report templates.

Query Report Templates

The Query Report Templates command lists the state—active or inactive—of one or more report templates in one or more Management Modules. You supply two regular expressions—one that specifies the report templates, and one that specifies the Management Modules.

Rename Report Template

The Rename Report Template command renames a report template. You specify the name of the report template and the Management Module that contains it, and a new name for the report template.

Delete Report Templates

The Delete Report Templates command deletes one or more report templates from one or more Management Modules. You supply two regular expressions—one that specifies the report templates, and one that specifies the Management Modules from which to delete them.

CA Technologies Prerelease Documentation. Authorized Use Only.

Generate Report

Use the Generate Report command to generate a report. You supply these arguments:

- report name
- start time
- end time
- management module name
- filename—This file is written to the current working directory, unless a path is specified. The file is written in the format associated with the file extension, which must be one of these: HTML, PDF, XLS, TEXT, XML, or CSV..

Commands

There are three Generate Report commands, basic, expanded, and an agent override command:

Basic command:

generate report named [EXACT MATCH] in management module named [EXACT MATCH] to [EXACT MATCH]

Expanded command:

generate report named [EXACT MATCH] in management module named [EXACT MATCH] starting at [EXACT MATCH] ending at [EXACT MATCH] to [EXACT MATCH]

Optional parameters

The optional parameters can be used with both the basic and expanded commands above:

{in agents matching [REGEX PATTERN]} {starting at [EXACT MATCH] ending at [EXACT MATCH]}

The syntax for the optional parameters is:

generate report named [EXACT MATCH] in management module named [EXACT MATCH] {in agents matching [REGEX PATTERN]} {starting at [EXACT MATCH] ending at [EXACT MATCH]} to [EXACT MATCH]

CA Technologies Prerelease Documentation. Authorized Use Only.

Report timestamp format

The report timestamp is presented in a variety of formats. The table below identifies the various possibilities.

Timestamp Format	Example
M/d/yy h:mm a	11/17/08 3:00 PM
M/d/yy h:mm:ss a	11/17/08 3:00:06 PM
M/d/yy h:mm:ss a z	11/17/08 3:00:06 PM PST
M/d/yy	11/17/08
MMM d, yyyy h:mm a	Nov 17, 2008 3:00 PM
MMM d, yyyy h:mm:ss a	Nov 17, 2008 3:00:06 PM
MMM d, yyyy h:mm:ss a z	Nov 17, 2008 3:00:06 PM PST
MMM d, yyyy	Nov 17, 2008
MMMM d, yyyy h:mm a	November 17, 2008 3:00 PM
MMMM d, yyyy h:mm:ss a	November 17, 2008 3:00:06 PM
MMMM d, yyyy h:mm:ss a z	November 17, 2008 3:00:06 PM PST
MMMM d, yyyy	November 17, 2008
EEE MMM dd HH:mm:ss zzz yyyy	Mon Nov 17 15:00:06 PST 2008
MM/dd/yy HH:mm:ss	11/17/08 15:00:06
yyyy-MM-dd HH:mm:ss	2008-11-17 15:00:06
yyyy-MM-dd HH:mm:ss.SSSS	2008-11-17 15:00:06.0840
yyyy-MM-dd HH:mm	2008-11-17 15:00
yyyy-MM-dd hh:mm:ss a	2008-11-17 03:00:06 PM
yyyy-MM-dd hh:mm a	2008-11-17 03:00 PM
yyyy/MM/dd hh:mm:ss a	2008/11/17 03:00:06 PM
yyyy/MM/dd hh:mm a	2008/11/17 03:00 PM
yyyy/MM/dd HH:mm:ss	2008/11/17 15:00:06
yyyy/MM/dd HH:mm	2008/11/17 15:00

CA Technologies Prerelease Documentation. Authorized Use Only

SNMP Collections commands

These commands define which collected Metrics are to be included in a published MIB.

List SNMP Collections

The List SNMP Collections command lists one or more SNMP collections in one or more Management Modules. You supply two regular expressions—one that specifies the SNMP Collections, and one that specifies the Management Modules.

Syntax

list snmp collections matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

list snmp collections matching (.) in management modules matching (.*)*

This command lists all SNMP Collections in all Management Modules.

Output

Each SNMP Collection is prefixed by the Management Module that contains it. For example:

Sample2.JDBC QueryUpdate

Activate SNMP Collections

The Activate SNMP Collections command activates one or more SNMP collections in one or more Management Modules. You supply two regular expressions—one that specifies the SNMP collections to activate, and one that specifies the Management Modules in which to activate those SNMP collections.

Syntax

activate snmp collections matching <REGULAR EXPRESSION> in management modules matching <REGULAR EXPRESSION>

Example

activate snmp collections matching (.) in management modules matching Sample2*

This command activates all SNMP collections in the Management Module named *Sample2*.

CA Technologies Prerelease Documentation
Authorized Use Only.

Output

Each SNMP collection activated by the command is listed, prefixed by the Management Module that contains it, suffixed by its state.

Sample2.JDBC QueryUpdate:Active

Deactivate SNMP Collections

The Deactivate SNMP Collections command deactivates one or more SNMP collections in one or more Management Modules. You supply two regular expressions—one that specifies the SNMP collections to deactivate, and one that specifies the Management Modules in which to deactivate those SNMP collections.

Query SNMP Collections

The Query SNMP collections command lists the state—active or inactive—of one or more SNMP collections in one or more Management Modules. You supply two regular expressions—one that specifies the SNMP collections, and one that specifies the Management Modules.

Rename SNMP Collection

The Rename SNMP collection command renames a SNMP collection. You specify the name of the SNMP collection and the Management Module that contains it, and a new name for the SNMP collection.

Delete SNMP Collections

The Delete SNMP Collections command deletes one or more SNMP collections from one or more Management Modules. You supply two regular expressions—one that specifies the SNMP collection, and one that specifies the Management Modules from which to delete them.

CA Technologies Prerelease Documentation. Authorized Use Only

Configuring Transaction Trace Options

The CLW gives you the option of configuring a variety of Transaction Trace options using these system properties:

- *-Dintroscope.clw.tt.console*—If set to "true", the transaction traces will be printed to a console.
- *-Dintroscope.clw.tt.dirname*—The local directory where transaction trace XML files are stored.
- *-Dintroscope.clw.tt.filename*—The file name used to store the transaction trace. This must not include the directory name.
- -Dintroscope.clw.tt.encoding—The encoding used when writing XML files to disk. Settings for this property include the default option *ISO-8859-1* and *UTF-8*, *UTF-16*, and *Windows-1252*

Starting a Transaction Trace session

The Transaction Tracer traces the activity of individual transactions as they flow through the boundaries of a single JVM.

Transaction Traces are automatically saved to the Transaction Events database, and by default are also captured in one or more XML files, which you can open and analyze in the Transaction Trace Viewer in the CA Introscope Workstation. See the *CA APM Workstation User Guide* for information about the CA Introscope Transaction Tracer.

CA Technologies Product Documentation.
Authorized Use Only.

Command syntax

There are many aspects of a Transaction Trace Session that a CLW command can specify. Some of these include:

- Execution threshold value—The minimum execution time for transactions to be traced. Tracing is performed for transactions that take longer than the threshold to execute. The threshold can be specified in seconds.
- Agent to trace—A Perl regular expression that specifies the agent or agents for which transactions are traced.
- Session duration—The duration of the Transaction Trace Session, in seconds.

  The syntax for starting a Transaction Trace Session, specifying the threshold in seconds, is:

  trace transactions exceeding <EXACT MATCH> (second|seconds|sec|secs|s) in agents matching <REGULAR EXPRESSION> for <EXACT MATCH> <second|seconds|sec|secs|s>

Other CLW command options are detailed in Command options, below (see page 173). For Transaction Trace CLW command examples, see Transaction Trace command examples (see page 176).

Specify Transaction Trace output file

You can specify the name of the Transaction Trace output file.

Syntax

Dintroscope.clw.tt.filename=<file name>

Example

Specifies the file name (but not the full path) of the transaction trace output file. Use in conjunction with a property to specify the output directory.

Output

Sends transaction trace output to the specified file.

Command options

This section describes Transaction Trace command options.

CA Technologies Prerelease Documentation Authorized Use Only.

Trace mode

The CLW supports two modes of capturing Transaction Trace Session results:

- Synchronous Mode—By default, the CLW writes the trace at the end of the session to the Transaction Events Database, and to a single XML file named *TransactionTraceData-yyyyMMddhhmmss.xml*, where the date-time string reflects the start of the session.

- Asynchronous Mode—In this mode, the CLW writes the trace for each transaction as soon as it is reported to the Transaction Events Database, and to its own XML file named *nnnn(sec)-TransactionTraceDatayyyyMMddhhmmss-<sequence number>.xml*, where *nnnn* is the transaction duration in seconds, and the date-time string reflects the start of the transaction. If you provide your own file name, it will append the sequence number before the last ".". If file name provided does not have ".", it will append the sequence number at the end of the file name.

Results recorded in Synchronous Mode are easier to view—all the data is in a single file. However, you cannot view the data until the end of the session, and for high-volume or complex transactions, synchronous capture can consume significant memory.

Recording transaction data asynchronously consumes less memory, and enables you to view transactions data before the end of the Transaction Trace Session.

-Dintroscope.clw.tt.mode=(Synch|Asynch)

Output to Transaction Events Database only

Use this option to run a transaction trace and save the output only to the Transaction Events Database (no XML file output). This command is useful when triggered by an Introscope Alert to automatically start tracing transactions when the Alert is triggered.

Precede the command with the word *quietly* as in:

```
quietly trace transactions exceeding 2 seconds in agents matching "(.*)" for 60 seconds
```

Transaction Trace file directory

Use this option to specify the directory to which the CLW writes Transaction Trace Session results. The directory must exist. If you specify a directory that does not exist, the Transaction Trace are written to the current working directory.

```
-Dintroscope.clw.tt.dirname=<directory>
```

Note: This property is deprecated with this release of Introscope.

CA Technologies Proprietary Document. Authorized Use Only.

Output to console

Use this option to redirect the output of a Transaction Trace Session to the command line shell. If you use this option, results are still sent to the Transaction Events Database, but not to a file.

```
-Dintroscope.clw.tt.console=<true|false>
```

Note: This property is deprecated with this release of Introscope.

Return only header data

In some cases, when retrieving a large amount of transaction trace data via CLW, the *traces.db* file may grow without bound, causing CLW to take a long time to return. Setting this property to false solves the problem by reducing the amount of data returned by the Enterprise Manager.

CLW will then return only header trace data. By default or when this property is set to true, getting transaction trace data via the CLW will return both the header and body.

```
-Dwily.ps.workstation.clws.gettracebody
```

Transaction Tracer Clamp

The size of an individual Transaction Trace is now clamped when it reaches 5000 trace components by default. Any Transaction Trace exceeding the clamp will be discarded at the agent and a warning message like the following will be logged in the agent log file:

```
2/23/07 05:33:18 PM PST [WARN] [IntroscopeAgent] Transaction trace component limit
of 5000 reached, recording will cease for this transaction.
```

To change the maximum number of components allowed in a transaction trace, change the following agent property:

```
introscope.agent.transactiontrace.componentCountClamp=max_TT_components (default
is 5000)
```

For example, the following line in the agent profile:

```
introscope.agent.transactiontrace.componentCountClamp=30000
```

will change the maximum components allowed in a Transaction Trace to 30,000.

WARNING!: If this clamp size is increased, the memory required for transaction traces may be higher and as such, it may be necessary for you to adjust the max heap size for the JVM accordingly or else the Enterprise Manager may run out of memory.

CA Technologies Pre-release Documentation. Authorized Use Only.

Transaction Trace command examples

This section provides examples of CLW commands to start a Transaction Trace Session. Each example command initiates a session 600 seconds long. Transactions that take longer than 5 seconds, from all agents, are traced.

Note: All example commands use default login settings to connect to the Enterprise Manager.

For examples of scripts that run Transaction Trace Sessions, see Sample scripts (see page 186).

Transaction Trace using default option values

This CLW command starts a Transaction Trace Session using default settings for:

- Trace Mode—Transaction Trace data is captured synchronously—all results are written to a single XML file.
- Directory Location—When *-Dintroscope.clw.tt.dirname* is not specified, trace results are written to the current working directory.

```
java -jar CLWorkstation.jar trace transactions exceeding 5 seconds in agents matching
"(.*)" for 600 seconds
```

Transaction Trace using asynchronous capture

This CLW command starts a Transaction Trace Session using the default directory location for session output:

```
java –Dintroscope.clw.tt.mode=Asynch -jar CLWorkstation.jar trace transactions
exceeding 5 seconds in agents matching "(.*)" for 600 seconds
```

Data is captured asynchronously—a separate XML file is created for each transaction.

Transaction Trace results written to a user-defined directory

This CLW command starts a Transaction Trace Session using the default setting for Trace Mode—Transaction Trace data is captured synchronously. All results are written to a single XML file:

```
java –Dintroscope.clw.tt.dirname=C:\mytraces -jar CLWorkstation.jar trace
transactions exceeding 5 seconds in agents matching "(.*)" for 600 seconds
```

This command specifies a non-default location for the XML file—the *C:\mytraces* directory.

CA Technologies Pre-release Documentation. Authorized Use Only.

Transaction Trace using parameter and error filters

You can initiate a Transaction Trace session that specifies the number of seconds to run and uses the same filters that are available in the Workstation—user ID, URL, URL query, session ID, request header, request parameter, session attribute, and error:

Filtering on User ID, URL, URL query, or session ID

```
trace transactions where (userid|url|url query|sessionid) (equals|not
equals|contains|not contains|starts with|ends with) [EXACT MATCH] in agents matching
[REGULAR EXPRESSION] for [EXACT MATCH] (second|seconds|sec|secs|s)
trace transactions where (userid|url|url query|sessionid) (exists|not exists) in
agents matching [REGULAR EXPRESSION] for [EXACT MATCH] (second|seconds|sec|secs|s)
quietly trace transactions where (userid|url|url query|sessionid) (equals|not
equals|contains|not contains|starts with|ends with) [EXACT MATCH] in agents matching
[REGULAR EXPRESSION] for [EXACT MATCH] (second|seconds|sec|secs|s)
quietly trace transactions where (userid|url|url query|sessionid) (exists|not
exists) in agents matching [REGULAR EXPRESSION] for [EXACT MATCH]
(second|seconds|sec|secs|s)
```

Filtering on request header, request parameter, or session attribute

```
trace transactions where (request header|request parameter|session attribute) [EXACT
MATCH] (equals|not equals|contains|not contains|starts with|ends with) [EXACT MATCH]
in agents matching [REGULAR EXPRESSION] for [EXACT MATCH] (second|seconds|sec|secs|s)
trace transactions where (request header|request parameter|session attribute) [EXACT
MATCH] (exists|not exists) in agents matching [REGULAR EXPRESSION] for [EXACT MATCH]
(second|seconds|sec|secs|s)
quietly trace transactions where (request header|request parameter|session
attribute) [EXACT MATCH] (equals|not equals|contains|not contains|starts with|ends
with) [EXACT MATCH] in agents matching [REGULAR EXPRESSION] for [EXACT MATCH]
(second|seconds|sec|secs|s)
quietly trace transactions where (request header|request parameter|session
attribute) [EXACT MATCH] (exists|not exists) in agents matching [REGULAR EXPRESSION]
for [EXACT MATCH] (second|seconds|sec|secs|s)
```

Filtering on errors

```
trace transactions with errors containing [EXACT MATCH] in agents matching [REGEX]
for [EXACT MATCH] (second|seconds|sec|secs|s)
quietly trace transactions with errors containing [EXACT MATCH] in agents matching
[REGEX] for [EXACT MATCH] (second|seconds|sec|secs|s)
```

Transaction Trace Queries from Transaction Events Database

Transaction Trace results are automatically stored in the Transaction Events Database. Use the CLW commands described here to retrieve that data.

CA Technologies Confidential/Proprietary Information. This Documentation is for your informational purposes only.

Get Historical Events Matching <VALUE>

This command retrieves a full event XML output file from the Transaction Events Database. A full event (as opposed to a summary event, below) is the full data available in the lower section of the Transaction Event Viewer.

```
java –Dintroscope.clw.tt.mode=Asynch -jar CLWorkstation.jar trace transactions get historical events matching "localhost"
```

Here is an example of the output:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<TransactionTracerSession Duration="78" Version="0.1"
EndDate="2005-06-20T17:40:24.146-08:00" StartDate="2005-06-20T17:40:24.068-08:00"
User="Admin">
<TransactionTrace Domain="SuperDomain" Duration="78" Process="UnknownProcess"
Host="maddleman-dt2" AgentName="Tomcat" StartDate="2005-06-20T17:40:24.068-08:00"
EndDate="2005-06-20T17:40:24.146-08:00">
<CalledComponent RelativeTimestamp="0" Duration="78" ComponentType="Browser"
ComponentName="Browser" MetricPath="Browser">
<CalledComponents>
<CalledComponent RelativeTimestamp="0" Duration="0" ComponentType="JNDI"
ComponentName="ProxyDirContext" MetricPath="JNDI|Context|ProxyDirContext">
<Parameters>
<Parameter Value="lookup" Name="Method"/>
</Parameters>
</CalledComponent>
<CalledComponent RelativeTimestamp="0" Duration="0" ComponentType="JNDI"
ComponentName="ProxyDirContext" MetricPath="JNDI|Context|ProxyDirContext">
<Parameters>
<Parameter Value="lookup" Name="Method"/>
</Parameters>
</CalledComponent>
<CalledComponent RelativeTimestamp="0" Duration="62" ComponentType="Servlets"
ComponentName="HelloWorldExample" MetricPath="Servlets|HelloWorldExample">
<CalledComponents>
<CalledComponent RelativeTimestamp="0" Duration="0" ComponentType="JNDI"
ComponentName="ProxyDirContext" MetricPath="JNDI|Context|ProxyDirContext">
<Parameters>
<Parameter Value="lookup" Name="Method"/>
</Parameters>
</CalledComponent>
<CalledComponent RelativeTimestamp="0" Duration="0" ComponentType="JNDI"
ComponentName="ProxyDirContext" MetricPath="JNDI|Context|ProxyDirContext">
<Parameters>
<Parameter Value="lookup" Name="Method"/>
</Parameters>
</CalledComponent>
</CalledComponents>
</TransactionTrace>
</TransactionTracerSession>
```

CA Technologies Prerelease Documentation. Authorized Use Only.

Get summaries matching <VALUE>

This command retrieves only the event summary XML from the Transaction Events database and outputs an XML file. An event summary is information available in the top pane of the Transaction Event Viewer. This command is useful to export data from the Transaction Events database for generating reports on the number of errors the system experiences by day, for example.

```
get historical event summaries matching [REGULAR EXPRESSION]
```

Here is an example of the output:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<TransactionTracerSession Duration="78" Version="0.1"
EndDate="2005-06-20T17:40:24.146-08:00" StartDate="2005-06-20T17:40:24.068-08:00"
User="Admin">
<TransactionTrace Domain="SuperDomain" Duration="78" Process="UnknownProcess"
Host="maddleman-dt2" AgentName="Tomcat" StartDate="2005-06-20T17:40:24.068-08:00"
EndDate="2005-06-20T17:40:24.146-08:00">
<CalledComponent RelativeTimestamp="0" Duration="78" ComponentType="Browser"
ComponentName="Browser" MetricPath="Browser">
<Parameters>
<Parameter Value="" Name="URL"/>
<Parameter Value="Normal" Name="Trace Type"/>
<Parameter Value="\20050620\17\1119314425302" Name="Trace ID"/>
<Parameter Value="" Name="User ID"/>
</Parameters>
</CalledComponent>
</TransactionTrace>
</TransactionTracerSession>
```

Application Triage Map Commands

Triage Map Alerts are specialized alerts associated with the AppMap and AppMap entities. Triage Map Alerts are located in the Triage Map Configurations Management Module in the *TriageMapConfigurationsManagementModule.jar* file.

These commands are available for this Management Module:

List Management Modules (see page 180)
Query the State of Management Modules (see page 180)
List Triage Map Alerts (see page 181)
Query the State of Management Modules (see page 182)
List Triage Map Entity Metric Groups (see page 184)
Query the State of Triage Map Entity Metric Groups (see page 185)

CA Technologies Confidential Document. Authorized Use Only.

List Management Modules

This command lists the name of the triage map Management Modules.

Syntax

list management modules matching <REGULAR EXPRESSION>

Example

list management modules matching (Triage Map.)*

This command lists Management Modules whose name begins with *Triage Map*.

Output

The name and state of each Management Module is listed. In this case Triage Map Configurations Management Modules matches the regular expression and it will be listed as follows:

```
Triage Map Configurations
```

Query the State of Management Modules

This command lists the statuses of the triage map Management Modules.

Syntax

query state of management modules matching [REGULAR EXPRESSION]

Example

query state of management modules matching (.)*

This command returns the name and state of all Management Modules.

CA Technologies Prerelease Documentation. Authorized Use Only

Output

Lists the name and state of each Management Module configured for your system. For example:

```
Supportability:Active
Trade Application Definitions:Active
Triage Map Configurations:Active
Browser Response Time Adapter:Active
Sample:Active
MOM Infrastructure Monitoring:Active
Reports Unit Tests:Active
Collector - 1:Active
ChangeDetector:Active
CEM BT Stats:Active
```

List Triage Map Alerts

This command lists the triage map alerts.

Syntax

list triage map alerts matching [REGULAR EXPRESSION] in management modules matching [REGULAR EXPRESSION]

Example

list triage map alerts matching (.) in management modules matching (Triage Map .*)*

This command will try to list all triage map alerts in the Management Module whose name begins with *Triage Map*.

CA Technologies Prerelease Documentation. Authorized Use Only.

Output

Lists the name of all the triage map alerts configured for your system. For example:

Triage Map Configurations.By Frontend|TradeService|Health:Average Response Time (ms)_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|loginAdmin.action|loginAdmin.action:Errors Per Interval_Summary_MetricAlert
Triage Map Configurations.By Frontend|ReportingService|Backend Calls|WebServices:Average Response Time (ms)_Location_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatients.action:Average Response Time (ms)_Location_MetricAlert
Triage Map Configurations.By Frontend|PatientFacadeServiceWebApp|Backend Calls|//localhost%9082/medrec(POINTBASE DB):Alert
Triage Map Configurations.By Frontend|ReportingEngine|Backend Calls|jdbc%hsqldb%file%report-records:Stall Count_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatients.action:Stall Count_Location_MetricAlert
Triage Map Configurations.Business Service|MedRecBS|Business Transactions|home.action:Alert
Triage Map Configurations.By Frontend|/medrec|Health:Errors Per Interval_Location_MetricAlert
Triage Map Configurations.Business Service|Trading Service|Business Transactions|View Orders:Total Defects Per Interval_Summary_MetricAlert
Triage Map Configurations.By Frontend|AuthenticationService|Health:Errors Per Interval_Summary_MetricAlert
Triage Map Configurations.By Business Service|Trading Service|Options Trading|Check Options:Concurrent Invocations_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|index.action|index.action:Stall Count_Location_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|index.action|index.action:Concurrent Invocations_Summary_MetricAlert
Triage Map Configurations.Business Service|MedRecBS|Business Transactions|viewPatients.action:Total Defects Per Interval_Summary_MetricAlert
Triage Map Configurations.By Business Service|MedRecBS|index.action|index.action:Average Response Time (ms)_Summary_MetricAlert

Query the State of Management Modules

This command lists the statuses of the triage map alerts.

CA Technologies Proprietary/Limited Use Only Documentation.

Syntax

query state of triage map alerts matching [REGULAR EXPRESSION] in management modules matching [REGULAR EXPRESSION]

Example

query state of triage map alerts matching (.) in management modules matching (Triage Map .*)*

This command lists all triage map alerts and their state in the Management Module whose name begins with *Triage Map*.

Output

Lists the name of all the triage map alerts along with the state configured for your system. For example:

```
Triage Map Configurations.By Frontend|TradeService|Health:Average Response Time
(ms)_Summary_MetricAlert:Active
Triage Map Configurations.By Business
Service|MedRecBS|loginAdmin.action|loginAdmin.action:Errors Per
Interval_Summary_MetricAlert:Active
Triage Map Configurations.By Frontend|ReportingService|Backend
Calls|WebServices:Average Response Time (ms)_Location_MetricAlert:Active
Triage Map Configurations.By Business
Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatien
ts.action:Average Response Time (ms)_Location_MetricAlert:Active
Triage Map Configurations.By Frontend|PatientFacadeServiceWebApp|Backend
Calls|//localhost%9082/medrec(POINTBASE DB):Alert:Active
Triage Map Configurations.By Frontend|ReportingEngine|Backend
Calls|jdbc%hsqldb%file%report-records:Stall Count_Summary_MetricAlert:Active
Triage Map Configurations.By Business
Service|MedRecBS|viewNewlyRegisteredPatients.action|viewNewlyRegisteredPatien
ts.action:Stall Count_Location_MetricAlert:Active
Triage Map Configurations.Business Service|MedRecBS|Business
Transactions|home.action:Alert:Active
Triage Map Configurations.By Frontend|/medrec|Health:Errors Per
Interval_Location_MetricAlert:Active
Triage Map Configurations.Business Service|Trading Service|Business
Transactions|View Orders:Total Defects Per Interval_Summary_MetricAlert:Active
Triage Map Configurations.By Frontend|AuthenticationService|Health:Errors Per
Interval_Summary_MetricAlert:Active
Triage Map Configurations.By Business Service|Trading Service|Options
Trading|Check Options:Concurrent Invocations_Summary_MetricAlert:Active
Triage Map Configurations.By Business
Service|MedRecBS|index.action|index.action:Stall
Count_Location_MetricAlert:Active
```

CA Technologies Confidential & Proprietary Documentation. For Internal Use Only.

List Triage Map Entity Metric Groups

This command lists the name of the triage map entity metric groups.

Syntax

list triage map entity metric groups matching [REGULAR EXPRESSION] in management modules matching [REGULAR EXPRESSION]

Example

list triage map entity metric groups matching (.) in management modules matching (Triage Map .*)*

This command lists all triage map entity groups in the Management Module whose name begins with *Triage Map*.

Output

Lists the name and state of all the triage map entity groups configured for your system. For example:

```
Triage Map Configurations.By Frontend|/medrec|Health:Average Response Time (ms)_Summary
Triage Map Configurations.By Frontend|AuthenticationEngine|Health:Average Response Time (ms)_Summary
Triage Map Configurations.Business Service|MedRecBS|Business Transactions|index.action:Total Defects Per Interval_Summary
Triage Map Configurations.By Business Service|Trading Service|View Orders|View my order:Concurrent Invocations_Location
Triage Map Configurations.By Frontend|PatientFacadeServiceWebApp|Health:Errors Per Interval_Summary
Triage Map Configurations.By Business Service|Trading Service|Transaction Summary|Request Transaction Summary:Average Response Time (ms)_Summary
Triage Map Configurations.By Frontend|ReportingEngine|Health:Errors Per Interval_Summary
Triage Map Configurations.By Business Service|Trading Service|Options Trading|Check Options:Errors Per Interval_Summary
Triage Map Configurations.By Frontend|AuthenticationEngine|Backend Calls|jdbc%hsqldb%file%customer-records:Average Response Time (ms)_Summary
Triage Map Configurations.By Frontend|OrderEngine|Health:Average Response Time (ms)_Location
Triage Map Configurations.Business Service|MedRecBS|Business Transactions|viewPatients.action:Total Transactions Per Interval_Summary
Triage Map Configurations.By Business Service|Trading Service|Options Trading|Check Options:Stall Count_Summary
```

CA Technologies Pre-release Documentation. Limited use Only.

Query the State of Triage Map Entity Metric Groups

This command lists the statuses of the triage map entity metric groups.

Syntax

query state of triage map entities matching [REGULAR EXPRESSION] in management modules matching [REGULAR EXPRESSION]

Example

query state of triage map entities matching (.*) in management modules matching (Triage Map .*)

This command will try to list all triage map entities and their state in the Management Module whose name begins with Triage Map.

Output

Lists the name of all the triage map entities configured for your system. For example:

```
Triage Map Configurations.By Frontend|OrderEngine|Backend Calls:Active
Triage Map Configurations.By Business
Service|MedRecBS|viewRecordSummary.action|kmuerxnjasdfgdfhjwerhgiusadfoiuoher
g:Active
Triage Map Configurations.By Frontend|/medrec|Backend Calls|System localhost on
port 9082:Active
Triage Map Configurations.By Business
Service|MedRecBS|viewPatient.action|viewPatient.action:Active
Triage Map Configurations.By Frontend|Apache-Axis:Active
Triage Map Configurations.By Business Service|Trading Service:Active
Triage Map Configurations.By Business Service|Trading Service|Balances:Active
Triage Map Configurations.By Frontend|AuthenticationEngine:Active
Triage Map Configurations.By Frontend|AuthenticationService|Backend
Calls|WebServices:Active
Triage Map Configurations.By Frontend|AuthenticationService|Health:Active
Triage Map Configurations.By Business
Service|MedRecBS|viewRecordSummary.action:Active
Triage Map Configurations.By Business
Service|MedRecBS|viewNewlyRegisteredPatients.action|kmuerxnjasdfgdfhjwerhgius
adfoiuoherg:Active
Triage Map Configurations.By Business Service|Trading Service|Place Order:Active
Triage Map Configurations.Business Service|Trading Service|Business
Transactions|Login:Active
Triage Map Configurations.Business Service|MedRecBS|Business
Transactions|index.action:Active
```

CA Technologies Pre-Release Documentation. Authorized Use Only.

Sample scripts

This section contains sample scripts that use CLW commands:

- Transaction Trace Session (see page 186)
- Transaction Trace session triggered by Alert (see page 190)

These scripts use other scripts—*setIscEnv.cmd* and *configureWily.cmd*—that are available on the Wily Community site at *https://community.wilytech.com/entry!default.jspa?categoryID=492&externalID=2103*. They are part of a Community Developed Tool named *Useful Introscope Scripts*, that can assist in startup, shutdown, restart, and using CLW. Download the appropriate ZIP file and follow the instructions to install the scripts:

1. Extract the appropriate file, *scriptsBin.unix.zip* or *scriptsBin.win.zip* into your Enterprise Manager installation. The package will create a *bin* subdirectory.
2. Go to the *bin* subdirectory, and invoke:

   .*configureScriptsEnv.cmd*

   or

   ./configureScriptsEnv.sh
3. This creates the file *setIscEnv.cmd* or *setIscEnv.sh*. This file contains all relevant Introscope folder locations from which the other scripts will read.

Additional sample scripts are available on the Wily Community site at http://community.wilytech.com.

CA Technologies Internal Release Document. Authorized Use Only.

Transaction Trace Session

This script initiates a Transaction Trace Session for all agents. At run-time, you supply two arguments:

- Execution threshold value—The minimum execution time, in seconds, for transactions to be traced.
- Session duration—The duration of the Transaction Trace Session, in seconds.

The script stores results in a date-oriented directory structure.

Name

clwTxnTracer

Syntax

clwTxnTracer <threshold> <duration>

where:

threshold is the transaction execution time threshold, in seconds

duration is duration of the session, in seconds

Example

clwTxnTracer 1 600

Causes Transactions Trace Sessions of 600 seconds duration to be run for any transaction that takes longer than 1 second to execute.

Results

The Transaction Trace data is written to a single file, in a directory that corresponds to current date:

<EM_Home>/traces\<year>\<month>\<day>\<agents>
_<threshold>secs_for<duration>secs.<start_time>.xml

For example:

C:\introscope\traces\2004\11\11\AllAgents_1secs_for600secs.
14091705.xml

```
@ECHO OFF
SETLOCAL
TITLE Introscope -- Transaction Tracer Session


REM ****
REM Sets the current working directory with quotes around it
REM This enables this file to execute correctly if it is
REM invoked somewhere other than its location.
REM ****
FOR /f "tokens=1 delims=" %%a in ('cd') DO SET PWD=%%a
```

CA Technologies Prerelease Documentation. Authorized Use Only.

```
REM ****
REM Set the command directory to the command argument itself,
REM namely the full path of this file or the relative path
REM to where the file was invoked.
SET CMD_DIR=%0
REM Now simply strip out the name of this file to get dir
REM !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
REM       Use current filename
SET CMD_DIR=%CMD_DIR:clwTxnTracer.cmd=%
REM !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
REM Strip the quotes
REM add quotes if they're not there
SET CMD_DIR="%CMD_DIR%"
REM if not add quotes, then this statement chokes
SET CMD_DIR=%CMD_DIR:"=%
REM if the current dir wasn't determined, just set it to working dir
IF NOT EXIST "%CMD_DIR%" SET CMD_DIR=.
IF "%CMD_DIR%" == %0 SET CMD_DIR=.
IF "%CMD_DIR%" == "%0" SET CMD_DIR=.



REM ****
REM Now set the home directory for Introscope, and others
IF NOT EXIST "%CMD_DIR%\setIscEnv.cmd" goto noSetEnvFile
call "%CMD_DIR%\setIscEnv.cmd"

SET TT_THRESHOLD_IN_SECS=%1
SET TT_DURATION_IN_SECS=%2
IF "X%TT_THRESHOLD_IN_SECS%" == "X" SET TT_THRESHOLD_IN_SECS=10000
IF "X%TT_DURATION_IN_SECS%" == "X" SET TT_DURATION_IN_SECS=600
IF %TT_THRESHOLD_IN_SECS% LEQ 0 GOTO badArg
IF %TT_DURATION_IN_SECS% LEQ 0 GOTO badArg
SET TT_AGENTS=(.*)

REM Java Settings (e.g. increase Max heap size )
SET JAVA_OPTS=-Xms64m -Xmx512m
SET JAVA_OPTS=%JAVA_OPTS% %CLW_USER_OPT% %CLW_PASSWORD_OPT% %CLW_EM_HOST_OPT%
%CLW_EM_PORT_OPT%
```

CA Technologies Internal/Prerelease Documentation.
Authorized Use Only.

```
SET COMMAND=trace transactions exceeding %TT_THRESHOLD_IN_SECS% seconds in agents
matching \"%TT_AGENTS%\" for %TT_DURATION_IN_SECS% seconds

SET AGENT_STR=%TT_AGENTS%
REM Remove the spaces, parentheses
SET AGENT_STR=%AGENT_STR:(=%
SET AGENT_STR=%AGENT_STR:)=%
SET AGENT_STR=%AGENT_STR: =%
if "X%TT_AGENTS%" == "X(.*)" SET AGENT_STR=AllAgents
FOR /f "tokens=1-4 delims=:." %%a in ("%time%") DO SET SECS_STR=%%a%%b%%c%%d
FOR /f "tokens=1-4 delims=/ " %%a in ("%date%") DO SET
TRACES_DIR=%ISC_TRACES%\%%d\%%b\%%c


REM Create the traces dir, if not exist
IF NOT EXIST "%TRACES_DIR%" MKDIR "%TRACES_DIR%"
SET
TT_FILE=%TRACES_DIR%\%AGENT_STR%_%TT_THRESHOLD_IN_SECS%secs_for%TT_DURATION_IN_SE
CS%secs.%SECS_STR%.xml

TITLE Introscope -- Transaction Tracer for %TT_DURATION_IN_SECS% secs with threshold
%TT_THRESHOLD_IN_SECS% secs
REM send stderr to null b/c there are meaningless disconnect excepts that might be
*scary-looking*
ECHO Saving session output to %TT_FILE%
"%ISC_JAVA_HOME%\bin\java" %JAVA_OPTS% -jar "%CLW_JAR%" %COMMAND% > "%TT_FILE%"
2>NUL
goto finish


:badArg
ECHO Arguments should be greater than zero.
ECHO Syntax: $0 "threshold in secs" "duration in secs"
goto finish

:noSetEnvFile
ECHO There is no environment setup file: setIscEnv.cmd
ECHO Please re-run configureWily.cmd in installation directory!
goto finish

:finish
ENDLOCAL
```

CA Technologies Prerelease Documentation. Authorized Use Only.

Transaction Trace session triggered by Alert

This script initiates a Transaction Trace Session as a result of an Alert being triggered. You use a Shell Command Action to invoke it, passing the Alert message text as an argument.

The script parses the Alert message text for the agent that reported the data that triggered the Alert, and the danger level value that was exceeded, which it uses as the agent and threshold values for the Transaction Trace Session. Session results are stored in a date-oriented directory structure.

Name

alert2TxnTracer

Syntax

For more information on creating shell command actions, see the *CA APM Workstation User Guide*.

Results

The Transaction Trace data is written to a single file, in a directory that corresponds to current date:

```
<EM_Home>/traces\<year>\<month>\<day>\<agents>
_<threshold>secs_for<duration>secs.<start_time>.xml
```

For example:

C:\introscope\\traces\2004\11\11\AllAgents_1secs_for600secs.14091705.xml

CA Technologies Prerelease Documentation
Authorized Use Only

Chapter 10: Using the WatchDog Command-Line Tool

For reliability and security, Enterprise Manager servers can be started and monitored by a separate WatchDog process. You can use the cross-platform WatchDog command-line tool to start, stop, manage, and restart a failed Enterprise Manager server.

This section contains the following topics:

About WatchDog (see page 191)
Running WatchDog (see page 192)
WatchDog commands (see page 193)
WatchDog command options (see page 197)
Viewing WatchDog log files (see page 202)
Sample scripts (see page 202)

About WatchDog

WatchDog is a cross-platform command-line tool that starts, stops, and manages an Enterprise Manager process. It also has the capability to restart a failed Enterprise Manager in support of the existing Enterprise Manager hot failover functionality.

The WatchDog command-line tool resides in the *<EM_Home>/bin* directory. You can invoke the commands in two ways—using the *WatchDog.bat* or *WatchDog.sh* helper or by running the Java application directly from *WatchDog.jar*, which allows you to specify more options, but is also more complex. For example, to start the Enterprise Manager with the WatchDog process in *watch* mode, type *WatchDog watch* from the bin directory.

Note: JRE 1.6 is required to use WatchDog.

CA Technologies Prerelease Documentation
Internal and Use Only.

Every WatchDog command returns control back to the command window or shell with a result code. This is because WatchDog is actually two programs:

- WatchDog manager—accepts lines from the command line, performs various actions, and returns the results.
- WatchDog server process—monitors the Enterprise Manager's health and automatically restarts a stopped Enterprise Manager.

Note: Because you can start the Enterprise Manager with mechanisms other than WatchDog, the *start* command cannot always know if the Enterprise Manager is in the process of starting. Because of this, WatchDog could, if invoked immediately after an Enterprise Manager is started without WatchDog's knowledge, attempt to start a second Enterprise Manager instance. The second attempt to start the Enterprise Manager would fail and an error result would not be returned, because the only command issued is *start.* An error, would be logged in the Enterprise Manager log file located in *<EM_Home>/logs*.

Eventually the first WatchDog *start* command attempt succeeds and WatchDog returns a *running* command.

Running WatchDog

Use a Windows command prompt or UNIX shell to view WatchDog result codes.

To run WatchDog using WatchDog.bat or WatchDog.sh

1. Open a Windows command prompt or UNIX shell.
2. Go to *<EM_Home>\bin* for Windows or *<EM_Home>/bin* for UNIX.
3. Type *WatchDog <command>*, where *<command>* is: *start*, *watch*, *stop*, *status*, or *help*. For more information about the commands, see WatchDog commands (see page 145).

Note: On z/OS only, in order to run WatchDog commands you must add a line to *runem.sh*, the script that start the Enterprise Manager in Z/OS environments. By default IHOME is blank. For example, if your Enterprise Manager is installed in the /u/users/wily/Introcoscope9.0.0 directory, include this line in the to *runem.sh* file: *IHOME=/u/users/wily/Introcoscope9.0.0; export IHOME*

For more information on runem.sh, see Starting and configuring the Enterprise Manager on z/OS (see page 39).

To run WatchDog by directly invoking WatchDog.jar:

1. Open a Windows command prompt or UNIX shell.
2. Go to *<EM_Home>\bin* for Windows or *<EM_Home>/bin* for UNIX.

CA Technologies Prerelease Documentation
Authorized Use Only.

3. Type these commands, depending on your operating system:

   - Windows:
     ..\jre\bin\java -jar WatchDog.jar command [-options]
   - UNIX:
     ../jre/bin/java -jar WatchDog.jar command [-options]

   For more information about the commands, see WatchDog commands (see page 145).

Running WatchDog on SAP

When using WatchDog on an Enterprise Manager running in an SAP environment, users are asked for administrator credentials before being able to use WatchDog.

To run WatchDog commands on an Enterprise Manager in an SAP environment:

- Supply your Enterprise Manager Admin userid and password to proceed.

WatchDog commands

WatchDog commands correspond to its four basic functions:

- Starting and restarting the Enterprise Manager process—see WatchDog start command (see page 194)
- Stopping the Enterprise Manager process—see WatchDog stop command (see page 195)
- Checking the Enterprise Manager's status—see WatchDog status command (see page 196)

Use one of two syntaxes when entering commands:

- Basic, which is simplified by the use of the *WatchDog.bat* for Windows or *WatchDog.sh* for UNIX helper scripts
- Full, which directly invokes the *WatchDog.jar* Java application and allows for all WatchDog options.

The WatchdDog is controlled on the command-line using java -jar WatchDog.jar command [-options].

CA Technologies Prerelease Documentation. Authorized Use Only

Note: On z/OS only, in order to run WatchDog commands you must add a line to *runem.sh*, the script that start the Enterprise Manager in Z/OS environments. By default IHOME is blank. For example, if your Enterprise Manager is installed in the /u/users/wily/Introcoscope9.0.0 directory, include this line in the to *runem.sh* file:

IHOME=/u/users/wily/Introcoscope9.0.0; export IHOME

For more information on runem.sh, see Starting and configuring the Enterprise Manager on z/OS (see page 39).

All commands return two result codes:

- one for the Enterprise Manager
- another for the WatchDog's auto-restart feature.

WatchDog result codes are preceded with *wd*.

WatchDog start command

The *start* command starts an Enterprise Manager process, but relies upon an existing command to launch the Enterprise Manager process.

You can start the Enterprise Manager using WatchDog with the automation restart of the WatchDog feature or using basic syntax. If you use basic syntax, it is not necessary to specify the *start* command, because it defaults to *Introscope Enterprise Manager.exe* for Windows or *Introscope_Enterprise_Manager* for UNIX. If you use the *watch* command, it puts the WatchDog server into *watch* mode, where it will periodically check to see if the Enterprise Manager has stopped running, and will restart it if it has to.

Use full syntax to type the *-watch* command option in conjunction with the *start* command, which tells puts WatchDog in *watch* mode.

Note: Because WatchDog is a synchronous process—that is, it returns a result without waiting for anything else—the *startcommandissued* and *stopcommandissued* are the only results possible after entering a *start* or *stop* command. These results mean only that the commands have been issued, not that the Enterprise Manager has actually started or stopped. To be certain, check the Enterprise Manager log, or check the operating system processes.

Start command syntax

This section gives you syntax for WatchDog start commands.

CA Technologies Pre-release Documentation. Authorized Use Only.

Basic syntax

WatchDog start

WatchDog watch

Full syntax

..\jre\bin\java -jar WatchDog.jar start -startcmd <cmd>

..\jre\bin\java -jar WatchDog.jar start -watch -startcmd <cmd>

Enterprise Manager return value

startcommandissued—the Enterprise Manager is not running.

alreadyrunning—the Enterprise Manager is already running.

starting—the Enterprise Manager is in the act of starting.

starterror—an error occurred when attempting to start the Enterprise Manager.

WatchDog return value

wdstartedwatching—WatchDog has entered watching mode.

wdalreadywatching—WatchDog is already running.

wdstartedsleeping—WatchDog server process is running, but it is not in watching mode.

wdalreadysleeping—WatchDog is already in sleep mode.

WatchDog stop command

The WatchDog stop command issues a *stop* command to the Enterprise Manager and WatchDog's server process. Because WatchDog may be running in *watch* mode, this command is the best way to ensure that the Enterprise Manager stops and remains stopped.

Stop command syntax

This section gives you syntax for WatchDog stop commands.

Basic syntax

WatchDog stop

CA Technologies Prerelease Documentation. Authorized Use Only

Full syntax

..\jre\bin\java -jar WatchDog.jar stop

Enterprise Manager return value

stopcommandisssued—the Enterprise Manager is running.

alreadystopped—the Enterprise Manager is not running.

stoperror—an error occurred when attempting to stop the Enterprise Manager.

WatchDog return value

wdstopped—WatchDog has been stopped.

wdstopcommandissued—WatchDog is running.

WatchDog status command

The *status* command reports the Enterprise Manager and WatchDog status.

Status command syntax

This section gives you syntax for WatchDog status commands.

Basic Syntax

WatchDog status

Full Syntax

..\jre\bin\java -jar WatchDog.jar status

Enterprise Manager return value

running—the Enterprise Manager is running.

stopped—the Enterprise Manager is not running.

starting—the Enterprise Manager is starting.

statusunknown—the status of the Enterprise Manager is unknown.

CA Technologies Prerelease Documentation. Authorized Use Only.

WatchDog return value

wdwatching—WatchDog is running.

wdsleeping—WatchDog server process is running (as it was when the WatchDog starts the Enterprise Manager), but it is not in watching mode.

wdstopped—WatchDog is not running.

WatchDog command options

WatchDog provides advanced options that are not available using basic syntax. If basic syntax is used, these options assume their default values.

The command options apply only to the *start* command, with the exception of the *-port* command option, which can be used for all commands.

Command options include:

- WatchDog -watch command option (see page 197)
- WatchDog -port command option (see page 198)
- WatchDog -startcmd command option (see page 199)
- WatchDog -interval command option (see page 199)
- WatchDog -startuptime command option (see page 200)
- WatchDog -jre command option (see page 200)

WatchDog -watch command option

The *-watch* command option can be used only with the *start* command (see WatchDog start command (see page 194)). It is the only way to enable the auto-restart feature in the Enterprise Manager.

You can turn the *watch* feature on or off after the Enterprise Manager has been started. For example, you can put the WatchDog server into *watch* mode if the Enterprise Manager is running without the WatchDog in *watch* mode by using the start command with the *-watch* option. Similarly, a *start* command without the *-watch* option can take the WatchDog out of *watch* mode, even if the Enterprise Manager is already running.

CA Technologies Proprietary Documentation. Authorized Use Only.

When you use the *-watch* option with the *start* command, it returns the following values:

- *startedwatching*—WatchDog was not watching before the command was entered. It is now monitoring.
- *alreadywatching*—WatchDog was already watching before the command was entered. It will continue monitoring.

When you do not use the *-watch* option with the *start* command, it returns the following values:

- *startedsleeping*—WatchDog was not sleeping before the command was entered. It has stopped monitoring.
- *alreadysleeping*—WatchDog was already sleeping before the command was entered. It will continue to refrain from monitoring.

-watch command option example:

```
<EM_Home>\bin>..\jre\bin\java -jar WatchDog.jar start -watch -startcmd
"..\Introscope Enterprise Manager.exe"
WatchMgr: iPort = 4321; watch = true; startcmd = ..\Introscope Enterprise Manager.exe;
startcommandissued & wdstartedwatching
```

WatchDog -port command option

Use the *-port* option to specify a value for the WatchDog's server port. If you do not enter a value, the default value 4321 is used.

When you specify a *-port* value for the first WatchDog start command, all subsequent commands will use the same value.

Note: Alternatively, you can run the WatchDog server on a non-default port, by modifying the *WatchDog.bat* file or creating a new batch file. This will ensure that the correct port is passed to WatchDog manager with every command.

-port command option example:

```
<EM_Home>\bin>..\jre\bin\java -jar WatchDog.jar
 start -port 5555 -startcmd "..\Introscope Enterprise Manager.exe"
WatchMgr:  iPort = 5555; watch = false; startcmd = ..\Introscope Enterprise
Manager.exe;
alreadyrunning & wdalreadysleeping
```

CA Technologies Prerelease Documentation
Authorized Use Only.

WatchDog -startcmd command option

Use this option to specify a command that starts the Enterprise Manager.

Note: This command option applies only to the *start* command. For more information, see WatchDog start command (see page 194).

WatchDog -interval command option

Use the *interval* option to specify the time in seconds that Enterprise Manager status checks occur to see if the Enterprise Manager has been stopped. The default value, which is used if this option is not specified, is 10 seconds.

This option applies when WatchDog's status is *watching*.

Note: This command option applies only to the *start* command. For more information, see WatchDog start command (see page 194).

-interval command option example:

```
C:\Program Files\CA Wily\Introscope9.0d140\bin>..\jre\bin\java -jar WatchDog.jar
start -watch -interval 300 -startcmd "..\Introscope Enterprise Manager.exe"
WatchMgr: iPort = 4321; watch = true; startcmd = ..\Introscope Enterprise Manager.exe;
startcommandissued & wdstartedwatching
```

CA Technologies Product Documentation. Authorized Use Only.

WatchDog -startuptime command option

Use this option to specify the time in seconds that WatchDog allows the Enterprise Manager to start. After WatchDog starts the Enterprise Manager, the amount of time specified for this option will pass and the Enterprise Manager's status will change to *starting*. From there, the normal mechanism configured to determine Enterprise Manager status will be put in place.

The Enterprise Manager status check is done immediately after an Enterprise Manager startup, so it will not give a *running* status. Instead, WatchDog's *status* command returns a *starting* status for a number of seconds after WatchDog starts the Enterprise Manager. This feature prevents WatchDog from accidentally starting two Enterprise Manager instances. The
-startuptime option allows for this delay to be configured to a value other than the default, which is 30 seconds.

Note: This command option applies only to the *start* command. For more information, see WatchDog start command (see page 194).

-startuptime command option example:

```
<EM_Home>\bin>..\jre\bin\java -jar WatchDog.jar  start -watch -startuptime 60
-startcmd "..\Introscope Enterprise Manager.exe"
WatchMgr: iPort = 4321; watch = true; startcmd = ..\Introscope Enterprise Manager.exe;
startcommandissued & wdstartedwatching
```

WatchDog -jre command option

Use this option to specify the JRE path if the JRE is not located in *jre/bin/java,* which is the default directory.

-jre command option example:

```
C:\Program Files\Java\jre1.6.0_07\bin\java" -jar WatchDog.jar start -watch -jre
"C:\Program Files\Java\jre1.6.0_07\bin\java"  -startcmd "..\Introscope Enterprise
Manager.exe"
startcommandissued & wdstartedwatching
```

CA Technologies Pre-release Documentation
Unauthorized Use Only.

WatchDog -emport command option

Use the *-emport* option to specify a value for the Enterprise Manager's server port when the Enterprise Manager is running on a port number other than its default port 5001. This command is required when the Enterprise Manager is running on a port number other than its default port 5001.

-emport command option example:

```
C:\Program Files\CA Wily\Introscope9.0\bin>..\jre\bin\java -jar WatchDog.jar
 start -emport 5003 -startcmd "..\Introscope Enterprise Manager.exe"
WatchMgr:  iPort = 5555; watch = false; startcmd = ..\Introscope Enterprise
Manager.exe;
alreadyrunning & wdalreadysleeping
```

WatchDog -emuser command option

Use the *-emuser* option when the Enterprise Manager has been configured with a user name other than its default of *Admin*.
This command is required when the Enterprise Manager has been configured with a user name other than its default of *Admin*.

In addition, use the *-emuser* option together with the *-empwd* option.

-emuser command option example:

```
<EM_Home>\bin>..\jre\bin\java -jar WatchDog.jar
 start -emuser EMUserName -empwd EMPassword -startcmd "..\Introscope Enterprise
Manager.exe"
WatchMgr:  iPort = 5555; watch = false; startcmd = ..\Introscope Enterprise
Manager.exe;
alreadyrunning & wdalreadysleeping
```

CA Technologies Internal / Documentation Authorized Use Only

WatchDog -empwd command option

Use the *-empwd* option when the Enterprise Manager has been configured with a password other than its default of "".
This command is required when the Enterprise Manager has been configured with a password other than its default of "".

In addition, use the *-empwd* option together with the *-emuser* option.

-empwd command option example:

```
<EM_Home>\bin>..\jre\bin\java -jar WatchDog.jar
 start -emuser EMUserName -empwd EMPassword -startcmd "..\Introscope Enterprise
Manager.exe"
WatchMgr:  iPort = 5555; watch = false; startcmd = ..\Introscope Enterprise
Manager.exe;
already running & wdalreadysleeping
```

Viewing WatchDog log files

WatchDogServer process log files are written to the *WatchDog.log* file.

To view WatchDog log files

1. Go to *<EM_Home>/logs.*
2. Open the *WatchDog.log* file.

Sample scripts

WatchDog.bat and *WatchDog.sh* implement basic syntax for Windows and UNIX platforms. These scripts are also useful as starting points to develop custom scripts that simplify full syntax of WatchDog. Use the *WatchDog.bat* and *WatchDog.sh* located in the *<EM_Home>/bin* directories of the Windows and UNIX installations.

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 11: Configuring WebView Options

This chapter describes how to configure optional behavior for the WebView application.

This section contains the following topics:

Configuring WebView for Non-ISO-8859-1 characters (see page 203)
Configuring the WebView application context path (see page 203)
Running WebView on z/OS (see page 204)

Configuring WebView for Non-ISO-8859-1 characters

To configure WebView to display non-ISO-8859-1 characters, such as Japanese:

1. In the *IntroscopeWebView.properties* file, set:
   `introscope.webview.svg.text.as.images=true`
2. In the *IntroscopeWebView.lax* file, add the *-Duser.language* and *-Duser.region* JVM options to the *lax.nl.java.option.additional* line. For instance, to configure WebView for Japanese characters, set the following JVM options:
   `-Duser.language=ja -Duser.region=JP`

Configuring the WebView application context path

You can optionally configure the application context path for the WebView installation. This path will be appended to the fully qualified hostname and port number of WebView, to form the URL for accessing WebView. This feature allows the user to define an "application name" for WebView.

To access WebView without a context path:

- Use this URL:

`http://<Server_name>:<Port_number>`

To access WebView when a context path is configured:

- Use this URL:

`http://<Server_name>:<Port_number>/<Context_path>/`

For example, if *Context_path* is defined as *webview*, and the *introscope.webview.tcp.port* property is set to the default value of *8080*, access WebView with this URL:

`http://<Server_name>:8080/webview/`

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

You configure a WebView context path using the *introscope.webview.context.path* property in the WebView properties file. For more information, see introscope.webview.context.path (see page 494). Restart WebView after defining the context path property.

Running WebView on z/OS

This section contains information and instructions for setting up Introscope WebView to run on z/OS.

Note: WebView runs *only* on z/OS 1.7 and later. To run WebView on z/OS, you must edit the WebView startup script to use a 1.5 JVM. See Editing runwv.sh to use a 1.5 JVM (see page 210).

You can run WebView on z/OS in one of these ways:

- as a USS shell script invoked from the command line
- as z/OS batch Job

Before you begin

Before you configure WebView, verify the prerequisites have been met.

Verify read/write access to configuration directory

Verify that WebView has both read access and write access to the <*EM_Home*>/*config* directory, so that any configuration changes can be properly saved.

Software prerequisites

The software prerequisites to run Java, and hence Introscope WebView, on z/OS are located in the Java for z/OS Program Directory GI10-0614.

These are the minimum software requirements:

- Introscope WebView requires Java 1.5 SR2 or later. Check your OS documentation to make sure it supports Java 1.5.
- UNIX System Services must be enabled.
- Language Environment 1.5 or higher.
- Java for z/OS.

CA Technologies Prerelease Documentation. Authorized Use Only.

Security and authorization

Authorization is determined by the RACF profile associated with a user's USERID. To successfully start the Introscope WebView, a user must have authority to access the following resources:

- UNIX System Services as specified by the OMVS parameter in the RACF profile
- EXEC PGM= BPXBATCH
- EXEC WVPROC

Virtual memory

If the USS/OMVS address space is not large enough, an out of memory condition will occur when an attempt is made to run WebView. The OMVS ASSIZEMAX (address-space-size) parameter in the RACF profile for the owner of the Introscope WebView job or started task must be large enough to execute the WebView. The maximum value for this parameter is 2G.

Copying and customizing WebView files

Before WebView can run, the following files (extracted from the z/OS archive) must be copied into the z/OS environment and customized to ensure compatibility with site-specific requirements:

- *IntroscopeWebView.properties*—the Introscope WebView properties file, which specifies information pertinent to the operation of the Introscope WebView
- *runwv.sh*—the shell script used to launch the Introscope WebView process
- *WVPROC*—the z/OS PROC used to execute z/OS BPXBATCH program that launches the shell script

IntroscopeWebView.properties and *runwv.sh* are UNIX files that must be copied into the Hierarchical File System (HFS) on the z/OS platform. These files will be used by UNIX System Services and Introscope WebView, respectively. Required changes to these files can be made using a Telnet-based editor such as vi. Using the USS/OMVS editor can be difficult, because it does not easily handle insertion of characters that increase line length beyond 72 characters.

Configure WebView for EBCDIC

Enable WebView to write out its messages and logs in EBCDIC (CP-1047) format by adding these properties to the *log4j.appender* section of the *IntroscopeWebView.properties* file:

```
log4j.appender.console.encoding=IBM-1047
log4j.appender.console=com.wily.org.apache.log4j.ConsoleAppender
log4j.appender.logfile.encoding=IBM-1047
log4j.appender.logfile=com.wily.org.apache.log4j.RollingFileAppender
```

CA Technologies Pre-release Documentation. Authorized Use Only.

Configuring WebView to run as a z/OS batch job

For ease of support in a z/OS environment, it is desirable to run WebView in a manner familiar to z/OS operations and development staff—either as a z/OS batch job or a started task. You can also run WebView as a shell process if desired. This section describes running the Introscope WebView as a batch job or started task.

Overview

The infrastructure required to run the Introscope WebView as an MVS batch job or started task is straightforward. The MVS utility program, BPXBATCH, enables users to invoke UNIX System Services to run shell commands, shell scripts and executable files in MVS batch. Specifically, WebView can be started using the BPXBATCH utility.

Copy WebView z/OS files to a z/OS Partitioned Dataset (PDS)

To copy WVPROC into the PDS:

1. Using ISPF, create a partition dataset named *userid*.JCL.CNTL, (where userid is the userid assigned to CA APM) with fixed block (FBA) 80 byte records.
2. Copy WVPROC into the PDS:

   userid.JCL.CNTL

   where *userid* is the userid assigned to CA APM and userid.*JCL.CNTL* is the name of the previously created PDS.

Note: Do not try to copy the shell script, *runwv.sh*, or the properties file, *IntroscopeWebView.properties*, into this PDS. These are HFS files, and have lines that are longer than 80 bytes. An attempt to copy them will either fail or lines will be truncated.

CA Technologies Prerelease Documentation.
Authorized Use Only

Customize the WebView z/OS PROC that invokes BPXBATCH

To customize the WebView z/OS process:

1. Edit the Introscope WVPROC using the TSO ISPF editor:
   - On the ISPF command line, type CAPS OFF and press Enter.
   - On the ISPF command line, type NUM OFF and press Enter.
   - Set the IHOME parameter equal to the fully qualified path to the directory that contains the runwv.sh.

   The following statement is coded so that the batch job and the UNIX shell script will end within 20 seconds:

   // PARM='SH nohup &IHOME/runwv.sh &IHOME & sleep 20'

   In the unlikely event that the z/OS is too busy to complete WebView initiation within 20 seconds, WebView will not be started. If you encounter this problem, try increasing sleep time.
2. Modify the STDERR and STDOUT file names contained in the following statements to satisfy installation naming conventions:

```
//STDOUT DD PATH='/<path>/stdout',
//STDERR DD PATH='/<path>/stderr',
```

Copy WVPROC into an active PROCLIB

Using ISPF option 3, copy the PROC to xxxx.PROCLIB where xxxx represents a site-dependent high level qualifier. If WVPROC is not an acceptable name in your installation, rename the PROC accordingly.

Running as a z/OS batch job

It is important to understand how the user is identified in various contexts:

- Batch—When a user submits a batch job through the TSO/E or ISPF interface, the user's TSO UserID and password are propagated to the batch job. Provided the submitted Batch JOB CARD does not explicitly contain a User and Password, the batch job acquires the same access authority as the TSO user. If User and Password are coded on the Batch JOB CARD, they override the UserID and Password of the TSO/E user.
- Started Task—If Introscope WebView is initiated as a started task, it runs under the authorization of the STC user ID. A Site-specific algorithm may be used to relate the name of the started PROC to a particular UserID. You must make sure that the algorithm produces a UserID with the correct authorization to run the Enterprise Manager.

CA Technologies Pre-release Documentation
Authorized Use Only

Running as a started task from the z/OS Console or SDSF

You may run the Introscope WebView as an MVS started task (STC). A started task is a familiar, administrator-friendly environment that allows the Enterprise Manager to be started, monitored and cancelled from the MVS console. In this case, WebView runs under the authorization of the STC user ID, thereby allowing assignment of specific authorities to the started task.

To run WebView as a started task, type the following command on the MVS console or on the command line of SDSF:

```
/start WVPROC
```

You must type the member name of the PROC stored in an active PROCLIB. If you were required to change the name of WVPROC, you must use the new name in the */start* command.

Running as a batch job

To run WebView as a batch job, submit the following JCL:

```
//jobname JOB installation jobcard parameters,NOTIFY=&SYSUID
ISCOPSTP EXEC WVPROC
/*
```

You must execute the member name of the PROC stored in an active PROCLIB. If you were required to change the name WVPROC, you must use the new name on the EXEC statement.

The z/OS batch job and the OMVS Shell will end shortly after WebView processes have started. If you submitted the batch JCL from your ISPF session and your JOB card contains NOTIFY=&SYSUID, your ISPF session will receive notification when your batch jobs ends.

CA Technologies Prerelease Documentation
Authorized Use Only

Verifying that the WebView job is running

The most accurate way to determine that the Introscope WebView has started correctly is to view the STDOUT file. To view STDOUT, you can telnet into the z/OS and view the STDOUT file defined in the WVPROC. The last line of the file should indicate that the Introscope WebView has started.

You can also determine that WebView has started from SDSF by following one of the procedures below:

- If you started WebView as a started task, on the DA screen you should see at least two active started tasks whose owner is the owner associated with the WebView PROC.
- If you started e WebView by submitting a batch job from your TSO session, look for two active tasks with your TSO USERID as the owner.

You may also see the job that invoked the WVPROC, in addition to the OMVS shell script. These two entries will disappear from the list once the WebView job has started.

The owner of your TSO session is also your TSO/E USERID. Do not mistake your TSO/E session for one of the Introscope WebView tasks.

Starting the Introscope WebView job automatically

To start the Introscope WebView automatically when JES2 starts:

1. Place the following statement in JES2PARM:

   $VS,'start WVPROC'

2. You must type the member name of the PROC stored in an active PROCLIB. If you were required to change the name of WVPROC, you must use the new name in the */start command*.

   You can also run Introscope WebView as a batch job using whatever job scheduling package you use to manage your batch workload.

CA Technologies Prerelease Documentation. Authorized Use Only.

Running WebView from a shell script

To start WebView as a USS/OMVS process from a shell script, use the *runwv.sh* file.

Examples of running WebView from the USS command line are:

runwv.sh &
nohup ./runwv.sh &

<WebView_Home> is the root of your WebView installation. If you run the script from a directory other than *<WebView_Home>*, identify the home directory as shown below:

nohup ./runwv.sh /<Actual_WebView_Home_Directory_Path> &

Editing runwv.sh to use a 1.5 JVM

To run WebView on z/OS, you must change the default Java Virtual Machine in the *runvw.sh* file to enable WebView to run on a 1.5 JVM.

To edit runsv.sh:

1. Open a command prompt and browse to the Introscope root directory.
2. In a text editor, open the file runwv.sh
3. Locate this line, which sets the location of your Java VM:

   JAVAPATH=/usr/lpp/java/IBM/J1.5; export JAVAPATH
4. Edit the line to point the JAVAPATH variable to the root directory of a 1.5 JVM.

   For example:

   JAVAPATH=/usr/lpp/java/IBM/J1.5; export JAVAPATH

   If you prefer, you can comment out the line and add a new line instead of editing it.
5. Save your changes and exit the text editor.
6. Use the edited *runwv.sh* script to start WebView.

Writing status information to log file

As WebView starts and runs, it writes status information to a log, which, by default, appears in the command prompt window (if started normally) and to the nohup.out file, if started in nohup mode.

CA Technologies Pre-release Documentation. Authorized Use Only.

Chapter 12: Configuring the Workstation

This section describes how to configure optional behavior for the Introscope Workstation.

This section contains the following topics:

Running Workstation in verbose mode (see page 211)
Redirecting Workstation output to a file (see page 211)
Configuring Workstation to supply login values (see page 212)
Configuring the Workstation for Asian-language reports (see page 212)

Running Workstation in verbose mode

Running the Workstation in verbose mode creates detailed log messages that are helpful for debugging or troubleshooting.

To run the Workstation in verbose mode:

1. Open the file *IntroscopeWorkstation.properties* in *<EM_Home>/config*.
2. In the property *log4j.logger.Workstation*, replace *"INFO"* with the following:

   VERBOSE#com.wily.util.feedback.Log4JSeverityLevel
3. Save changes.

Redirecting Workstation output to a file

To redirect Workstation output to a file:

1. Open the file, *IntroscopeWorkstation.properties*, in *<EM_Home>/config.*
2. In the *log4j.logger.Workstation* property, replace *"console"* with *"logfile."* For example, this property setting causes the Workstation to log verbose messages to a logfile:
   `log4j.logger.Workstation=VERBOSE#com.wily.util.feedback.Log4JSeverityLevel,logfile`
3. You can change the change the name and location of the Workstation logfile using the *log4j.appender.logfile.File* property.

CA Technologies Pre-release Documentation. Authorized Use Only

Configuring Workstation to supply login values

You can modify the *Introscope Workstation.lax* file to supply login values and bypass the login screen.

To set up command line arguments to bypass the Workstation login screen:

1. Open the file, *<EM_Home>/Introscope Workstation.lax*.
2. In the *lax.command.line.args* property, add *-login* commands for each login command to which to supply a value. For example, using local authentication, this property might look like this:
   ```
   lax.command.line.args=$CMD_LINE_ARGUMENTS$ -loginimmediate
   -loginhost foos -loginport 4503 -loginresponse sanderson,45tst87
   ```

Note: Do not delete the default value, *$CMD_LINE_ARGUMENTS$* from the setting.

If the login is successful (the correct information was supplied in the commands using the correct syntax), you are logged into the Workstation and a Console opens.

If login fails, the Workstation does not start, and an error message is logged.

Configuring the Workstation for Asian-language reports

If you plan to generate Introscope reports in languages that use multibyte character sets such as Chinese and Japanese, follow the instructions here to add components to your Workstation installation.

To generate Asian-language reports in RTF and HTML formats:

- Install the Supplemental Language Support to East Asian Languages in Windows. See http://www.microsoft.com/globaldev/handson/user/xpintlsupp.mspx.

To generate Asian-language reports in PDF formats:

- Download and install the Acrobat Reader Asian Font Pack from http://www.adobe.com/products/acrobat/acrrasianfontpack.html.

CA Technologies Pre-Release Documentation. Authorized Use Only

Chapter 13: System Configurations

This section describes overall CA APM system configurations.

This section contains the following topics:

Configuring the clock skew tolerance (see page 213)
Setting the Maximum Data Point Retrieval Property (see page 215)
Limiting the data points returned by a query (see page 216)
Increasing the number of query results returned (see page 217)
Configure APM Status Console Clamps and Important Events (see page 217)

Configuring the clock skew tolerance

The clock skew tolerance is the variation of the clock setting adjustments that are typically done in small increments. For example, setting a slightly fast running computer clock back a few seconds account for the small alterations you see in the logs. The variations in this delay cause the clock to timestamp the same metric at different times, which cause the system to get two metric values for the overlapping time period.

Turning back time is something that can happen both explicitly (the operator changes the clock time), or implicitly via automatic clock synchronizations, which is likely to occur with large installations.

Enterprise Manager addresses clock skew tolerance one of two ways:

- For large clock adjustments the system issues an error and exits.
- For small backwards clock adjustments, a warning is issued and the harvest cycle is stopped until the timestamp is as far along as it was previously.

  The value of this threshold is set at a default of 60 seconds, but can be configured. A negative clock adjustment of less than seconds results in a warning and temporary harvest suspension, while anything greater results in a system exit.

CA Technologies Proprietary / Confidential Documentation. Authorized Use Only.

If the Enterprise Manager detects that your system clock has been set back, which is a common occurrence when a fast CPU clock is reset by a network time service, the Enterprise Manager delays the harvest cycle until the clock catches up to its previous time. This delay is usually accompanied by a warning message. If you want to suppress warning messages, you can set the optional *introscope.enterprisemanager.clock.skewsuppresswarn* property. The value of this property is the number of milliseconds the metric harvest will be delayed before it is considered serious enough to be logged.

To configure the clock skew tolerance threshold:

1. Stop the Enterprise Manager.
2. Go to *<EM_Home>\config* directory and open the *IntroscopeEnterpriseManager.properties* file.
3. Add the following property, where <XXX> is a number in seconds:
   introscope.enterprisemanager.clock.skewtoleranceseconds=<XXX>
4. Optional: Set the *introscope.enterprisemanager.clock.skewsuppresswarn* property. The default value expressed in milliseconds is *0*.

   Note: This value cannot be set to more than the introscope.enterprisemanager.clock.skewtoleranceseconds property, or it will suppress the warnings completely and result in a system exit.
5. Start the Enterprise Manager, which will apply the new settings.

CA Technologies Prerelease Documentation. Authorized Use Only.

Setting the Maximum Data Point Retrieval Property

You can set the maximum data point retrieval property as a performance measure for JDBC queries. A data point is the value of one metric for one time period.

This Enterprise Manager clamp property limits the maximum number of metric data points any Enterprise Manager (either a clustered Collector or standalone Enterprise Manager) retrieves from SmartStor for any single query. However, the Enterprise Manager will always return the full data for at least one metric. If this property is not set, it could cause an Enterprise Manager overload, which will cause your system to perform slowly and could potentially cause an out of memory error. Additionally, if not set, this property may prompt the Command Line Workstation or Workstation to extract too much data, which in turn may cause the Enterprise Manager to stop responding.

To set the maximum data point retrieval property:

1. Go to *<EM_Home>\config* directory and open the *IntroscopeEnterpriseManager.properties* file.
2. Add this property:

   introscope.enterprisemanager.query.datapointlimit=100000

   where *100000* limits the data points retrieved from SmartStor per query at approximately 100,000. If no numeric value is entered, no limitation is applied.
3. Save the *IntroscopeEnterpriseManager.properties* file.

   The setting is applied when the Enterprise Manager checks the property file during a hot deploy, which occurs about once every 60 seconds.

CA Technologies Internal/Restricted Use Document
Authorized Use Only.

Limiting the data points returned by a query

You can set the *introscope.enterprisemanager.query.returneddatapointlimit* property as a JDBC query performance measure. This Enterprise Manager clamp property limits the number of data points returned by a query. A data point is the value of one metric for one time period.

The number of data points returned by a query depends on the number of matching metrics, the time range you are requesting, and the resolution of the SmartStor data. For example, suppose you're querying for 1 hour of data for a single metric with a resolution of 1 minute. If the data exists in SmartStor with a resolution of 15 seconds (this will depend on the configuration of the SmartStor data tiers), then the query matches 240 data points. However, because you are asking for 1 minute resolution, the 15-second data points are aggregated into 1-minute data points, so the query only returns 60 data points.

For a given metric, the query returns either the full set of data requested or nothing. You never see partial data for a metric: it's either all or nothing. The effect of the clamp is to reduce the number of metrics for which data is returned.

If you have a memory scarce environment and this property is not set, large data queries could cause an Enterprise Manager out of memory condition.

To limit the data points returned by a query:

1. Go to *<EM_Home>\config* directory and open the *IntroscopeEnterpriseManager.properties* file.
2. Add this property:

   introscope.enterprisemanager.query.returneddatapointlimit= 100000

   where 100000 limits the data points retrieved from SmartStor per query at approximately 100000. If no numeric value is entered, no limitation is applied.
3. Save the *IntroscopeEnterpriseManager.properties* file.

   The setting is applied when the Enterprise Manager checks the property file during a hot deploy, which occurs about once every 60 seconds.

CA Technologies Internal Document
Authorized Use Only

Increasing the number of query results returned

When *RangeQuery*, *PrefixQuery*, *WildcardQuery*, or *FuzzyQuery* queries are performed they are expanded before the search is performed. For example, if the indexed documents contain the terms "car" and "cars" the query "ca*" will be expanded to "car OR cars" before the search takes place. This might cause a system error if this number exceeds the default value of the term threshold, which is 2048. If this occurs, a *TooManyClauses* exception will be returned.

To increase the number of query results returned:

1. Go to *<EM_Home>\config* directory and open the *IntroscopeEnterpriseManager.properties* file.
2. Add the following property with a higher value then its default value of 2048:
   introscope.enterprisemanager.lucene.search.BooleanQuery.SetMaxClauseCount

   The setting is applied when the Enterprise Manager checks the property file during a hot deploy, which occurs about once every 60 seconds.

Configure APM Status Console Clamps and Important Events

You can configure threshold values for CA Introscope clamp, CA CEM clamp, and Important Event properties in the apm-events-thresholds.xml file. When these property thresholds are met or exceeded, the Enterprise Manager notifies the APM Status console. Clamp information displays in the Active Clamps tab. Important Events information displays in the Important Events table view.

Follow these steps:

1. Go to *<EM_Home>/config* directory and back up the apm-events-thresholds.xml file.

   Important: Do not disregard backing up the file. If a configured file gets corrupted, the property values revert to their defaults and all your configurations are lost.
2. Configure the properties as appropriate for your environment.

   Note: Configure *only* the values in the *threshold value* fields. Do *not* change any other fields such as clamp ID, importantEvent ID, and description because CA Introscope uses these fields for system monitoring.

   a. Open the apm-events-thresholds.xml file.

   b. Modify the clamp properties.

   c. Modify the important event properties.
3. Save and close apm-events-thresholds.xml.

   The Enterprise Manager applies your configuration settings during a hot deploy, which occurs about once every 60 seconds.

CA Technologies Pre-release Documentation. Authorized Use Only.

If the Enterprise Manager finds a syntax error in the apm-events-thresholds.xml file, it

- Notifies the APM Status console to display an apm.events.thresholds.config.error important event in the Important Events table view.
- Reverts all the clamp ?and Important Event? property values to their default values.

4. If necessary, troubleshoot the apm-events-thresholds.xml file.
   a. Open the apm-events-thresholds.xml file.
   b. Review the file for syntax errors and make any needed corrections.
   c. Save and close the file.
   d. If you cannot find or fix the apm-events-thresholds.xml file error, replace the corrupt apm-events-thresholds.xml file with the backup file you saved in Step 1.
   e. Repeat steps 2 and 3.
5. If you configuring Enterprise Managers across a cluster, perform steps 1 through 4 on all Collectors, the MOM, and CDV, if deployed.

   All the Enterprise Managers are configured with CA Introscope and CA CEM clamp and Important Event properties. The apm-events-thresholds.xml file is deployed and clamp and Important Event clamps activate when property threshold levels are met or exceeded.

CA Technologies Pre-Release Documentation. Authorized Use Only.

Chapter 14: Configuring the Application Triage Map

This chapter describes how to configure optional behavior for the application triage map.

This section contains the following topics:

The application triage map (see page 219)
Application triage map data clamping (see page 223)
Application triage map data pruning (see page 224)

The application triage map

The application triage map presents a graphical visualization of the components that make up your application, showing application health and errors. This map is automatically generated from performance and analysis of Introscope metrics, errors, and events. It presents applications in the business-centric terms that you've defined. The application triage map enables you to instantly grasp the structure of and relationships between the applications in your environment in a visual manner to help you identify and triage current and emerging problems.

For more information about the appearance and use of the application triage map, see the *CA APM Workstation User Guide*.

Introscope agents detect business transactions and frontends and detect their dependencies. Agents then report this information as data to the Enterprise Manager, which stores the data in the APM database and sends the data to the Workstation. The Workstation displays the data as the application triage map in the Map tab.

You can configure:

- whether or not agents collect application triage map data.
- the rate at which Introscope samples frontend or business service data for the purpose of drawing the map. See Transaction sampling (see page 220).
- how long application triage map data is stored in the APM database. See Map element aging (see page 220).
- whether or not, and how often, application triage map data is pruned from the APM database. See Application triage map data pruning (see page 224).

CA Technologies Proprietary Documentation. Authorized Use Only.

Transaction sampling

Introscope uses a technique known as *transaction sampling* to provide data for the application triage map display. After an agent is started, Introscope takes (by default) the first hundred transactions to draw an initial map of a frontend or business transaction and its dependencies. Subsequently, it gradually reduces the number of transactions it samples until it reaches a level where only every tenth transaction provides continuing data support to the map. After each ten thousand transactions, this pattern repeats.

You can adjust the sampling rate by editing these properties in the *IntroscopeAgent.profile*:

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

For information about editing properties in *IntroscopeAgent.profile*, see the *CA APM Java Agent Implementation Guide* or the *CA APM .NET Agent Implementation Guide* (depending on your environment).

Be aware that adjusting these properties to increase the number of transactions supporting the map display may affect overhead.

Map element aging

As data about the connection between applications ages out or is pruned from the APM database, agents rediscover frontends and business services and their dependencies, which the Workstation displays on the application triage map.

If you enable agents to collect application triage map data, then you can configure how often agents perform this data collection. For example, you can configure agents to collect application triage map data once every 24 hours if there are frequent changes to your organization's environment, or once a week if there are infrequent changes.

You can also configure the time period over which the Enterprise Manager delivers data to the Workstation for map display. For example, you can configure application triage map data to be displayed for a week or a month after being discovered by an agent. The time period is based on two map time range properties: *aging time* and *window time*.

CA Technologies Prerelease Documentation. Authorized Use Only

Aging is a characteristic of application triage map elements. When Introscope no longer receives data for a frontend or business transaction, the map element representing that frontend or business transaction is said to be aging. The *aging time* is the amount of time during which Workstation will continue to display the element in the application triage map while it is aging. When this time has expired, the element is removed from the agent cache and the element appears in the application triage map as aged. (See the *CA APM Workstation User Guide* for a description of the appearance of aged elements.) The *aging time* is set in the *introscope.apm.data.agingTime* property. The default is one day.

The *time window* is the amount of time the Workstation will display an aging element before the element disappears from the map altogether. This is set in the *introscope.apm.data.timeWindow* property. The default is three days.

For example, with default settings:

ApplicationA appears in the application triage map as a frontend, with dependency ApplicationB. While agents continue to report data for ApplicationB, ApplicationB appears in the map display normally. When data ceases to be reported from ApplicationB, the clock starts on both the aging time and the window time. After one day, the map node for ApplicationB will appear as aged. After three days, the ApplicationB map node will disappear from the map.

Things to notice:

- The timeWindow *includes* the period of agingTime.
- If you set timeWindow and agingTime to the same duration, map elements will not be displayed as aged, but will simply disappear at the end of the timeWindow.
- If you set timeWindow to a value greater than agingTime, aged map elements will disappear at the end of the agingTime period.

Tip: When deciding on what value to set for agingTime and timeWindow, choose a value based on how often your organization adds and modifies applications. If you set lower values, then the map displays application dependencies on a more real-time basis, but this can use more agent and Enterprise Manager overhead. If you set higher values, then the map displays application dependencies that can be less current, but this can use less agent and Enterprise Manager overhead. For information about agent and Enterprise Manager overhead, see the *CA APM Sizing and Performance Guide.*

The explanation in this section to this point refers to the Live view of the application triage map. See also How aging-related properties affect the historical view (see page 223).

CA Technologies Prerelease Documentation. Authorized Use Only.

To configure application triage map data aging properties:

1. Open the *IntroscopeEnterpriseManager.properties* file in the *<EM_Home>/config* directory.
2. If you are changing the value for *introscope.apm.data.agingTime:*
   a. Stop the agents.
   b. Shut down the Enterprise Manager.

Important! If you change the value for the *introscope.apm.data.agingTime* property, see Implications of reducing the application triage map aging time setting (see page 222).

3. Configure the *introscope.apm.data.agingTime* property. The default is 1 day. The value cannot be less than 1 minute.
4. Configure the *introscope.apm.data.timeWindow* property. The default is 3 days. The value cannot be less than 1 minute.
5. Save and close the *IntroscopeEnterpriseManager.properties* file.
6. If you changed the *introscope.apm.data.agingTime* property, restart the agents.
7. Restart the Enterprise Manager.

   When the Enterprise Manager starts, you can go to *<EM_Home>/logs* to review the Enterprise Manager log file messages associated with the map data aging properties. For example:

```
11/23/09 12:02:13 PM PST [INFO] [Manager] [ The APM Data  property :
introscope.apm.data.agingTime set to value : [ PropertyWithTimeUnit : 1 Min]]
```

Implications of reducing the application triage map aging time setting

If you plan to reduce the value of the property *introscope.apm.data.agingTime* in the *IntroscopeEnterpriseManager.properties* file, CA Technologies recommends you clean the application triage map data to avoid unexpected effects on historical data. Reducing the aging time means a smaller data sample appears in the map display, and this may result in a midleading display based on data in the APM database, because of the way agents report data to the Enterprise Manager. Cleaning the data from the database prevents this problem.

To clean the data, you must use a PostgreSQL client. Such a client can be downloaded from the PostgreSQL site at http://www.pgadmin.org.

To use the PostgreSQL client to clean the database:

- Execute the following SQL statement:

  ```
  truncate apm_edge;
  ```

  Increasing the value of the property *introscope.apm.data.agingTime* has no such unexpected effects. Cleaning the database is only necessary if you want to decrease the value of the property.

CA Technologies Pre-Release Documentation. Unauthorized Use Only

How aging-related properties affect the historical view

When viewing historical data with the application triage map, keep in mind these concepts:

- Regardless of how you set the beginning and end points of your historical time period, the data displayed on the application triage map will always span at least the period of time set in the *introscope.apm.data.timeWindow* property. For example:
  - If you set the beginning and end dates of the historical period to a 10-day span, but *introscope.apm.data.timeWindow* is set to 3 days, the map will display data for the entire 10-day span.
  - If you set the beginning and end dates of the historical period to a 2-day span, but *introscope.apm.data.timeWindow* is set to 3 days, the map will still display 3 days worth of data, with the 3-day period ending at the same time as the ending of the 2-day period. Whatever the *introscope.apm.data.timeWindow* property is set to is the minimum time period for what the map displays, even if the manually set historical period is less.
- The aging behavior of map elements will follow the period of time set in the *introscope.apm.data.agingTime* property as if that aging time occurs at the end of the historical window. For example:
  - If *introscope.apm.data.agingTime* is set to 1 day, and your historical time period is set to 10 days, the map element will be displayed as aged if no agent reported this element for the final (1) day of the period.

Application triage map data clamping

To control data flow to the application triage map for the purpose of assuring good performance, Introscope uses a threshold, set using the property introscope.apm.query.max.results (found in *APMEnterpriseManager.properties*), to clamp the amount of data the Workstation will attempt to display in the map. When you select a tree node or attempt to unroll a map element which would exceed this threshold, you see a message stating "The map is too large to display." You can adjust this threshold by adjusting the value of the property.

Note: Adjusting the value for the introscope.apm.query.max.results property may affect Workstation performance.

For more information about this property, see introscope.apm.query.max.results (see page 447).

CA Technologies Proprietary Documentation. Authorized Use Only.

To edit the application triage map display threshold:

1. Open the *APMEnterpriseManager.properties* file in the *<EM_Home>/config* directory.
2. Shut down the Enterprise Manager.
3. Find the introscope.apm.query.max.results property and adjust its value.
4. Save and close the *APMEnterpriseManager.properties* file.
5. Restart the Enterprise Manager.

Application triage map data pruning

The application triage map data in your APM database can grow over time. Excessive data in the APM database can affect Enterprise Manager retrieval of application triage map data as well as Workstation and Enterprise Manager performance. You can prune application triage map data to avoid collecting excessive map data in the APM database.

Application triage map data pruning is enabled by default to be done on the APM database in a cluster environment for standalone Enterprise Managers and MOMs.

Important! You cannot configure a Collector for application triage data pruning. If you set the *introscope.apm.pruning.enabled* property to *true* for a Collector, errors are returned.

You can configure how long map data remains in the database until it is pruned and how often you want to scheduling the pruning to occur. Map data pruning is controlled by a scheduler, which can be configured to prune data, on a daily, weekly, or monthly basis.

If application triage map data is pruned while a user is viewing the map, the pruned data remains on the map until the user changes the map time range, which triggers the map to be redrawn.

To configure map data pruning:

1. Open the *IntroscopeEnterpriseManager.properties* file in the *<EM_Home>/config* directory.
2. Shut down the Enterprise Manager.
3. Verify that the *introscope.apm.pruning.enabled* property is to *true*.

   This property enables the agents to prune application triage map data.

   If you reset *introscope.apm.pruning.enabled* from *false* to *true*, restart the Enterprise Manager for this property to take effect.

CA Technologies Pre-release Documentation. Authorized Use Only.

4. Optional: Configure these properties:
   - *introscope.apm.data.preserving.time*
   - *introscope.apm.pruning.cron.trigger.expression*

   a. Configure the *introscope.apm.data.preserving.time* property, which controls how long application triage map data is stored in the APM database. Map data older than the preserving time value is considered aged out, and can be pruned. The default is 365 days.

      During pruning, map data that is older than the sum of the *introscope.apm.data.preserving.time* and *Introscope.apm.data.agingTime* values is deleted. For example, if the *introscope.apm.data.preserving.time* value is 30 days and the *introscope.apm.data.agingTime* value is 1 day, map data older than 30 + 1 = 31 days is deleted.

      Update the value based on how often your organization adds and modifies applications, which adds application triage map data to the APM database. For more information, see introscope.apm.data.preserving.time (see page 435). The default value is 365 days.

   b. Configure the *introscope.apm.pruning.cron.trigger.expression* property, which schedules how often aged-out application triage map data is pruned from the APM database. The default value is daily at 6:00 AM.

      Update the value based on how often your organization adds and modifies applications, which adds application triage map data to the APM database. For more information, see introscope.apm.pruning.cron.trigger.expression (see page 436).

5. Save and close the *IntroscopeEnterpriseManager.properties* file.

6. Restart the Enterprise Manager.

   When the Enterprise Manager starts, you can go *<EM_Home>/logs* to review the Enterprise Manager log file messages associated with the map data pruning properties. For example:

```
11/23/09 12:02:13 PM PST [INFO] [Manager] [ The APM Data  property :
introscope.apm.pruning.enabled set to value : true]
11/23/09 12:02:13 PM PST [INFO] [Manager] [ The APM Data  property :
introscope.apm.pruning.cron.trigger.expression set to value : 0 0 6 * * ?]
```

CA Technologies Proprietary Documentation. Authorized Use Only.

Chapter 15: Configuring APM to Load Catalyst Data

The [assign the value for wacc for your book] lets CA Service Operations Insight ([set the varname value at the book level]) to connect to CA Application Performance Management (CA APM), providing visibility within [set the varname value at the book level] to the configuration items and information collected by CA APM. Connector lets CA APM to display data from other CA Technologies applications coming through [set the varname value at the book level].

CA APM

The CA APM solution provides 24x7, end-to-end monitoring and management of all transactions to identify, triage, prioritize, and resolve problems before they affect your end users and your business. CA APM enables you to proactively monitor and manage your increasingly complex web application services environment today and as you grow. Real-time detailed data provided by the solution helps you quickly understand and address problems before customers are impacted, meet service level agreements, and provide the online experience your customers expect.

CA APM's central component is the **Enterprise Manager**. The **CA Introscope Workstation** is a thick client user interface to the metrics, alerts and other elements provided by the Enterprise Manager. **WebView** is a thin-client, slimmed-down version of the information provided by the Workstation. The Customer Experience Manager interface, sometimes referred to as CA CEM, is a web-based interface to transactions, incidents and other data collected about website user experience by the TIM component of the Enterprise Manager.

To learn more about how CA Technologies products work together to help you monitor, detect, and triage application management problems, see the following demonstration:

http://www.ca.com/media/apm/wily/ca-wily-shell.html

This chapter describes how to configure CA APM to load data from CA Catalyst.

This section contains the following topics:

Configure CA APM to Load CA Catalyst Data (see page 228)
Configure the CatalystPolicy.xml File (see page 229)

CA Technologies Confidential Document
Authorized Use Only

Configure CA APM to Load CA Catalyst Data

If you configure the *Catalyst.properties* file, you can load data into CA APM from CA Catalyst, process it, and view it in the Workstation.

Follow these steps::

1. Open the *Catalyst.properties* file in the *<EM_Home>/config* directory.
2. Shut down the Enterprise Manager.
3. Edit these required properties:
   - *catalyst.host=<enter your CA Catalyst RESTful APIs location>*
   - *catalyst.protocol = http*
   - *catalyst.port = 7070*
   - *catalyst.login =<enter your CA Catalyst user name>*
   - *catalyst.password =<enter your CA Catalyst password>*

     Note: Setting the catalyst.login and catalyst.password properties allows the Enterprise Manager to log in to CA Catalyst. Upon restart (see final step in this procedure), the Enterprise Manager encrypts the Catalyst password using password-based encryption (PBE) with MD5 and Data Encryption Standard (DES) algorithms.
   - *catalyst.rest.entrypoint = /ssaweb/search*
4. (Optional) Set the CA Catalyst alert filtering properties.
5. Save and close the *Catalyst.properties* file.
6. (Optional) Open the *Catalyst.Policy.xml* file in the *<EM_Home>/config* directory and configure rules for the entity alerts association.
7. Save and close the *Catalyst.Policy.xml* file.
8. Restart the Enterprise Manager.

More Information:

Configure the CatalystPolicy.xml File (see page 229)
Catalyst.properties (see page 464)

CA Technologies Pre-release Documentation. Authorized Use Only

Configure the CatalystPolicy.xml File

The *CatalystPolicy.xml* file is an optional configuration that provides maximum flexibility for policy-based configurations with connectors. The *CatalystPolicy.xml* file lets you define rules such as USM entities, list of associated entities, and a list of associated relations. The configurations refine searches for alerts on any given entity based on the defined rules.

The following is an example of the XML file structure:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<policy xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" version="0.1"
xsi:noNamespaceSchemaLocation="policy0.1.xsd">
    <entity name="ComputerSystem" alertable="true">
        <include_entities>Processor Memory File Port</include_entities>
        <include_relations>IsHostFor IsComposedOf HasMember</include_relations>
    </entity>
    <entity name="Database" alertable="true">
        <include_entities>Tablespace</include_entities>
        <include_relations>IsManagerFor</include_relations>
    </entity>
    <entity name="DatabaseInstance" alertable="true">
        <include_entities>Tablespace</include_entities>
        <include_relations>IsManagerFor</include_relations>
    </entity>
    <entity name="Group" alertable="false">
        <include_entities>ComputerSystem VirtualSystem</include_entities>
        <include_relations>HasMember</include_relations>
    </entity>
</policy>
```

CA Technologies Reference Documentation. Authorized Use Only.

More Information:

Configure CA APM to Load CA Catalyst Data (see page 228)
Catalyst.properties (see page 464)

Chapter 16: Introducing CA CEM

Welcome to CA Customer Experience Manager (CA CEM), a tool that gives business management, customer service, application support, and IT operations a new level of real-time visibility into your users' (internal and external) transactions. This visibility enables IT organizations to deliver high quality transactions to their users at superior levels of operating efficiency.

CA CEM monitors your users' transactions to isolate the causes of problems in the data center. CA CEM measures the performance and quality of transactions, identifies defects and variance, and quantifies the impact on users and the business. By proactively detecting trends in degraded transaction response times and providing a variety of actionable reports, CA CEM enables you to take action before a problem occurs or Service Level Agreements (SLAs) are out of compliance.

This section contains the following topics:

Accessing the CEM console and setup pages (see page 231)
CA CEM administrator responsibilities (see page 235)

Accessing the CEM console and setup pages

You can set up and administer CA CEM from any computer that has a web browser.

This topic describes the interfaces of CA CEM and how to access them:

- Accessing the CEM console (see page 231)
- Accessing the Wily System Setup page for the TIM (see page 233)
- Accessing the TIM from the CEM console (see page 234)

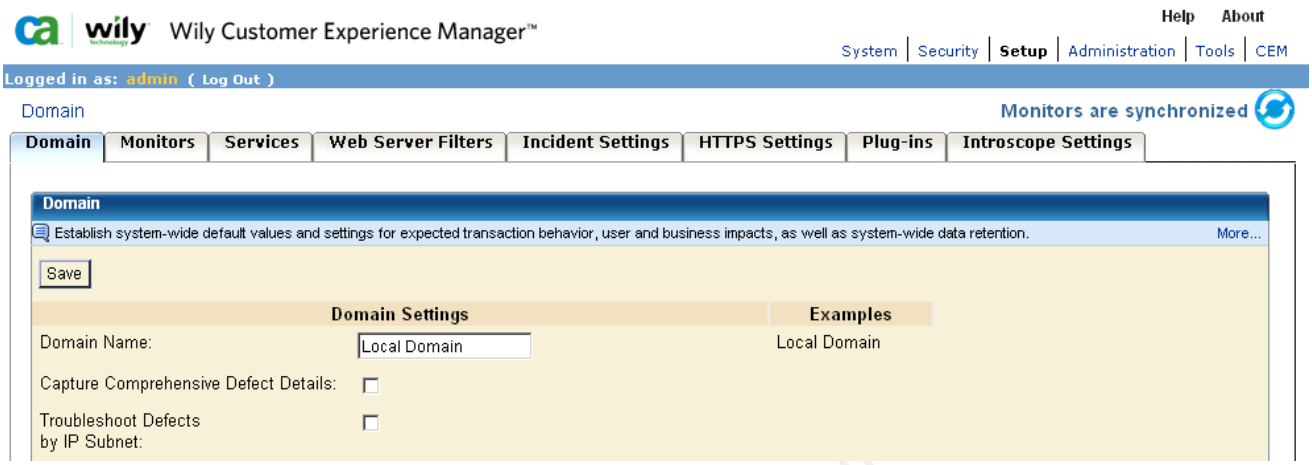
Access the CEM Console

You can set up and administer business applications, business services, and transaction definitions from any computer that has a web browser.

The CEM console is the main user interface of CA CEM. Administrators use it to set up and configure CA CEM, including transaction recording and creating transaction definitions. Analysts use it to produce and view reports.

CA Technologies Pre-release Documentation. Authorized Use Only.

If you are an administrator, the CEM console looks like this:



(If no TIMs are enabled, the Tools and CA CEM links do not appear.)

Unless otherwise indicated, this guide assumes that you are using the CEM console as an administrator.

In a clustered environment, the CEM console is on the Manager of Managers (MOM).

Follow these steps:

1. Open a web browser and type the address:
   `http://<IP_Address>:8081/wily`

   where *<IP_Address>* is the IP address or DNS name for the MOM or a standalone Enterprise Manager. For example:
   `http://192.168.1.54:8081/wily`
   `http://cem.company.com:8081/wily`

   To use a DNS name, your DNS administrator must have configured it.

   Note: The default port is 8081. It is defined in the IntroscopeEnterpriseManager.properties file as *introscope.enterprisemanager.webserver.port=8081* and can be changed.

2. Enter the user name and password.

   The default user name for the CA CEM administrator is *cemadmin*.

   The CEM console appears.

For information about security and logins, see the *CA APM Security Guide*.

CA Technologies Prerelease Documentation
Authorized Use Only.

Troubleshooting Login to CA CEM

Symptom:

If you are unable to log in to CA CEM, verify that Active Scripting is enabled on your browser.

If Active Scripting is disabled, when you click Log In on the CA CEM login page, you are redirected to a CEM Console / Back page. This happens even when your User Name and Password have been entered correctly.

Solution:

To enable Active Scripting, see the user documentation for your browser. For example, in Internet Explorer 8, the setting is under Internet Options > Security > Custom level > Scripting > Active scripting.

Access the TIM System Setup Page

The System Setup page for the TIM provides access to various configuration and status pages. CA CEM administrators rarely use these pages, typically to check status information or for troubleshooting.

CA Technologies Prerelease Documentation. Authorized Use Only.

Follow these steps:

1. Open a web browser and enter the address:

   `http://<IP_Address>/cgi-bin/wily/system/index`

   where *<IP_Address>* is the TIM IP address or DNS name.

   For example:
   `http://192.168.1.54/cgi-bin/wily/system/index`
   `http://cem.company.com/cgi-bin/wily/system/index`

   To use a DNS name, your DNS administrator must have configured it.

2. Enter the user name and password.

   The default user name for the System Setup page is *admin*.

   The System Setup page appears.

Access the TIM from the CEM Console

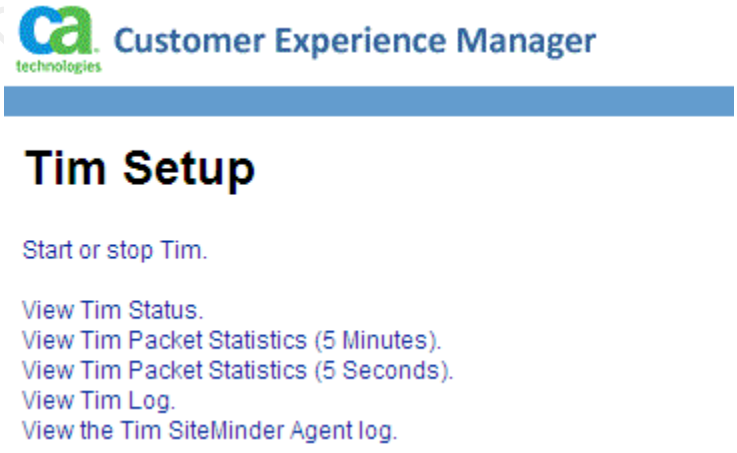
If you know the IP address and path for the TIM, you can access the TIM Setup page from the System Setup page. However, if you have the CEM console open, you can access the TIM Setup page from the console.

Follow these steps:

1. In the administrator console, select Setup > Monitors.
2. Click the IP address of the TIM (far right column).
3. Enter the user name and password.

   The default user name for the System Setup page is *admin*.

   The TIM Setup page appears.



CA Technologies Prerelease Documentation. Authorized Use Only.

CA CEM administrator responsibilities

The CA CEM administrator's responsibilities vary at each site, depending on environment complexity, type of integration, number of users, support resources, and the growth rate of the CA CEM-monitored applications.

- Daily tasks (see page 235)
- Weekly tasks (see page 235)
- Monthly tasks (see page 236)
- Quarterly tasks (see page 236)

Daily tasks

These are typical daily tasks for the CA CEM administrator:

- Perform and verify the daily backup.
- Review incidents and if appropriate defects. Resolve issues using Introscope.
- Report any discovered CA CEM issues to the appropriate groups.
- Review the TIM and IntroscopeEnterpriseManager logs.
- Review the system events online or after receiving email notification. For example, are defects being dropped. (See Viewing (see page 311)CA CEM events (see page 311) and Receiving email notification of events. (see page 312))
- Review the health of the TIMs. (See Reviewing TIM machine health (see page 314).)
- Review the health of CA CEM on the Enterprise Manager using Introscope Workstation.

Weekly tasks

These are typical weekly tasks for the CA CEM administrator:

- Test the backup and restore procedures.
- Test that the integration with CA SiteMinder, CA NSM, CA Service Desk, etc. is working as expected. For example, is single sign-on working with SiteMinder web agent?
- Save off site a copy of the CA CEM configuration and backup files.
- Review CA CEM administrator list in Embedded Entitlements Manager for unneeded administrators.
- Review CA CEM audit log for unexpected or unauthorized changes.

CA Technologies Proprietary Documentation. Authorized use only.

- Review CA CEM CPU, disk, and system utilization. (See Reviewing TIM machine health (see page 314).)
- Close resolved or old incidents as appropriate.
- Review reports for accuracy and trends.
- Check that email reports are being created and distributed as expected.
- Review the real-time business transaction metrics and custom dashboards in Introscope.
- Review integration logs, such as the logs for CA Service Desk and CA SiteMinder.
- Review the evidence collection logs. (See Setting business impact and collecting evidence for incidents (see page 276).)

Monthly tasks

These are typical monthly tasks for the CA CEM administrator:

- Check failover procedure (if installed).
- Clean up old/unneeded users and user groups.
- Review if new business services/transactions/etc. need to be defined.
- Review if transaction specification settings need to be modified.
- Review if baseline or SLA settings need to be adjusted.
- Review monthly reports for accuracy and trends.
- Review the enabled transactions and the business transaction statistics; disable definitions for transactions that no longer need to be monitored. (See the *CA APM Transaction Definition Guide* for information about managing transaction definitions.)
- If an HTTP analyzer plug-in is enabled, review its operation. Is the plug-in meeting the expected need? Are there additional requirements that it needs to meet? Has there been any performance degradation since the last JAR file update? (See the *CA APM Transaction Definition Guide*.)

Quarterly tasks

These are typical quarterly or yearly tasks for the CA CEM administrator:

- Review CA CEM architecture for performance, scalability, failover, geographic distribution, and other factors.
- Determine if CA CEM Professional Services are required for CA CEM customization or integration.

CA Technologies Prerelease Documentation. Authorized Use Only

- Beta test (if appropriate) next CA CEM release for possible problems, new features.
- Review user group definitions. (For example, do the IP Subnet user group rules need modification?)

CA Technologies Prerelease Documentation.
Authorized Use Only.

Chapter 17: Networking Alternatives for CA CEM

This section contains the following topics:

Process for CA CEM networking (see page 239)
About CA CEM network options (see page 239)
About monitoring network traffic (see page 240)
About CA CEM insertion locations (see page 241)
Connections from browser to the Enterprise Manager (see page 245)
IPv6 and CA CEM (see page 245)

Process for CA CEM networking

This list shows the decisions involved with CA CEM networking:

1. Learn about the network options available for CA CEM (see page 239).
2. Understand how network taps and mirrored ports work with (see page 240)CA CEM. (see page 239)
3. Understand the pros and cons for each (see page 241)CA CEM (see page 239)network insertion option. (see page 241)
4. Learn about running (see page 245)CA CEM (see page 239)in an IPv6 network and about monitoring IPv6 network traffic. (see page 245)

About CA CEM network options

There are multiple locations within your network where it is possible to locate CA CEM and alternative methods for connecting the system to the network. This chapter describes the options.

TIMs and the Enterprise Manager reside on separate hardware. The primary consideration is connectivity to the TIMs that report to the Enterprise Manager (Collectors and Manager of Managers (MOM)).

There are multiple options for where you place a TIM; they are described in About CA CEM insertion locations (see page 241). The location of a TIM determines what information you can collect.

CA Technologies Prerelease Documentation. Authorized Use Only.

The preferred insertion point for TIMs is logically between the network load balancer and the web server aggregation switch, or between the load balancer and the web servers if no aggregation switch is used. For more information, see Between the aggregation switch and the network load balancer (see page 244).

When you install a TIM in the location of your choice, you must also select a method of connecting it to your network.

About monitoring network traffic

The TIMs read traffic from the network via a monitoring interface, which can be connected to a mirrored port or to a network tap.

The following sections describe the trade-offs between using a mirrored port versus a network tap to forward packets to the TIMs.

- Using network taps (see page 240)
- Using mirrored ports (see page 241)

Using network taps

Network taps are hardware devices that tap directly into the infrastructure cabling, creating copies of packets and forwarding them on to one or more destinations. Taps tie into the infrastructure at the physical layer, and do not use resources on routers or switches. Taps require that a connection be temporarily suspended in order to insert the tap.

Taps have no effect on network latency. Many taps forward transmitted and received traffic separately on two physical interfaces, requiring two physical interfaces in the TIM. These data streams must be merged by the TIM. This merging of streams has a slight performance impact on the TIM. Some taps are full-duplex and do not require additional interfaces or merging on the TIM.

Taps are transmit-only devices, so they might be more secure than mirrored ports. Another advantage of taps is that they are relatively simple to configure. Because a tap is an inline device, failure can disrupt operation of the network or require an outage to replace the defective tap (unless it is a passive tap). Advances in tap technology have greatly reduced the probability of failure.

CA Technologies Prerelease Documentation. Authorized Use Only.

Using mirrored ports

Mirrored ports (SPAN ports, in Cisco terminology) are a software feature of network routers and switches. In modern switches, such as the 6500 series switches from Cisco, enabling SPAN ports has no impact on switch performance. Mirrored ports are often used to attach network troubleshooting devices such as protocol analyzers. There is a limit to the number of mirrored ports a switch supports, so in some cases network engineers are reluctant to allow permanent attachment of devices to these maintenance ports.

Because mirrored ports are a software feature of a router or switch, they do not require an infrastructure change and are less disruptive to enable than a tap. Unlike taps, which require additional capital expense, switches that are already deployed can usually support mirrored ports with no additional expense.

Mirrored ports have no effect on network latency. Mirrored ports often combine send and receive traffic prior to forwarding so that fewer interfaces are required on the TIMs. TIMs do not have to merge data streams from multiple interfaces so performance is increased somewhat in this configuration.

Mirrored ports are transmit-only ports in most cases, although incorrect configuration changes could cause a mirrored port to be configured as a bidirectional switch port. For this reason mirrored ports are sometimes considered less secure than taps. The TIM software is designed so that it does not transmit on monitoring ports, ensuring that the TIM introduces no security risk.

Note: If your application has end-user traffic coming in through two different switches then it is possible for the traffic from both switches to be mirrored to the same TIM. You need to procure and deploy the network card (multi-port NIC) to do this. See the *CA APM Compatibility Guide* for information on network card support.

About CA CEM insertion locations

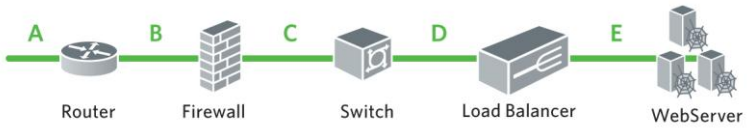
The insertion location of the TIMs in the infrastructure influences the information collected.

TIMs read traffic from the network via a monitoring interface. As described in the previous section, the monitoring interface can be connected to a mirrored port or to a network tap. In both cases the TIM receives copies of packets of interest, analyzes them, and discards them. The TIM does not transmit on the monitoring port. The traffic is unidirectional into the TIM, and the TIM is configured to receive packets only from the monitoring port.

CA Technologies Proprietary/Trade Secret Documentation. Authorized Use Only.

In a typical network infrastructure for a large enterprise there are a number of locations in which to insert the TIM monitoring port. The diagram illustrates different insertion locations.

Note: Most environments make use of a redundant infrastructure. For simplicity however, only one path is shown in the diagram.



Reading the diagram left to right, packets enter the enterprise via an Internet router. Insertion location A is on the outside of the router, between the Internet and the router.

This router typically forwards traffic to a firewall, and insertion location B is between the router and this firewall. The firewall can inspect the inbound and outbound traffic and block or accept the traffic per the enterprise's security policy. The firewall can also perform network address translation (NAT) to protect the enterprise's network-attached devices from direct connection to the Internet.

Note: In some cases, the aggregation switch and the network load balancer are housed in the same device.

Insertion location C is between the firewall and the aggregation switch / network load balancer. Acceptable traffic is forwarded from the firewall to a aggregation switch / network load balancer. In many cases the aggregation switch / network load balancer can terminate SSL sessions as well as instantiate a virtual IP address (VIP) for clients to connect to. Many aggregation switches / network load balancers act as a proxy, terminating connections from clients and establishing connections, on behalf of the clients, to the web servers. Some aggregation switches / network load balancers establish and use persistent connections to the web servers to decrease the load on the web servers.

In many instances the aggregation switch forwards packets to a network load balancer. Insertion location D is between the aggregation switch and the network load balancer.

Finally, the web servers are connected to the aggregation switch, or directly to the network load balancer. Insertion location E is directly in front of the web servers, either between an aggregation switch and the web servers or between a network load balancer and the web servers.

Packets leave the enterprise in the reverse path, passing through each of the insertion locations in reverse order. (It is also possible for packets enter taking one path and exit via a different path (that is, a triangle). Or they might leave the web server and go directly to the Internet.

CA Technologies Pre-release Documentation. Limited Use Only.

Between the Internet and the Internet router

Insertion location A, between the Internet and the Internet router, always provides visibility into the client IP address.



Insertion location A almost never provides visibility into the server's IP address or the server's MAC address. CA CEM uses the web server IP or MAC address in its exclusion-based relationship discovery mechanisms to limit the set of evidence collected for faster troubleshooting. Without this information evidence collection continues to work, but the evidence set is not reduced, and technicians must analyze more information than they would without the use of exclusion-based discovery.

Because the TIM does not use the client IP address for identifying the user, there are very few advantages to using insertion location A, and there are a number of disadvantages.

Between the Internet router and the firewall

Insertion location B, between the Internet router and the firewall, nearly always provides visibility into the client IP address.

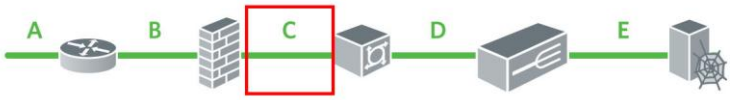


Insertion location B almost never provides visibility into the server's IP address or the server's MAC address. CA CEM uses the web server IP or MAC address in its exclusion-based relationship discovery mechanisms to limit the set of evidence collected for faster troubleshooting. Without this information evidence collection continues to work, but the evidence set is not reduced, and technicians must analyze more information than they would without the use of exclusion-based discovery.

Because the TIM does not use the client IP address for identifying the user, there are very few advantages to using Insertion location B, and there are a number of disadvantages.

Between the firewall and the aggregation switch

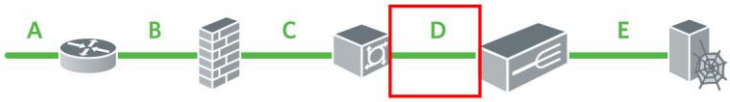
Insertion location C, between the firewall and the aggregation switch, nearly always provides visibility into the client IP address. Insertion Location C almost never provides visibility into the server's IP address.



Because the TIM does not use the client IP address for identifying the user, there are very few advantages to using insertion location C, and there are a number of disadvantages. In some environments, insertion location C provides marginal benefits over insertion locations A and B.

Between the aggregation switch and the network load balancer

Insertion location D, between the aggregation switch and the network load balancer, often provides visibility into the client IP address. In order for this to occur, "preserve source address" must be enabled in the network load balancer if it is acting as a proxy. Insertion location D almost always provides visibility into the server's IP address.



If the network load balancer uses persistent connections, the load on the TIMs is decreased when the TIMs are located here. The decrease in load is a function of the reduction in open connections, which might be significant. If the network load balancer terminates SSL on behalf of clients and does not re-encrypt the traffic, then the load on the TIMs is further reduced in nearly all cases by a significant amount.

In some load balancing environments, such as Resonate, it might be possible to identify the server by its MAC address. In many environments, however, the server IP address and server MAC address returned is that of the network load balancer virtual IP address (VIP).

Because the TIM uses the server IP address for evidence collection, there are significant advantages to using insertion location D, and virtually no disadvantages. Insertion location D is the preferred insertion location for the TIMs.

Between the network load balancer and the web server

Insertion location E, between the network load balancer and the web server (or servers), often provides visibility into the client IP address. In order for this to occur, "preserve source address" must be enabled in the network load balancer if it is acting as a proxy. Insertion location E almost always provides visibility into the server's IP address.



If the network load balancer uses persistent connections, the load on the TIMs is decreased when the TIMs are located here. The decrease in load is a function of the reduction in open connections, which might be significant. If the network load balancer terminates SSL on behalf of clients and does not re-encrypt the traffic, then the load on the TIMs is further reduced in nearly all cases by a significant amount.

Because the TIM uses the server IP address for evidence collection, there are significant advantages to using insertion location E. However, insertion location E has some disadvantages that insertion location D lacks. Since there might be many web servers, there might be many insertion locations. This might make it impractical to use taps for insertion location E.

Insertion location D and insertion location E are functionally equivalent when port mirroring is used. In many cases spanning a virtual LAN (VLAN) is all that is necessary to use insertion location E. For insertion location D, simple port mirroring might be sufficient.

Connections from browser to the Enterprise Manager

By default, HTTP access is allowed between the browser and the Enterprise Manager.

However, you can restrict access to the Enterprise Manager to allow SSL connections alone. See the *CA APM Security Guide*.

IPv6 and CA CEM

This section briefly discusses Internet Protocol version 6 (IPv6), the machine settings pages for CA CEM on an IPv6 network, and how IPv6 addresses are formatted in CA CEM.

- About IPv6 (see page 246)
- IPv6 in CA CEM (see page 246)

- Monitoring IPv6 network traffic (see page 246)
- Running CA CEM on dual-stack (IPv6 and IPv4) hosts (see page 247)
- Allowed formats for IPv6 addresses (see page 248)

About IPv6

The number of Internet Protocol version 4 (IPv4) addresses is expected to run out within the next few years. IPv6 increases the number of Internet addresses available.

The IPv6 address is 128 bits (16 bytes). It is typically written in base 16 (hexadecimal), with each 16-bit group separated by a colon (:). For example, 2001:0db8:30d2:55c9:06ff:e579:5aa4.

By contrast, the IPv4 address is 32 bits (4 bytes). It is typically written in base 10, with each byte separated by a point (.). For example, 172.24.131.197.

IPv6 in CA CEM

There are two aspects to IPv6 support in CA CEM:

- Monitoring IPv6 network traffic
- Configuring CA CEM on an IPv6 network

These two aspects are not dependent on each other. For example, you can place your CA CEM machines on an IPv4 network and monitor IPv6 *and* IPv4 network traffic. Similarly, you can place your CA CEM machines on an IPv6 (or dual-stack) network and monitor IPv6 *and* IPv4 network traffic.

Monitoring IPv6 network traffic

CA CEM supports both IPv4 and IPv6 addresses. When you enter an IP address or subnet mask in a field, you can enter either an IPv4 or an IPv6 address.

To monitor IPv6 network traffic:

- Set up CA CEM in the same way as for monitoring IPv4 traffic. Enter IPv6 addresses (and host names) instead of IPv4 addresses (and host names) where appropriate.

  See Allowed formats for IPv6 addresses (see page 248).

CA Technologies Prerelease Documentation
Authorized Use Only.

Running CA CEM on dual-stack (IPv6 and IPv4) hosts

When you set up and install CA CEM, follow the instructions provided by your SE or support personnel or given in the *CA APM Installation and Upgrade Guide*. As dual-stack hosts, IPv4 *and* IPv6 addresses are assigned to the Enterprise Manager and TIM machines.

Important! If the TIMs are assigned IPv6 addresses, then the Enterprise Manager that runs the TIM Collection service must also be assigned an IPv6 address. Similarly, if the TIMs have IPv4 addresses, the Enterprise Manager that runs the TIM Collection service must have an IPv4 address.

This important constraint is enforced in the CA CEM console; for example:

- If enabled TIMs have IPv6 addresses, you cannot run the TIM Collection service on an Enterprise Manager with an IPv4 address (or vice versa).
- If an enabled TIM has an IPv6 address, you cannot enable a TIM with an IPv4 address (and vice versa).
- You cannot change the IP address of the Enterprise Manager running the TIM Collection service if it will cause a mismatch of address type between the Enterprise Manager and the enabled TIMs.
- You cannot move the TIM Collection service to another Enterprise Manager if it will cause a mismatch of address type between the Enterprise Manager and the enabled TIMs.

If you need to confirm, modify, or troubleshoot the configuration, follow these procedures.

To review or modify IPv6 settings for CA CEM on a TIM:

1. Access the Wily System Setup page on the TIM. (Accessing the Wily System Setup page for the TIM (see page 233).)
2. From the Wily System Setup page, select Machine Settings > Configure IPv6. Verify that the IPv6 addresses are correctly assigned and that IPv6 is enabled.

   Points of note about the IPv6 Setup page are:

   - The global scope IPv6 addresses are analogous to the IPv4 address.
   - RADV and DHCPv6 can assign more than one global IPv6 address to the machine. However, only one address can be assigned manually.

CA Technologies Prerelease Documentation
Authorized Use Only.

- If you assign the IPv6 address manually, use the slash notation to indicate the length of the network prefix. (This is similar to specifying the subnet mask for the IPv4 network.)
- The Click to disable buttons are reserved for troubleshooting. For example, if you aren't able to communicate with a CA CEM machine over the IPv6 network, verify that IPv6 is enabled both system-wide and for the interface that is connected to the management network (eth*n*).

3. Verify that your IPv4 settings are correctly assigned: return to Machine Settings and select Network Setup.

To change the IP address for dual-stack Enterprise Manager:

- See Editing the IP address for an Enterprise Manager (see page 258).

To point your browser to the CA CEM console:

- If you use the literal IP address in the URL for the console, use the [ ]s (square bracket) notation to denote the IPv6 address.

  See the path examples in Allowed formats for IPv6 addresses (see page 248).

Allowed formats for IPv6 addresses

The notation used for IPv6 addresses varies depending on whether you are in the CEM console or the Wily System Setup pages:

- CEM console — When entering IPv6 addresses in the CEM console, you *cannot* use the compressed notation except for specifying:
  - the subnet mask, where compressed notation *and* the slash notation together are allowed
  - the Client IP Address on the Administration > Recording Sessions page (where both compressed notation and mixed notation can be used).

  Make sure to follow the examples given in the UI.
- Wily System Setup pages — In the Wily System Setup pages (such as the IPv6 Setup page), you *can* use the compressed notation as well as the longer notations to specify IPv6 addresses.

  You must use a prefix length (separated from the address with a slash) instead of a subnet mask in the IPv6 Setup page.

This table lists examples of accepted notation for IPv6 addresses in CA CEM.

Example IPv6 addresses	Notation and Notes
3fff:ffff:0ed5:0000:0000:0000:2448:78eb	Standard notation: Use anywhere in CA CEM.

CA Technologies Prerelease Documentation
Limited Use Only.

Example IPv6 addresses	Notation and Notes
3fff:ffff:ed5:0:0:0:2448:78eb	Leading zeros removed within a 16-bit group: Use anywhere in CA CEM.
3fff:ffff:fed5::2448:78eb	Compressed notation: Do not use in the CEM console UI except with slash notation in the subnet mask and on the Administration > Recording Sessions page.
3fff:ffff:fed5:0::/64	Slash notation: Use for subnet mask (Setup > Domain > Subnet Mask and Administration > User Groups) and for the Manual IPv6 Address (Wily System Setup > Machine Settings > Configure IPv6).
ffff:ffff:ffff:ffff:0:0:0:0	Subnet Mask: Use for the Subnet Mask (Setup > Domain > Subnet Mask and Administration > User Groups).
2001:0db8:85a3:08d3:0:0:172.31.124.132	Mixed notation IPv4 and IPv6: Do not use in CA CEM, except on the Administration > Recording Sessions page.

This table lists examples of accepted notation for IPv6 paths in CA CEM.

Example IPv6 paths	Notation and Notes
http://cemhost_ipv6:8081/wily	Example path to the CEM console, where cemhost_ipv6 is defined as a hostname.
http://[2001:db8:85a3:8d3:0:0:370:7348]:8081/wily	Example path to the CEM console, where the Enterprise Manager IP address is IPv6.

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 18: Administering CA CEM

This section contains the following topics:

Process for CA CEM administration (see page 251)
Distributing Enterprise Manager services (see page 252)
Defining your CA CEM domain (see page 262)
Incident cost versus business impact (see page 268)
Defining incident-related settings (see page 270)
Configuring SMTP server settings (see page 281)
Setting the security groups for new business services (see page 283)
Enabling monitors (see page 287)
Synchronizing monitors (see page 288)
Viewing configuration changes since last synchronization (see page 289)

Process for CA CEM administration

This list shows the main tasks involved in administering CA CEM:

1. Distribute Enterprise Manager services across multiple Enterprise Managers. See Distributing Enterprise Manager services (see page 252)
2. Define your domain, including settings for user recognition, user groups, transactions, and data retention. See Defining your CA CEM domain (see page 262)
3. Set incident settings for incident generation, business impact and evidence collection, incident age-out rules, and email on incidents. See Defining incident related settings (see page 270)
4. Define email settings for reporting and event notifications. See Configuring SMTP server settings (see page 281)
5. Set the default security groups for new business services. See Setting the security groups for new business services (see page 283)
6. Enable the TIMs—this task is required at setup and any time the TIMs have been disabled. This enables communication between a Collector and a TIM. See Enabling monitors (see page 287)
7. Synchronize monitors—this task is required to send transaction and defect configuration information to the TIMs and have the TIMs begin monitoring with the updated configuration.

   It is also required to send configuration changes related to business services and transactions to the Introscope agents. See Synchronizing monitors (see page 288)

CA Technologies Proprietary/Confidential Documentation. Authorized Use Only.

Distributing Enterprise Manager services

If you have more than one Enterprise Manager (EM), that is, Collectors and a MOM, then you can distribute the Enterprise Manager services across the Enterprise Managers. If you do not distribute services, the Enterprise Manager services all run on one Enterprise Manager; in some situations, this results in sub-optimal performance.

Note: If you have only one (standalone) Enterprise Manager in your system, then you can skip this topic. In the case where the Enterprise Manager is standalone, all Enterprise Manager services must run on the one Enterprise Manager.

When you start CA APM in a multi-Enterprise Manager environment, the Enterprise Manager services all start on the first Collector. After all the Collectors have started, then you can assign the Enterprise Manager services to the appropriate Enterprise Managers.

This table indicates when to perform the Enterprise Manager services tasks:

If you are…	Do these tasks…
Setting up CA APM for the first-time	Distributing Enterprise Manager services (see page 252) Modifying the Enterprise Manager configuration (see page 258)
Adding Enterprise Managers to your environment	Distributing Enterprise Manager services (see page 252) Modifying the Enterprise Manager configuration (see page 258)
Changing the IP address of an Enterprise Manager in your environment	Modifying the Enterprise Manager configuration (see page 258)
Removing Enterprise Managers from your environment	Distributing Enterprise Manager services (see page 252) (to move Enterprise Manager services from the Enterprise Manager before you shut down the Enterprise Manager) Removing an Enterprise Manager (see page 259)

CA Technologies Prerelease Documentation. Authorized Use Only.

You can distribute these Enterprise Manager services:

Note: You can have one (and only one) instance of each service running in a cluster. For example, you would never have the TIM Collection service running on two Enterprise Managers at the same time.

DB Cleanup service

Deletes old data from the database.

Stats Aggregation service

Aggregates statistics in the database once a day. By default, aggregation is around midnight. For more information, see About CA CEM report aggregation (see page 319).

TIM Collection service

If you do not have any TIMs (that is, if you are running Introscope but not CA APM), then the location of this service does not matter.

Is made up of a group of services used to process data collected by the TIMs. Includes the services for processing logins, defects, recording, CA CEM events, hourly statistics, customer experience metrics, and transaction discovery.

Note: The TIM Collection service is the most computationally intensive of the Enterprise Manager services.

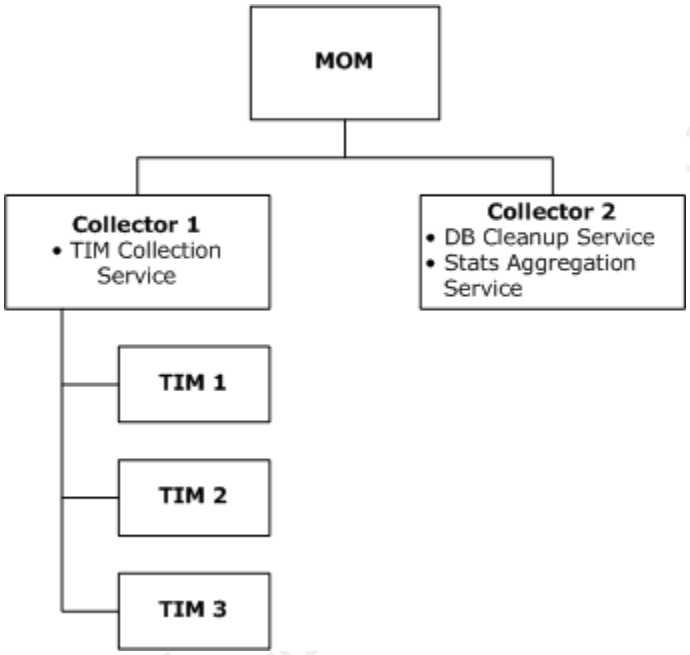
Other Enterprise Manager services are assigned to all Enterprise Managers. Because these services cannot be distributed, they are not shown on the Services pages.

CA Technologies Pre-release Documentation. Authorized Use Only

Typical distributions of the Enterprise Manager services

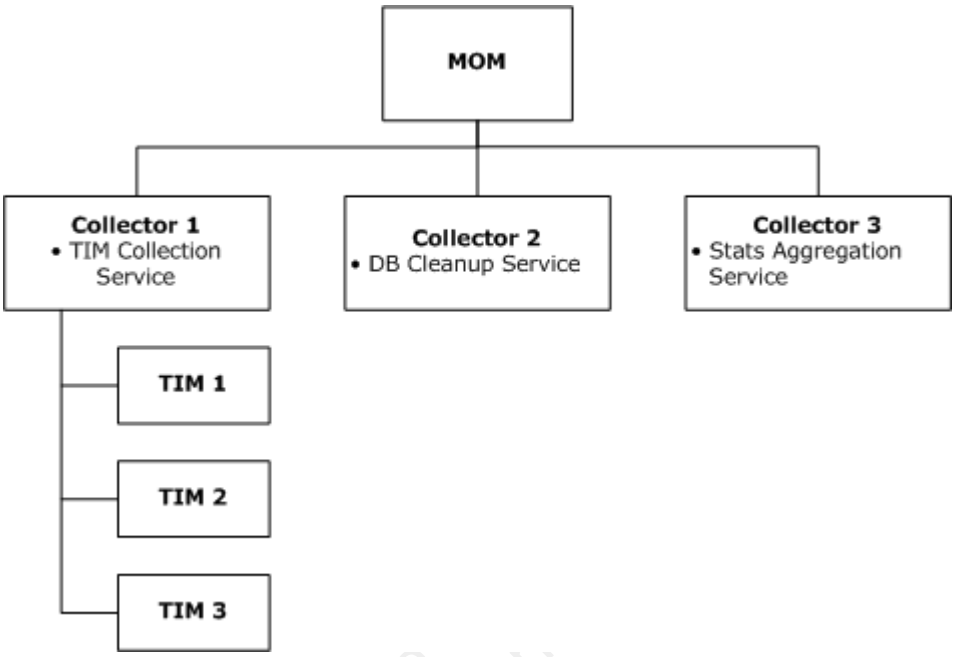
Note: Although you can run Enterprise Manager services on a MOM, it is not recommended.

This shows a typical distribution of Enterprise Manager services in a two-Collector environment:



This shows a typical distribution of Enterprise Manager services in a three-Collector environment:

CA Technologies Authori... ...ation.



Distributing Enterprise Manager services in a clustered environment

This is the procedure to redistribute the Enterprise Manager services across the Enterprise Managers. For example, do this the first time you start CA APM and when you add or remove Enterprise Managers from your system.

Important! Remove Enterprise Manager services from an Enterprise Manager before you shut it down.

The purpose of this procedure is to distribute workload across the Enterprise Managers.

CA Technologies recommends that you choose an off-peak time to redistribute Enterprise Manager services. Shutting down and moving services at peak time can result in processing delays.

Note: When the TIM Collection Service is moved to a different Enterprise Manager, any unprocessed data files remaining on the original Enterprise Manager do not get processed (unless the service gets moved back to the original Enterprise Manager).
See Changing the directory locations for startup, download and bad files (see page 530).

CA Technologies Pre-release Authorized Use Only

To distribute Enterprise Manager services in a clustered environment:

1. Select Setup > Services > Service Configuration.

   The Service Configuration page shows an overview of which Enterprise Managers are currently running each service.

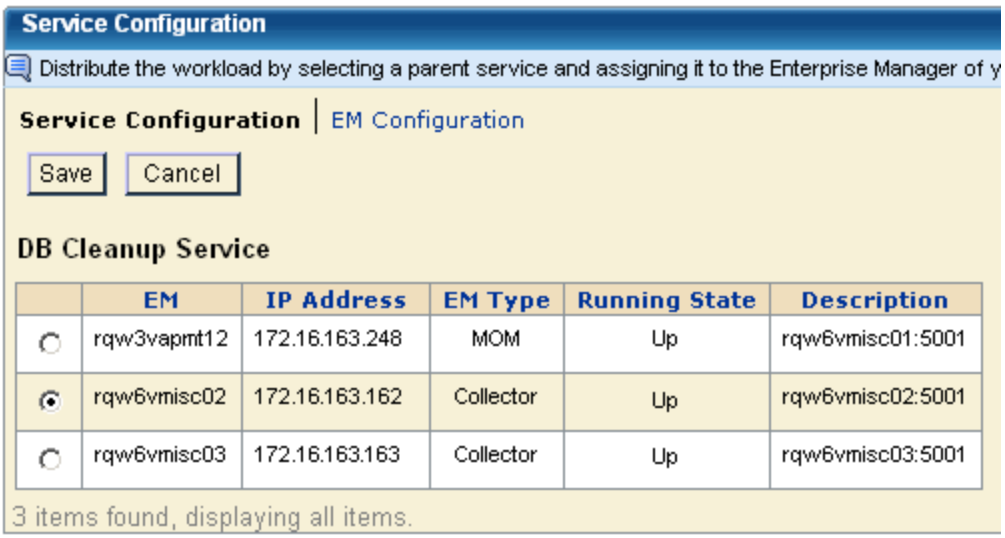
2. Click the Service that you want to assign to a different Enterprise Manager.

   For example, the DB Cleanup service.

3. Select the Enterprise Manager to which you want to assign the service and click Save.

   All Enterprise Managers in the cluster appear in the list. You can determine—at a glance—if an Enterprise Manager is up or down, and if it is a MOM or Collector.

   CA Technologies recommends that you do not run antivirus software on the Enterprise Manager that is running the TIM Collection service because antivirus can slow the service excessively.



4. Repeat steps 2 and 3 until each of the Enterprise Manager services is running on the Enterprise Manager you want.

5. Click Save.

Selecting the TIMs for real-time transaction metrics

By default, real-time transaction metrics data is collected from all TIMs. However, you can specify that only some of the TIMs in your environment participate. You might want to do this in a testing or troubleshooting situation.

To view or specify which TIMs participate in real-time transaction metrics:

1. Select Setup > Services.
2. Click TIM Collection Service.
3. In the TIM Collection Service menu, select RTTM Collection Service, if it is not already selected.

   The TIMs that are participating in real-time transaction metrics are marked enabled in the TIMs list.
4. If necessary, select or clear TIMs from the list.
5. Click Save.

Selecting the TIMs for automatic transaction discovery

By default, HTTP traffic from all TIMs is scanned when automatic transaction discovery is running. However, you can specify that only some of the TIMs in your environment participate in automatic transaction discovery.

You might want to do this if you know that the transactions of interest are only being accessed by a particular TIM. Limiting transaction discovery to this TIM means no overhead related to discovery is placed on the other TIMs.

To view or specify which TIMs participate in automatic transaction discovery:

1. Select Setup > Services.
2. Click TIM Collection Service.
3. In the TIM Collection Service menu, select Transaction Discovery Collection Service.

   The TIMs that are participating in transaction discovery are marked enabled in the TIMs list.
4. If necessary, select or clear TIMs from the list.
5. Click Save.

CA Technologies Proprietary & Confidential Documentation. Authorized Use Only.

Modifying the Enterprise Manager configuration

If you want to modify the description or IP address for an Enterprise Manager, follow these procedures.

Editing the description

By default, the description is set to include the hostname and port number of the Enterprise Manager.

For example, you might edit this description to add the machine's asset label or a location ID to help identify it.

To modify the description for an Enterprise Manager:

1. Select Setup > Services > EM Configuration.
2. Select an Enterprise Manager.
3. Edit the Description and click Save.

Editing the IP address for an Enterprise Manager

In most cases, the correct IP address for the Enterprise Manager is automatically supplied in the EM IP address field.

Here are some examples of when you might need to edit the IP address for an Enterprise Manager:

- The Enterprise Manager machine has two IP addresses because it is on two networks and the EM Configuration page is showing one IP address and the TIM Settings page shows the other address for the TessCollectorIpAddr or the TessIpAddr value. To view the TIM Setting page, see the optional step below.
- The Enterprise Manager is a dual-stack host with both an IPv6 and an IPv4 address. If the enabled TIMs have IPv6 addresses, then the Enterprise Manager that runs the TIM Collection Service needs to show an IPv6 address in the EM IP Address field.

To modify the IP address for an Enterprise Manager:

1. Select Setup > Services > EM Configuration.
2. Select an Enterprise Manager.
3. Select an alternate Enterprise Manager IP Address and click Save.

   After you have edited this address once, you must manually update it if the IP address changes again. (The field is no longer automatically updated when the Enterprise Manager restarts.)
4. Select Setup > Monitors.

CA Technologies Proprietary Documentation. Authorized Use Only.

5. Disable and Enable all TIMs:
   a. Select the TIMs.
   b. Click Disable.
   c. Click Enable.
6. (Optional) If the Enterprise Manager is the MOM or is running the TIM Collection Service, confirm that the enabled TIMs have received the updated IP address:
   a. Click the IP address of the TIM.

      The TIM System Setup page appears.
   b. Click Configure TIM Settings.
   c. Confirm that the TessCollectorIpAddr and TessIpAddr settings are correct. Do not edit these settings directly.

Removing an Enterprise Manager

This procedure explains how to remove an Enterprise Manager from the Services pages.

Note: You should do this *after* you shut down an Enterprise Manager.

Enterprise Managers automatically appear on the Services pages on start up. However, the names of the Enterprise Managers do *not* automatically get removed from the Services pages. The purpose of this procedure is to remove the names of unused Enterprise Managers from the Services page.

To remove an Enterprise Manager:

1. If the Enterprise Manager has any Enterprise Manager services running on it, move them to a different Enterprise Manager. See Distributing Enterprise Manager services (see page 252).
2. Shut down the Enterprise Manager. See the *CA APM Configuration and Administration Guide*.
3. Select Setup > Services > EM Configuration.
4. Select an Enterprise Manager.
5. Click Delete.

   The Enterprise Manager no longer appears on the Services pages.

If you later restart this Enterprise Manager, it reappears on the Services pages.

CA Technologies Pre-release Documentation. Authorized Use Only.

Troubleshooting Enterprise Manager services

This table provides some tips that might help you troubleshoot Enterprise Manager service problems.

Symptom:

CA CEM events related to services appear on the events page:

- Tim Collection Service has stopped
- Stats Aggregation Service has stopped
- Database Cleanup Service has stopped

Solution:

One of the Enterprise Manager services has gone down and needs to be restarted as soon as possible.

Restart the Enterprise Manager that is running the service.

Symptom:

It appears that daily statistics are not being aggregated.

Reports do not show the daily data.

Solution:

It could be that the Stats Aggregation service is not running correctly.

Review the IntroscopeEnterpriseManager log file that is on the Enterprise Manager running the Stats Aggregation service.

Symptom:

It appears that the Enterprise Manager is not receiving data from the TIMs in a timely way.

Solution:

It could be that the TIM Collection service is not running correctly.

Review the IntroscopeEnterpriseManager log file that is on the Enterprise Manager running the TIM Collection service.

Symptom:

File processing (of logins, defects and real-time transaction metrics) is falling behind, and there is a backlog of files in the <EM_Home>/cem/tess/downloadfiles.

CA Technologies Prerelease Documentation. Authorized Use Only.

Solution:

Other processes could be competing with the TIM collection service for resources (CPU, memory, I/O) on the Enterprise Manager.

If there is antivirus software running on the Enterprise Manager, then turn off the antivirus protection or at least verify that it is not scanning temporary data storage areas.

By default the storage areas are <EM_Home>/cem/tess/downloadfiles, <EM_Home>/cem/tess/startupfiles, and <EM_Home>/cem/tess/badfiles directories. If you have changed them, see Changing the directory locations for startup, download and bad files (see page 530) in the *CA APM Configuration and Administration Guide*.

Symptom:

The Collector or the MOM are not receiving information from the TIMs.

For example, there are error messages like the following in the TIM log:

```
ERROR There is no collector service for this TIM
```

Solution:

The IP address of the Collector or the MOM might have changed and the Enterprise Manager Configuration page does not reflect the change.

Restart the Enterprise Manager or follow the procedure Editing the IP address for an Enterprise Manager (see page 258) in the *CA APM Configuration and Administration Guide*.

Symptom:

Same as above. (The Collector or the MOM are not receiving information from the TIMs.)

For example, there are error messages like the following in the TIM log:

```
ERROR There is no collector service for this TIM
```

Solution:

The Collector or the MOM might have two IP addresses, and one IP address appears on the Enterprise Manager Configuration page and the other on the TIM Settings page.

Look for a mismatch of addresses in the EM Configuration page and the TIM Settings page.

If the addresses are not the same, follow the procedure Editing the IP address for an Enterprise Manager (see page 258) in the *CA APM Configuration and Administration Guide*.

CA Technologies Proprietary Documentation
Authorized Use Only.

Symptom:

The Collector running state appears as Down on the Service Configuration page even though the Collector is up.

Solution:

The time is out of sync between the MOM and the Collector.

Synchronize the time between the MOM and the Collector.

For more information, see about Enterprise Manager clock synchronization (see page 102) in the *CA APM Configuration and Administration Guide*.

Symptom:

The Collector running state appears as Unknown.

Solution:

Verify that the host name and IP address of the Collector is correct. If it is not correct, edit the *introscope.enterprisemanager.clustering.login.em1.host* property in the *IntroscopeEnterpriseManager.properties* file in the *<EM_Home>**config* directory. Verify that the IP address is resolvable. Do not use localhost.

Defining your CA CEM domain

A CA CEM domain provides a way for the administrator to establish CA CEM-wide default values and settings for expected transaction behavior, user and business impacts, as well as CA CEM-wide data retention.

Your company needs to have one domain defined for CA CEM. Business services, business transactions, and business applications are created below the domain level. The default domain name is Local Domain.

As you create user groups, business services, and transactions, you have the option of inheriting the CA CEM domain default values or creating specific values (for example, SLA settings for users).

CA Technologies Pre-Release Documentation. Authorized Use Only.

To define your CA CEM domain:

1. Select Setup > Domain.
2. Define the Domain Settings:

   Domain Name

   Give a meaningful name to the CA CEM domain. The domain name is not case sensitive.

   Capture Comprehensive Defect Details

   If you select this, the CEM > Incident Management > Defect details page can display additional data about query and post parameters and the request and response bodies. See the *CA APM Security Guide* for more information.

   Troubleshoot Defects by IP Subnet

   If you select this, the incident troubleshooting chart Network tier shows the distribution of defects by originating subnet. (It does not re-group users by IP address.)

   Subnet Mask

   Use this feature with Troubleshoot Defects by IP Subnet to define the subnets.

   The mask is usually set to 255.255.255.0. However, for a large network, 255.255.0.0 might be more appropriate.

3. Define the User Group settings:

   Create User Groups by IP Subnet

   This feature is used to automatically create new user groups by the new user's client IP subnet.

   Important! If your installation is a large-scale e-commerce site (hundreds of thousands or millions of users) do not use this feature. The number of user groups created can become unmanageable. See Managing users and groups (see page 291) for more information.

   Subnet Mask

   Use this feature with Create User Groups by IP Subnet to define the subnets to create the user groups.

   The mask is usually set to 255.255.255.0. However, for a large network, 255.255.0.0 might be more appropriate. User groups can be further refined by creating user groups based on specific IP addresses and more restrictive subnet masks. See Managing the new users group (see page 298).

CA Technologies Proprietary Documentation. Unauthorized Use Only.

4. Define the Identification Settings:

Ignore Applications in User Recognition

If you do not select this option, then users are unique for each business application, with statistics history counted for *each* UserID (business application + login name). For example:

- ApplicationOne_MyLoginName = UserID
- ApplicationTwo_MyLoginName = another UserID
- ApplicationThree_MyLoginName = a third UserID
- eCommApplication_MyLoginName = a fourth UserID, but because it is an e-commerce business application, defects are kept; statistics go into unspecified users history.

If you do select the Ignore Applications in User Recognition option, then users are counted only once, even when they log in to multiple business applications.

Important! If you already have the original scenario (that is, one user with multiple UserIDs gathering statistics) and then select this option to ignore business applications, you will create another UserID:

- MyLoginName = UserID (for all future logins by this user, regardless of business application)

In this example, there would be statistics history for one user, kept in five buckets. The statistics history remains in the database as configured in the data retention rules.

Path Parameter Delimiters

If you use path parameters to identify users, sessions, user groups, or transactions, specify delimiters here. For more information about path parameters, see the *CA APM Transaction Definition Guide*.

CA Technologies Pre-release Documentation. Authorized Use Only

5. Define the Default Transaction Settings:

   a. Define the Impact Level, which determines the business priority of the defect type for the transaction.

      In most cases, the default value of Medium is sufficient. This value is inherited by the business services in the CA CEM domain by default, but can be edited at the transaction level.

   b. Define the transaction Service Level Agreement (SLA) levels.

      These values are inherited by the business services in the CA CEM domain by default, but can be edited at the transaction level. This table describes the CA CEM domain level SLA options.

      Success Rate SLA

      Success Rate Service Level Agreement as a percentage. The default is 95 percent defect-free transactions.

      Sigma SLA

      Sigma Service Level Agreement as a number. The Sigma score is a measure calculated from defects per million opportunities (DPMO). The default is 4.00, where near-perfection is 6.00 and perfection (zero transaction defects) is infinity.

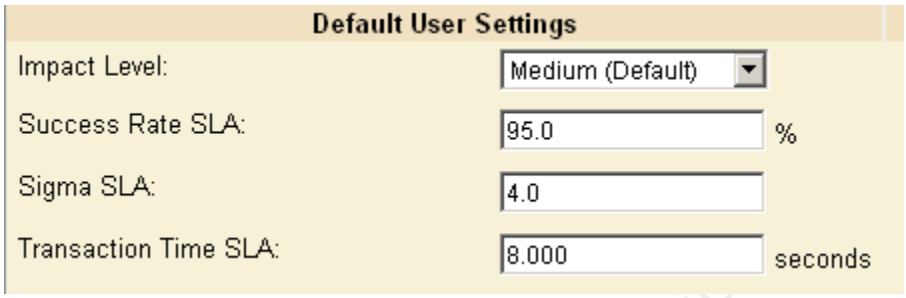
      Transaction Time SLA

      Transaction Time Service Level Agreement in seconds. The default is 8 seconds total elapsed time.

      Tip: Times in CA CEM can be specified to the 1/1000th of a second (0.001). If you specify less than that (for example, 0.0001), CA CEM sets the value to 0.000.

CA Technologies Pre-release Documentation. Authorized Use Only

6. Define the Default User Settings:



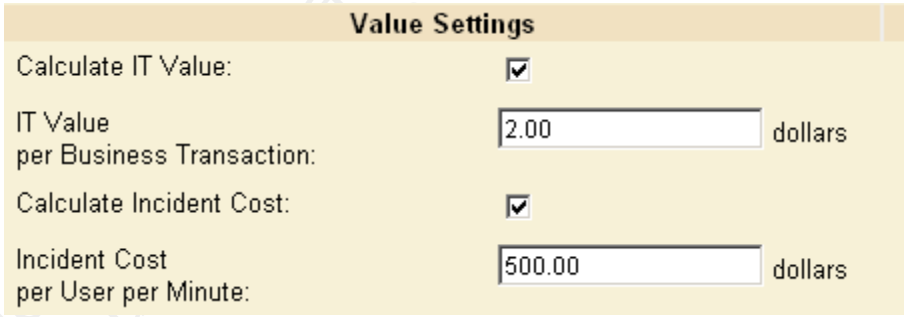
   a. Define the Impact Level, which determines the business priority of the defect type.

   In most cases, the default value of Medium is sufficient. This value is inherited by the business services in the CA CEM domain by default, but can be edited at the user group level.

   b. Define the default user Service Level Agreement (SLA) levels.

   These values are inherited by the business services in the CA CEM domain by default, but can be edited at the user group level.

7. Define the Value Settings.



   Performance reports (Business Services and Business Transactions pages) can display an IT value column (or not), which is calculated as:

   IT value = (IT value per business transaction) * (number of business transactions)

   For example, $2.00 per good business transaction observed; this means that every successful transaction contributes $2.00 to your business.

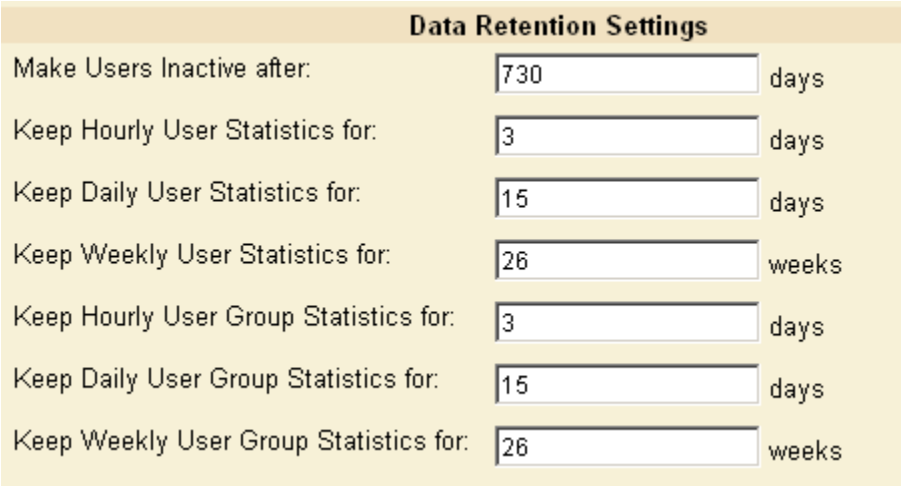
   a. To display IT value: Select the Calculate IT Value check box and type the IT Value per Business Transaction, in currency.

      If you do not select the check box, the IT Value column does not appear on Performance reports.

   b. For information about setting the Calculate Incident Cost check box, see Incident cost versus business impact (see page 268).

CA Technologies Product Documentation Only

8. Define the Data Retention Settings:



Make Users Inactive after

Users are active as long as they are generating transactions on the network. The default is to make users inactive when no transactions have been observed for them for the past 730 days.

Tip: Consider factors like the length of a vacation or a leave of absence, versus a termination, which can help you decide the appropriate number of days.

Statistics aggregation

Statistics are aggregated per hour per user, per user group, and for all users. You can select how long to keep user and user group statistics:

Keep Hourly User Statistics for

The default is to keep hourly user statistics for 3 days. After 3 days, you can see daily or weekly user statistics, but not hourly statistics prior to 3 days ago.

Keep Daily User Statistics for

The default is to keep daily user statistics for 15 days. After 15 days, you can still see weekly user statistics, but not daily statistics prior to 15 days ago.

Keep Weekly User Statistics for

The default is to keep weekly user statistics for 27 weeks. Weekly user statistics that are more than ~6 months old are discarded.

Keep Hourly User Group Statistics for

The default is to keep hourly user group statistics for 3 days. After 3 days, you can see daily or weekly user group statistics, but not hourly statistics prior to 3 days ago.

CA Technologies Internal / Prerelease Documentation
Authorized Use Only

Keep Daily User Group Statistics for

The default is to keep daily user group statistics for 15 days. After 15 days, you can still see weekly user group statistics, but not daily statistics prior to 15 days ago.

Keep Weekly User Group Statistics for

The default is to keep weekly user group statistics for 27 weeks. Weekly user group statistics that are more than ~6 months old are discarded.

Warning! If you set these data retention settings too high, daily aggregation may lag. If the IntroscopeEnterpriseManager log file shows that the Enterprise Manager is processing statistics records from a few days ago instead of from yesterday, then daily aggregation is lagging.

9. Click Save to keep all your settings.

Incident cost versus business impact

This topic explains the effect of selecting the Calculate Incident Cost check box on the Domain page (Setup>Domain).

The Incident Management > Incidents reports can display either an Incident Cost column or a Business Impact column. CA CEM defines incident cost as:

```
Incident cost = (cost-per-user-per-minute) * (number of users) * (number of minutes)
```

— where number of minutes is the incident's duration in minutes, and number of users is the number of unique users, not including any user who is categorized as an unspecified user.

Note: Enterprise transactions that occur before a user logs in and all e-commerce transactions are attributed to the unspecified users group. When CA CEM transactions are executed by unspecified users only, CA CEM displays a symbol in the User column, and the incident cost is $0.00 (cost-per-user-per-minute * 0 * number of minutes).

To display Incident Cost:

1. Select the Calculate Incident Cost check box.

   If you select this check box, the Business Impact column does not appear on the Incidents report; the Incident Cost column appears instead.

2. Type the Incident Cost per User per Minute, in currency.

   For example, $0.50; this means that if a business incident remains open for 5 users for 10 minutes, the cost to your business is $25.00.

CA Technologies Pre-Release Documentation. Authorized Use Only.

To display Business Impact:

- Clear the Calculate Incident Cost check box.

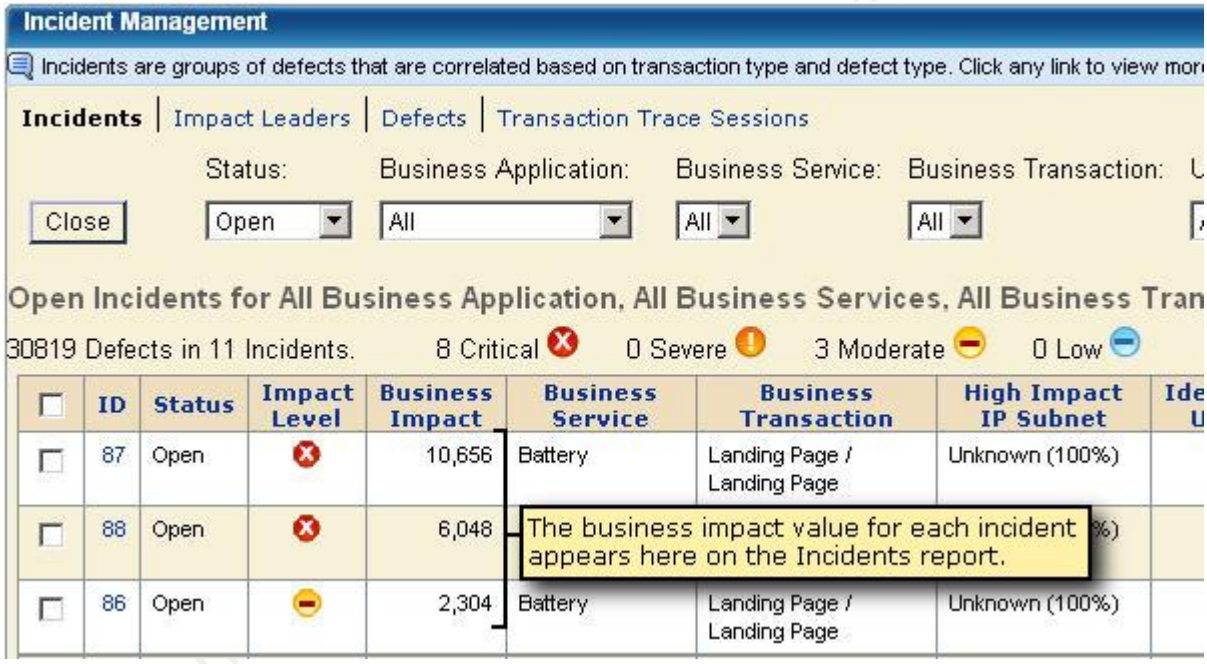
  If you do not select this check box, the Incident Cost column does not appear on the Incidents report; the Business Impact column appears instead.

  No entry is required in the Incident Cost per User per Minute text box.

Note: For information on business impact limits, see Defining incident related settings (see page 270).

Incidents report

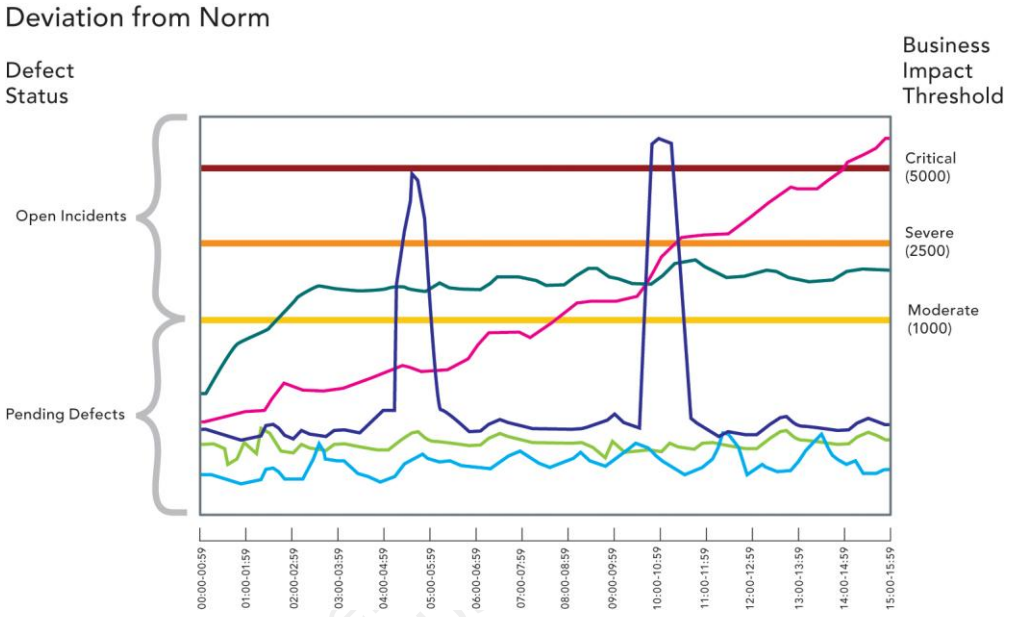
The incidents report looks like this:



Note: If you selected the Calculate Incident Cost check box on the Setup > Domain page, the Incident Cost appears on the Incidents report rather than the Business Impact.

Defining incident-related settings

In many environments, some number of defects is normal. But when there is a change from the norm, you need to know.



You can manipulate incident management settings, incident generating rules, and impact settings to generate incidents when there is a deviation from the norm.

Defect business impact

CA CEM assigns a business impact to each incoming defect, based on the defect's business transaction, defect type, and user.

A *defect's* business impact is calculated as follows:

Business transaction weight * defect type weight * user weight

An impact level is assigned to each:

- business transaction
- defect type
- user

This impact level …	Is assigned this weight …
Minimum	1
Very low	2

This impact level …	Is assigned this weight …
Low	3
Medium	4
High	5
Critical	6
Trigger immediately	7

Incident business impact

CA CEM aggregates incoming defects into incidents. Incidents are aggregated by business transaction and defect type. An *incident's* business impact is defined as the sum of the business impact of all its aggregated defects. This business impact appears in the Incident Management reports.

If business transaction, defect type, and user impacts are all medium (4), then 4 * 4 * 4 = 64 is the defect impact. This is default, shown in the graphic in Incidents report (see page 268). If you use these default business impacts, then:

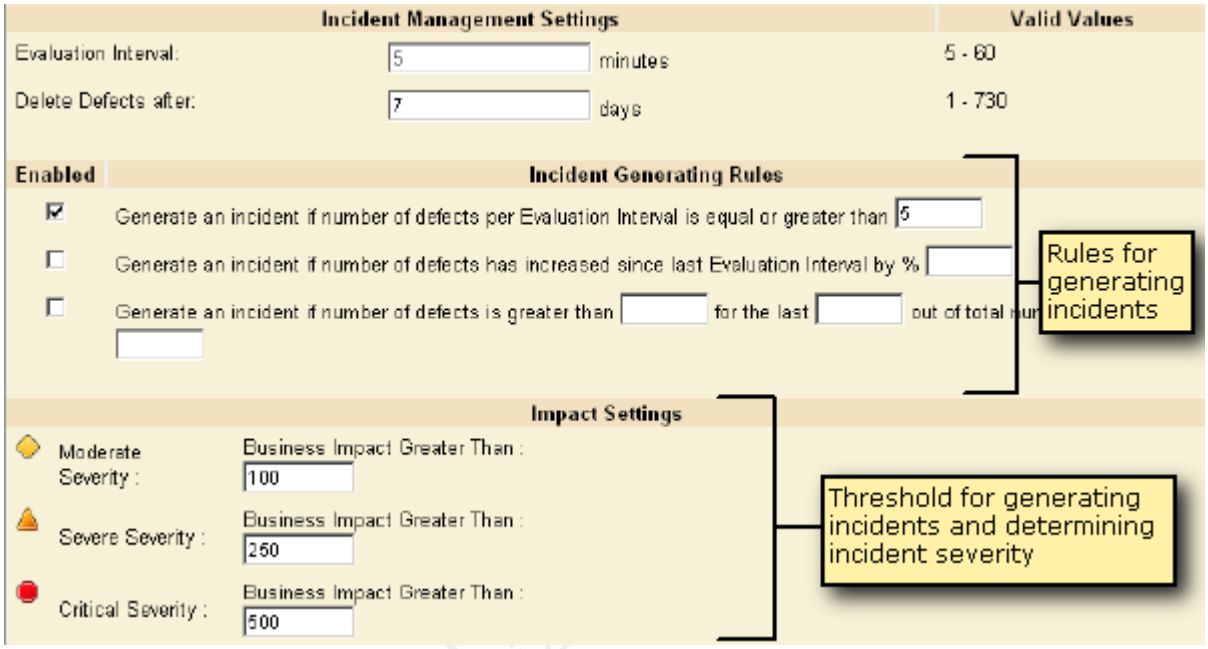
- 1000 for Moderate = 16 defects open a moderate incident
- 2500 for Severe = 40 escalates to a severe incident
- 5000 for Critical = 79 escalates to a critical incident

CA Technologies Preliminary Documentation. Authorized Use Only.

How incidents are generated

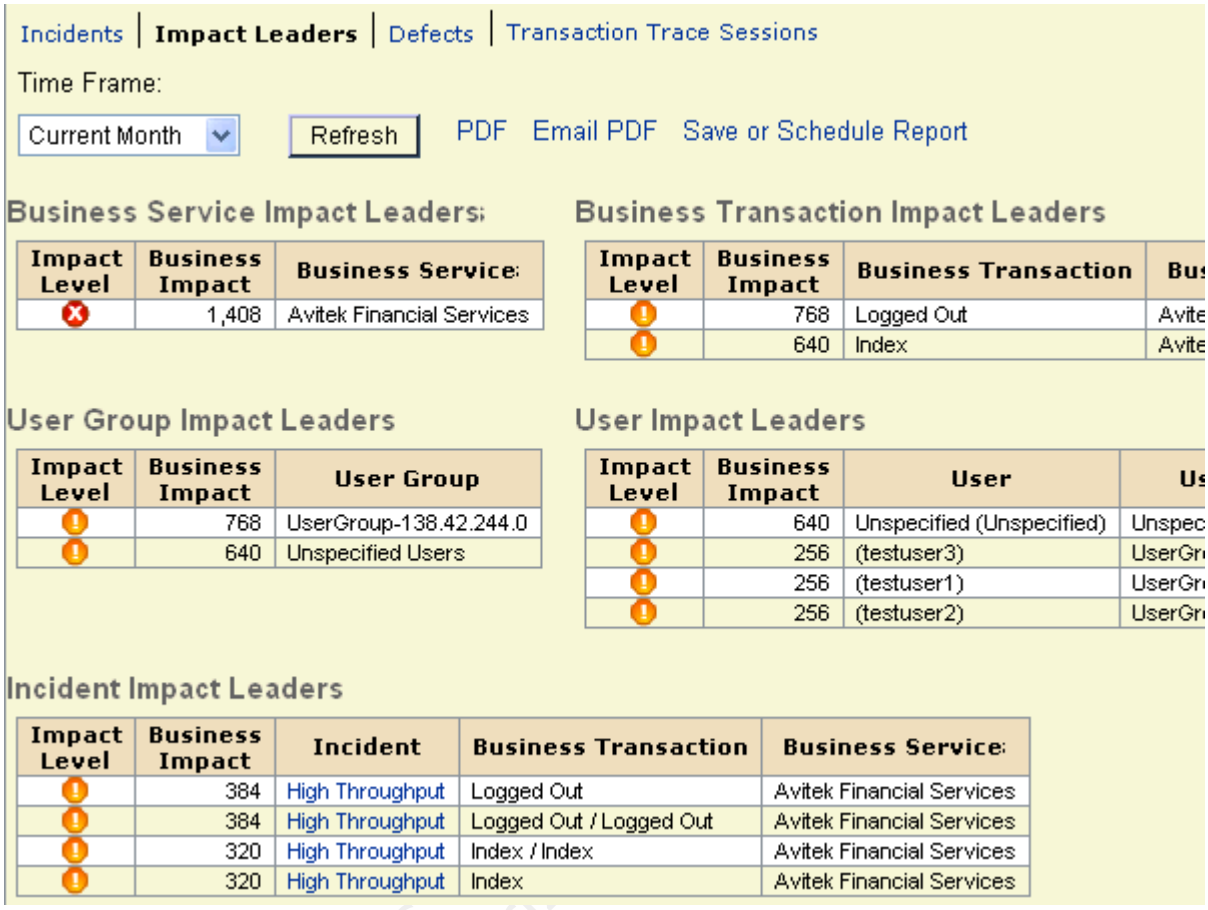
There are two ways that incidents are created:

- By having a business impact threshold met (this is the default method described above)
- By triggering the enabled incident generating rules



Impact leaders report

The Impact Leaders report summarizes the top five impacts, including business services, business transactions, users/groups, and incidents.



Incidents | **Impact Leaders** | Defects | Transaction Trace Sessions

Time Frame:

Current Month | Refresh | PDF | Email PDF | Save or Schedule Report

Business Service Impact Leaders:

Impact Level	Business Impact	Business Service:
	1,408	Avitek Financial Services

Business Transaction Impact Leaders

Impact Level	Business Impact	Business Transaction	Bus
	768	Logged Out	Avite
	640	Index	Avite

User Group Impact Leaders

Impact Level	Business Impact	User Group
	768	UserGroup-138.42.244.0
	640	Unspecified Users

User Impact Leaders

Impact Level	Business Impact	User	Us
	640	Unspecified (Unspecified)	Unspec
	256	(testuser3)	UserGr
	256	(testuser1)	UserGr
	256	(testuser2)	UserGr

Incident Impact Leaders

Impact Level	Business Impact	Incident	Business Transaction	Business Service:
	384	High Throughput	Logged Out	Avitek Financial Services
	384	High Throughput	Logged Out / Logged Out	Avitek Financial Services
	320	High Throughput	Index / Index	Avitek Financial Services
	320	High Throughput	Index	Avitek Financial Services

Defects report

The defect report shows the business impact of each defect.

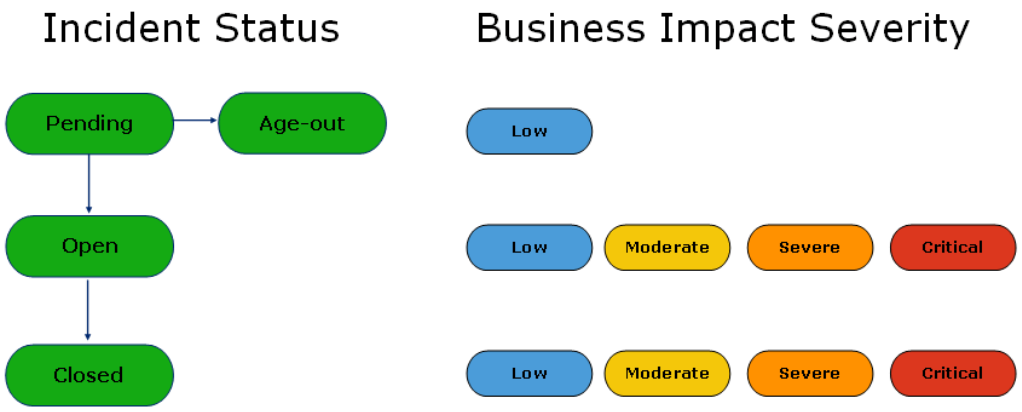
Business Transaction	Defect Name	Condition	Value	Business Impact	Login Name
Logged Out	High Throughput	100.0kB/s	193.8kB/s	64 (4 - 4 - 4)	testuser3
Logged Out / Logged Out	High Throughput	100.0kB/s	193.8kB/s	64 (4 - 4 - 4)	testuser3
Index / Index	High Throughput	100.0kB/s	190.4kB/s	64 (4 - 4 - 4)	Unspecified
Index	High Throughput	100.0kB/s	190.4kB/s	64 (4 - 4 - 4)	Unspecified
Logged Out	High Throughput	100.0kB/s	190.4kB/s	64 (4 - 4 - 4)	testuser2
Logged Out / Logged Out	High Throughput	100.0kB/s	190.4kB/s	64 (4 - 4 - 4)	testuser2

The business impact value appears here, also showing the weights (business transaction weight * defect type weight * user weight).

Incident status versus business impact severity

An incident can traverse several states, depending on its business impact value. Any CA CEM user can close incidents. The boundaries between each state can be defined on the Setup > Incident Settings page using the moderate, severe, and critical severity settings.

This chart compares incident status in comparison to business impact severity.



Configuring incident management settings

Your incident management configuration directly affects the incident generating rules and age-out rules and their logic.

To configure incident management settings:

1. Select Setup > Incident Settings.
2. Type an Evaluation Interval in minutes—the default value is 5 minutes.

   This value is used in the incident generating and age-out rules. For example, if you set it to 5 minutes, then every 5 minutes, each active (enabled) rule is reviewed to see if an incident should be generated or if pending incidents should be aged-out.
3. Type the number of days in Delete Defects after—the default is to delete defects after 7 days.
4. Click Save to keep all your settings.

CA Technologies Pre-release - Authorized Use

Setting incident generating rules

You can configure customized rules to cause incidents to be generated when it makes sense to your business.

CA CEM also generates incidents when severity thresholds set in Impact Settings are reached; see Setting business impact and collecting evidence for incidents (see page 276).

To configure incident generating rules:

1. Select Setup > Incident Settings.
2. Type the number of defects that would require an incident to be generated. This number means an immediate increase in the number of defects over a short period of time.

   For example, you might configure a rule for 1000 defects within 5 minutes (the example evaluation period) because that number of defects in that period of time would cause a big impact to your business.
3. Type the percentage of an increase in defects that would require an incident to be generated. This percentage means an immediate increase in the percentage of defects over a short period of time.
4. Type a combined rule:
   - The number of defects that would require an incident to be generated.
   - A number of evaluation intervals, for example 2 (of 3), as shown above.
   - Another number of evaluation intervals, for example (2 of) 3, as shown above.

     In this example, you might configure a rule for 500 defects in the last 2 of 3 evaluation intervals (or 10 of 15 minutes, if the evaluation interval is 5 minutes).
5. Select the Enable check box next to each rule that you want to activate.

   CA CEM checks every enabled rule based on the evaluation interval (for example, every 5 minutes).

   Note: Incident generating rules use AND logic. In other words, every enabled rule must be true before CA CEM generates an incident.
6. Click Save to keep all your settings.

CA Technologies Prerelease Documentation. Authorized Use Only.

Setting business impact and collecting evidence for incidents

Impact settings allow you to set thresholds that both generate incidents and determine incident severity levels. Evidence collection settings allow you to deploy advanced evidence collection scripts.

CA CEM can be configured to execute an evidence collection script each time the incident severity changes.

This script...	is executed when...
Collect evidence on Open	the incident is opened.
Collect evidence on Moderate Severity	the incident severity changes from Low to Moderate.
Collect evidence on Severe Severity	the incident severity changes from Moderate to Severe.
Collect evidence on Critical Severity	the incident severity changes from Severe to Critical.

Before you can collect evidence in CA CEM, you install Python or Perl on your system. Then, you modify the CA CEM property python.path or perl.path to specify the installation location of Python or Perl.

Note: This topic describes collecting evidence for incidents. A CA Technologies Professional Services offering allows you to collect and display evidence for individual defects. For information, see Setting the directory path for the Additional Defect Information link (see page 525).

To set up Python or Perl for use with scripts:

1. Install Python or Perl on the MOM or the standalone Enterprise Manager.
2. Edit the *tess-customer.properties* file, and modify the python.path or perl.path property.

   For example, *python.path=/usr/bin/python* or *perl.path=C:\\perl\\bin\\perl.exe*

   If a backslash is used for a Windows platform, a double slash (\\) is required because a single slash (\) is interpreted as an escape sequence.

   Note: If you do not already have a *tess-customer.properties* file, create one.

   See CA CEM default properties (see page 523) for more information.

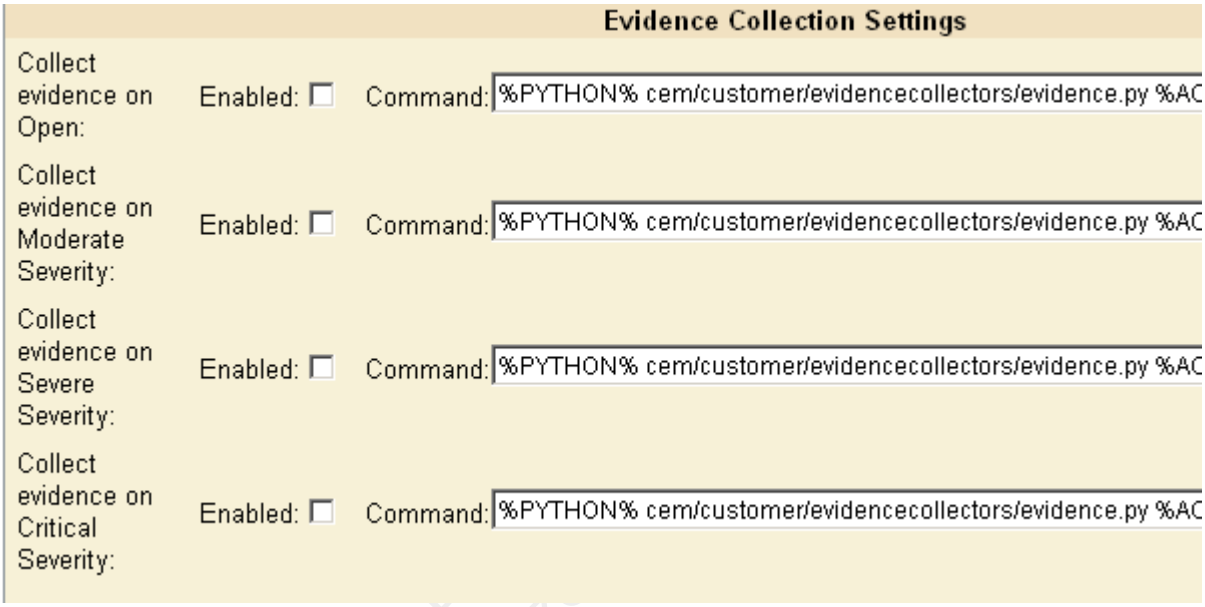
CA Technologies Internal Use Only. Authorized Documentation.

After you have set up Python or Perl, you can create and enable evidence collection scripts to run when an incident is opened or when its status changes.

To define the business impact levels for your incidents and set evidence collection scripts:

1. Select Setup > Incident Settings.

   The default settings appear.



2. Determine the impact settings for moderate, severe, or critical incidents.

   These settings determine both when an incident is created and when an incident severity is changed.

   The defaults are 1000 for Moderate, 2500 for Severe, and 5000 for Critical. See Defining incident-related settings (see page 270).

CA Technologies Proprietary Authorized

3. Determine the evidence collection settings. Select a check box if you want to collect evidence when an incident is opened, or on moderate, severe, or critical incidents. For example, the following scripts can be run when a moderate incident is detected:

   ```
   %PYTHON% C:\Introscope\cem\customer\evidencecollectors\evidence.py %ACTION%
   %SEVERITY% %DEFECTID% %TRANID% %TRANDEFID% %SERVER% %CLIENTIP% %OUTPUT%
   %PERL% cem/customer/evidencecollectors/evidence.pl %ACTION% %SEVERITY%
   %DEFECTID% %TRANID% %TRANDEFID% %SERVER% %CLIENTIP% %OUTPUT%
   ```

   For information about these arguments, see Arguments for use in evidence collection scripts (see page 278).

   The location you specify for the python.path property (step 2 of "To set up Python or Perl for use with scripts" above) is substituted for *%PYTHON%*. If you have installed Perl, use *%PERL%* instead of *%PYTHON%* in the command line.

4. Click Save to save your changes to the impact levels and evidence collection settings.

Arguments for use in evidence collection scripts

These are the arguments that can be used in evidence collection scripts. An example is shown in step 3 of "To set up Python or Perl for use with scripts" above.

ACTION

This is OPEN or CLOSE, depending on whether the incident is being opened or being closed.

SEVERITY

The business impact (Low, Moderate, Severe, or Critical).

DEFECTID

The database ID of the incident.

TRANID

The database ID of the business transaction.

TRANDEFID

The database ID of the business service.

SERVER

The IP address of the server where the most recent defect was detected. (The defect might be happening on more than one server. However, only the most recent server's IP address is captured in this argument.)

CA Technologies Prerelease Documentation
Authorized Use Only.

CLIENTIP

IP address of the client (i.e., your user's PC) where the most recent defect was detected. (The defect might be happening on many clients. However, only the most recent client's IP address is captured in this argument.)

OUTPUT

Where to put the output of the evidence collection script.

An example script file for evidence collection is installed with the product at this location:*cem/customer/evidencecollectors/evidence-example.py*
The script illustrates the use of the arguments.

Configuring incident age-out rules

Incident age-out rules allow you to get rid of low-level pending incidents.

You can configure customized rules to determine when pending incidents are aged out.

Pending incidents are defect events that have occurred, but have not reached the impact thresholds and are therefore not yet reported as incidents in CA CEM. For example, if a business impact threshold is 1000 and a business impact of only 999 has been observed, the incident is pending.

To configure pending incident age-out rules:

1. Select Setup > Incident Settings.
2. Type the number of minutes that must pass without related defects before a pending incident can be aged out. For example, for one day, type 1440 minutes (60 minutes * 24 hours = 1440).
3. Type the number of minutes that must pass with an incident in a pending state before a pending incident can be aged out.
4. Type a combined rule:
   - The number of defects that the pending incident must not yet have.
   - A number of evaluation intervals, for example 9 (of 10).
   - Another number of evaluation intervals, for example (9 of) 10.

   In this example, you might configure a rule for 500 defects in the last 9 of 10 evaluation intervals (or 45 of 50 minutes, if the evaluation interval is 5 minutes).

CA Technologies Proprietary Documentation. Authorized Use Only

5. Select the Enable check box next to each rule that you want to activate.

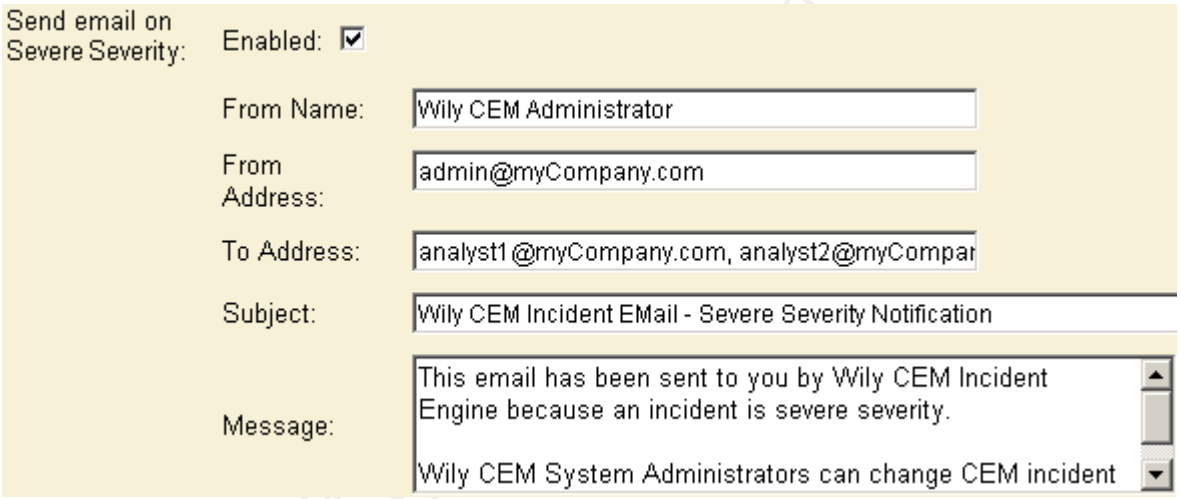
   CA CEM checks every enabled rule based on the evaluation interval (for example, every 5 minutes).

   Note: Incident age-out rules use OR logic. In other words, any enabled rule can cause CA CEM to age-out incidents.

6. Click Save to keep all your settings.

Sending emails when incidents open or change severity

You can configure CA CEM to send an email when an incident opens or its severity changes. For example, you might configure CA CEM to send an email to the incident analysts whenever an incident becomes critical.



To configure email to be sent when an incident opens or its severity changes:

1. Select Setup > Incident Settings.
2. In the Incident Email Settings section, edit the fields as necessary.

   If you want to send emails to more than one address, use a comma (,) to separate the email addresses.
3. Click Save to keep all your settings.
4. Ensure that CA CEM is configured to send email. See Configuring SMTP server settings (see page 281).

   CA CEM sends an email notification to the CEM administrator with the following information:

   - Incident ID
   - Business Application Impacted
   - Business Service Impacted
   - Business Transaction Impacted
   - Number of Affected Groups
   - Number of Affected Users
   - Number of Defects
   - Defect Name
   - Defect First Occured
   - Defect Last Occured

Configuring SMTP server settings

You must supply your SMTP server's name and port number so that CA CEM can send reports by email. You can also configure the subject and message text for the emails that accompany attached reports. For information about reports, see Reporting with CA CEM (see page 319).

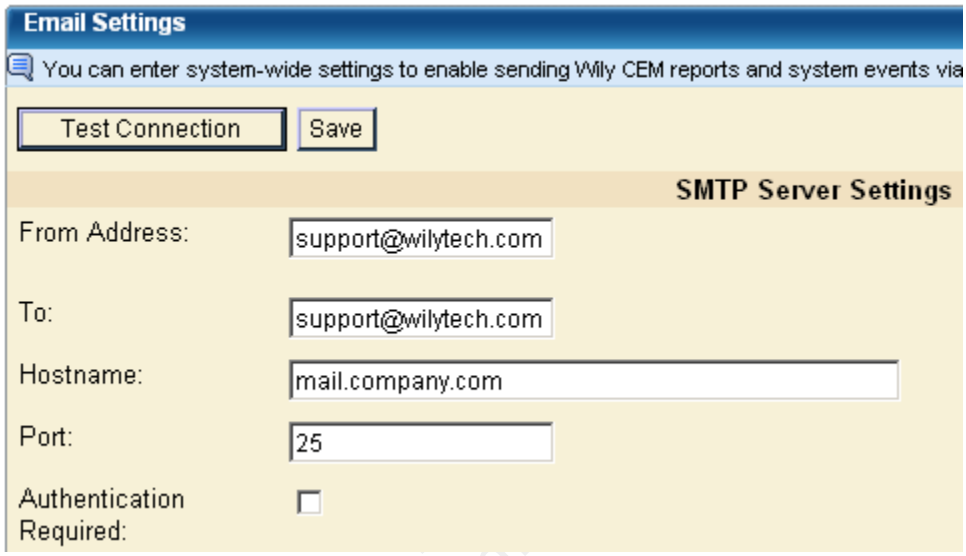
These SMTP setting are also used for email notification of CA CEM events and incidents. (See Receiving email notification of events (see page 312) and Sending emails when incidents open or change severity (see page 280).)

You must test the email connection before saving it.

CA Technologies Prerelease Documentation
Authorized Use Only.

Follow these steps:

1. Select System > Email Settings.



2. Define SMTP server settings:

Hostname

Enter the host name of the SMTP server.

Port

Enter the port number on the SMTP Server.

From Address

This is the address that the test-connection email will appear to be sent from.

To

This is the address that the test-connection email will be sent to. By default, this address and the From Address are taken from the CA CEM user (operator) details in Embedded Entitlements Manager. If this address is missing, the *email.defaultFromAddress* property in the tess-default.properties file is used; see email.defaultFromAddress (see page 534) to change this default.

Authentication Required

Select this box if the SMTP server requires a user name and password.

Username

If authentication is required, enter the user name.

Password

If authentication is required, enter the password.

CA Technologies Prerelease Authorized Use Only

3. Define email settings:

   Subject

   Enter a standard subject line for sending reports via email.

   Message

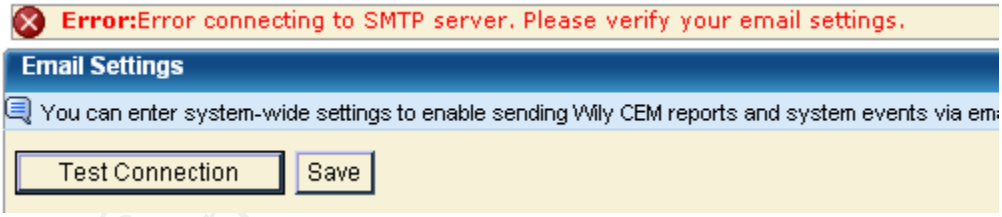
   Enter a standard message for sending reports via email.

   Footer Text

   Enter standard text to appear at the bottom of reports that are sent via email.

4. Click Test Connection. One of three things happens:

   - If a confirmation message appears and you receive an email indicating a successful test, then the SMTP settings are correct and the To email address is correct.
   - If a confirmation message appears but you do not receive an email, then the SMTP settings are correct but the To email address is incorrect.
   - An error message appears, then the SMTP settings are incorrect. (It might take time for the message to appear if you have to wait for the SMTP server to time out.)



5. Click Save.

Note: If you do not set up email for reports, CA CEM users receive error messages when they attempt to create reports.

For more information on types of reports and CA CEM user options for configuring and scheduling reports, see Reporting with CA CEM (see page 319).

Setting the security groups for new business services

If your organization has deployed Embedded Entitlements Manager, you can use access policies to control which security groups have access to a business service and its associated data.

A *security group* is a grouping of users created for the Embedded Entitlements Manager APM application. In the Embedded Entitlements Manager, security groups are called *application groups*.

CA Technologies Product Documentation: Author Use Only

The administrator configures which security groups get automatic access to all newly created business services.

By default, four access policies have been set up:

Business Service All Permissions

CEM System Administrator group has all permissions to all business services.

Business Service read and write

CEM Configuration Administrator group has read and write permissions for all business services.

Business Service read sensitive data

CEM Incident Analyst group has permission to read sensitive data for all business services.

Business Service read

CEM Analyst and CEM Incident Analyst groups have permission to read all business services.

If you are using standard security groups (CEM System Administrator, CEM Configuration Administrator, CEM Incident Analyst, CEM Analyst), then these access policies might be sufficient. However, if you have created additional security groups and you want the members of these groups to have access to all business services, you need to create or modify access policies to include these groups.

Note: The procedures in this topic describe how to modify or create access policies for *all* business services. If you want to modify or create access policies for *specific* business services, see the topics about access policies the *CA APM Transaction Definition Guide*.

To view or edit an existing access policy for all business services:

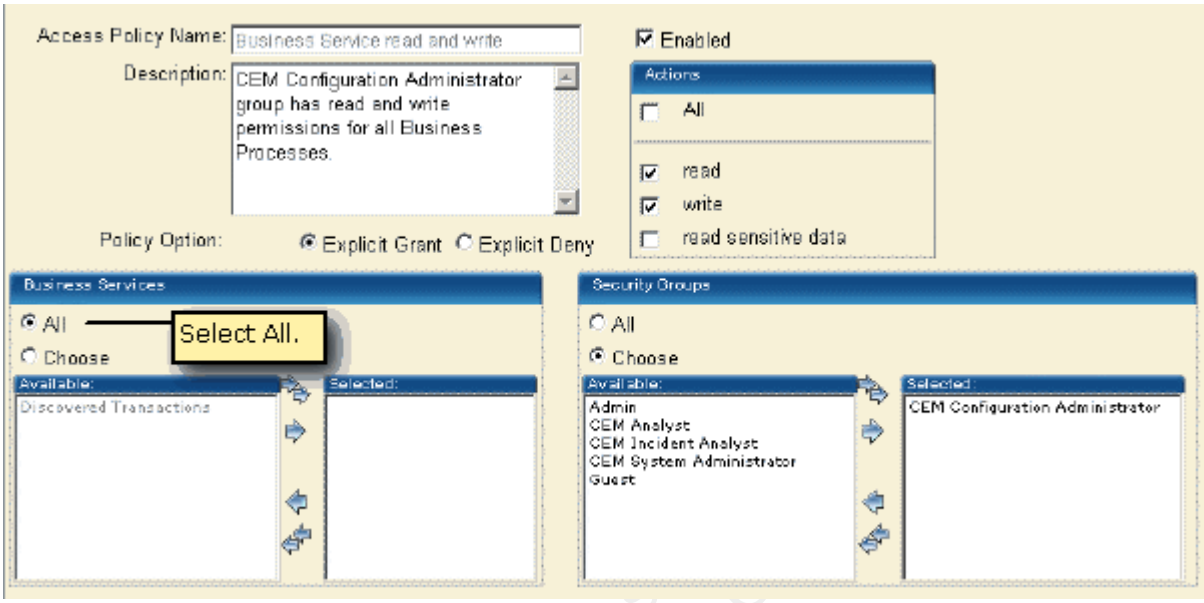
1. Select Security > Access Policies.

   If no access policies have been added since CA CEM was installed, you see the access policies listed in the table above.

CA Technologies Proprietary and Confidential Documentation. Authorized Use Only.

2. Click the name of a policy to view or edit it.

   The Update Access Policy page appears.



3. Modify fields on the Update Access Policy page and click Save.

   The fields are described in the list below.

   Ensure that you select All for Business Services. This means that the selected security groups are associated with all subsequently created business services and applied to all existing business services.

To create an new access policy for all business services:

1. Select Security > Access Policies.
2. Click New.
3. Complete the fields on the New Access Policy page and click Save.

   The fields are described in the list below.

   Note: Ensure that you select All for Business Services and Choose for Security Groups. This means that the selected security groups are associated with all business services (existing and future).

The fields in the New Access Policy page are described here.

Access Policy Name

Give a meaningful name to the policy.

Description

Add a descriptive comment.

CA Technologies Proprietary Authorized Use Only

Policy Option

Choose to create a policy to explicitly grant or deny actions for the business service.

If there is a conflict between grant and deny, the deny policy takes precedence.

Enabled

Enable a policy if you want it to be in effect.

Disable a policy if you don't want to use it now, but might want to use in the future. Delete policies that you don't want to use again.

Write

Select this to allow the security group members to view, edit, and delete the business service and all its child objects (business transactions, transactions, components, and so on) from the Business Service tab.

Users who do not have access to at least one business service, do not see the Business Service tab.

Read

Select this to allow the security group members to view the graphs and reports associated with that business service in the CEM tab.

Note: Ensure that Embedded Entitlements Manager is the only realm configured. If the local realm is configured (in addition to the Embedded Entitlements Manager realm), then all users can access all data in *scheduled* reports. For information on realms, see the *CA APM Security Guide*.

Read sensitive data

Select this to allow the security group members to see the Host, URL path, TCP port, cookies, and HTTP headers associated with a defect. These appear on the CEM > Incident Management > Defect details page.

If the Capture Comprehensive Defect Details check box (on the Setup > Domain page) is selected, then query and post parameters and request and response body information for a defect can also appear on the Defect details page. For more information, see the CA CEM chapter in the *CA APM Security Guide* for information.

Use this option with the read check box. The read action is required to access the defect details page.

Business Services

Select All. This applies the access policy to *all* business services and to *all* new business services when they are created.

Note: To modify the access policies for *specific* business services, see the *CA APM Transaction Definition Guide* for details. For example, if you might want to deny the analysts access to one particular business service by setting an explicit deny for that business service for the CEM Analyst and the CEM Incident Analyst security groups.

CA Technologies Proprietary Documentation. Authorized Use Only.

Security Groups

Select Choose and move the security groups for the policy to the Selected column.

The security groups shown here correspond to the APM application-specific user groups in Embedded Entitlements Manager. Embedded Entitlements Manager global user groups cannot be selected in CA CEM.

Use Embedded Entitlements Manager to view or modify the security groups. See the *CA APM Security Guide* for information about using Embedded Entitlements Manager. Security groups cannot be managed from the CEM console.

Enabling monitors

You need to make sure that the TIMs are enabled so that transactions can be monitored. The Enterprise Manager does not receive data from disabled TIM monitors.

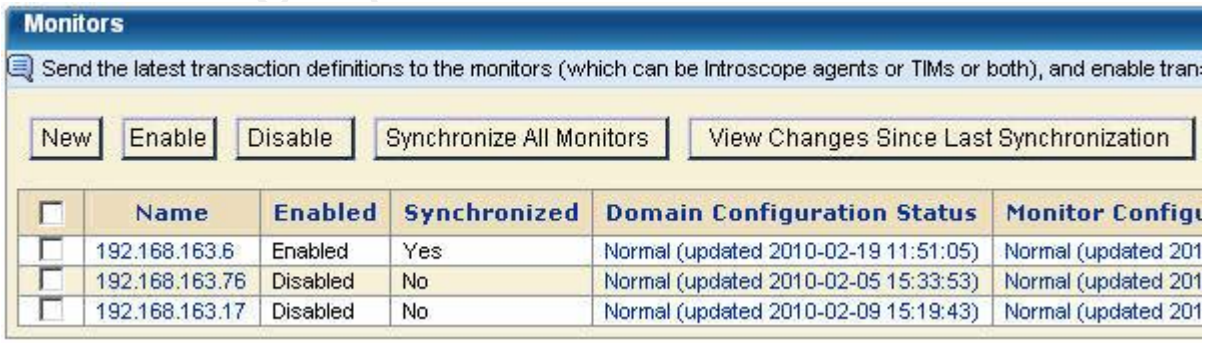
These are some examples of when you need to enable TIMs:

- When setting up a TIM and associating it with an Enterprise Manager. (See the *CA APM Installation and Upgrade Guide*.)
- After upgrading the TIMs and/or Enterprise Manager. (See the *CA APM Installation and Upgrade Guide*.)
- After a configuration is imported. (See Importing the APM database configuration (see page 71).)

To enable TIM communication with an Enterprise Manager:

1. Select Setup > Monitors.
2. Select the check box next to each required TIM and click Enable.

   Communication between the Enterprise Manager and the TIMs is then enabled.



Synchronizing monitors

Synchronizing monitors sends the current domain configuration file to all enabled TIMs and to the agents. The domain configuration file includes transaction definitions, defect specifications, and, Introscope transaction trace parameters.

These are some examples of when you need to synchronize monitors:

- After creating or updating business services in CA CEM. Changes to business services and transaction in CA CEM do not appear in the Workstation or WebView Investigator tree until monitors have been synchronized and some metrics for the changed business services have been generated.
- When setting up a TIM and associating it with an Enterprise Manager. (See the *CA APM Installation and Upgrade Guide*.)
- After upgrading the TIMs and/or Enterprise Manager. (See the *CA APM Installation and Upgrade Guide*.)
- After creating or modifying transaction definitions, to send the updated definitions to the agents and to enabled TIMs. (See the *CA APM Transaction Definition Guide*.)
- After importing a configuration file, to send the imported definitions to the agents and enabled TIMs. (See Importing the APM database configuration (see page 71).)

To synchronize the monitors on your network:

1. Select Setup > Monitors.
2. If any required monitors are disabled, enable them. (See Enabling monitors (see page 287).)
3. If you want to review the changes that have been made to configuration since the last synchronization, then see Viewing configuration changes since last synchronization (see page 289).
4. Click Synchronize All Monitors.

   All the enabled monitors are synchronized and the related icon appears in the upper right. This sends the updated transaction definitions to all enabled TIMs on your network. The TIMs immediately start monitoring the transactions you have defined.

   Synchronizing all monitors also causes CA CEM to send transaction and tracing information to the Enterprise Manager.



CA Technologies Proprietary Documentation
Authorized Use Only

Viewing configuration changes since last synchronization

You can list the changes that have been made to the configuration before you synchronize. This allows you to review the changes before synchronizing.

Any changes that have been made to these properties (since the last synchronization), appear in the Domain Configuration History table:

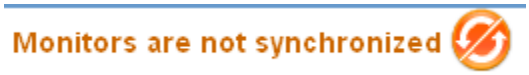
- business application definitions
- business service definitions
- defect definitions
- business transaction definitions
- defect importance definitions
- transaction definitions
- defect impact levels
- user identification parameters
- transaction parameter definitions
- session identification parameters
- business transaction definitions
- interim session identification parameters
- transaction importance definitions
- user group identification parameters
- transaction component definitions
- private parameters

If you delete certain types of parameters, such as a private parameter, the object name appears blank and the property changes display as "Nothing found to display." This is because the object name and changes have been deleted from the database.

CA Technologies Prerelease Documentation. Authorized Use Only.

To view configuration changes made since last synchronization:

1. Confirm that monitors are not synchronized.

2. Select Setup > Monitors.

3. Click View Changes Since Last Synchronization.

   The change history table appears.

Note: The change history table lists the configuration changes made since the last synchronization. This is not the same as listing the differences between the last synchronization and the current configuration. For example, if you changed a value from 5 to 2 and back to 5; this results in two rows in the table—not none, as you might expect.

By default, the maximum number of rows in Domain Configuration History table is 1,000. To change this, see Limiting the number of changes listed since last synchronization (see page 456).

Not all property values listed are meaningful to you. For example, the Config Manager value, displayed when a parameter is created, shows the internal class name for the object.

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 19: Managing users and groups on CA CEM

This list shows the main tasks involved in managing users and groups:

1. Learn about the types of user groups available with CA CEM (see page 291).
2. Learn about automatically creating user groups via information within the transactions (see page 294).
3. Learn about automatically creating user groups via IP subnet (see page 295).
4. Reassign users from the new users group to other groups (see page 298).
5. Configure user groups and SLA (service level agreement) settings (see page 299).
6. Manage users and user groups with bulk editing (see page 301).
7. Correlate user groups SLAs with business transactions (see page 302).
8. Delete a user group (see page 305) or deactivate or delete an individual user (see page 306).

Note: The *users* who are discussed in this chapter are the users of the *monitored application*. For example, in a banking application, the users are the bank's clients, signing in to view and manage their accounts. These users might be grouped according to their account type.

However, if you are looking for information about grouping CA CEM users into *security user groups*, for example, to configure the tabs, menu items, and business services the CA CEM users can access on the CEM console, see the *CA APM Security Guide*.

CA Technologies Pre-release Documentation
Authorized Use Only

About user groups

User groups provide the convenience of configuring settings for a collection of users instead of having to configure the settings for each individual user separately. User groups can be defined so that you can easily identify user populations that might be experiencing problems.

New user groups can be created manually, and users can be moved manually into their user groups. CA CEM also offers ways to automatically create user groups. A user belongs to only one CA CEM user group.

There are multiple types of user groups:

Request attribute user groups

Users grouped by the content within the request they are making.

Subnet user groups

Groups that have automatic assignment of new users by user subnet.

Manually created user groups

Users must be manually moved from New Users group (or another user group) into this type of user group.

New Users

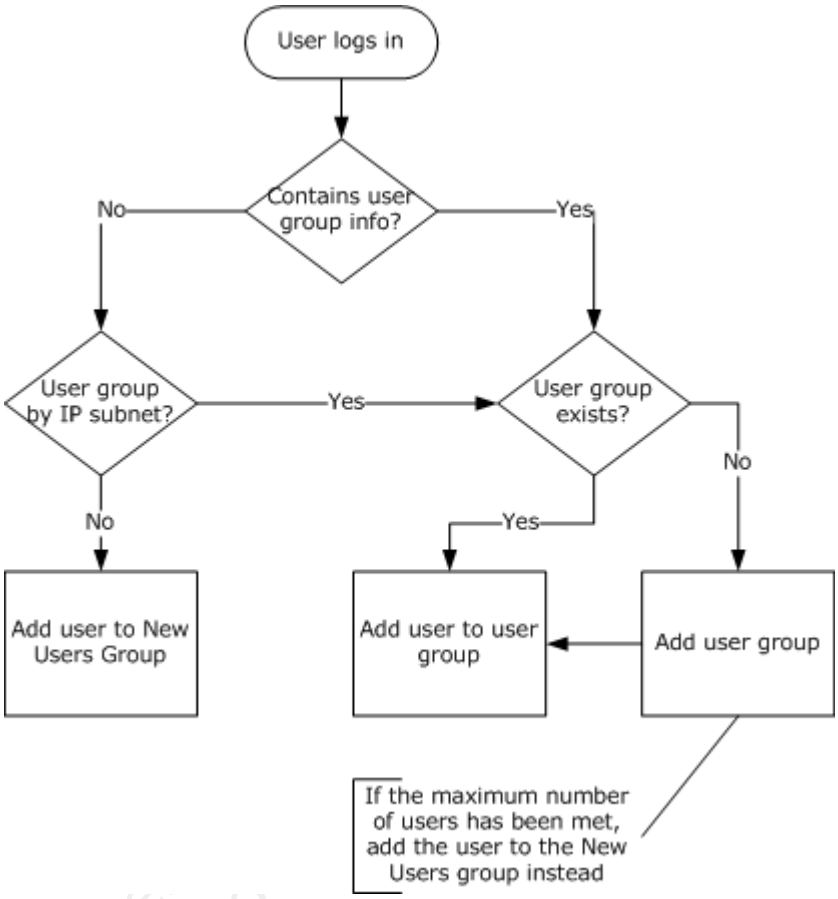
When no automatic user group generation method is selected, all new user logins are assigned to the new users group.

Unspecified Users

All e-commerce transactions are assigned to the unspecified users group. Also enterprise transactions that occur before a user logs in are assigned to the unspecified users group.

CA Technologies Prerelease Documentation.
Authorized Use Only

This flow illustrates the logic behind automatic user group creation and assignment.



The maximum number of user groups

CA CEM provides the ability to control the number of automatically created user groups. This can be helpful in unanticipated situations where the number of groups can increase dramatically before the CEM administrator realizes.

By default, the maximum number of automatically created user groups is 5000. If you need to adjust the default value, see The Default tess-default.properties File (see page 523).

Automatically assigning new users by request attributes

Users can be grouped by the content within the request they are making. The most common cases are:

- URL strings — for example, property code or customer code in query string
- HTTP headers — for example, users of different phone types or different gateways
- CA APM TG headers — for example, synthetic users using CA APM TG (x-wtg-info) header

Using this method, CA CEM can report performance information by user location or company. For example, this can help you manage outages by affected populations. User group statistics are available for e-commerce business applications.

Note: Grouping by user group parameter takes precedence over automatic user group classification by subnet.

Important! If your installation is a large-scale e-commerce site (hundreds of thousands or millions of users) use caution with this feature. The number of user groups created can become unmanageable.

Important! If you are identifying users by client IP address instead of user name, do not assign users by request attribute. See the *CA APM Transaction Definition Guide* for information about using the client IP address to identify users.

To assign new users by request attribute:

1. Select Administration > Business Applications.
2. Select the business application under which you want to identify users by request attribute.
3. Select User Group Identification.

   A list of all existing user group parameters appears.
4. Click New to create a new user group parameter.
5. Select the appropriate parameter Type.

CA Technologies Pre-Release Documentation
Authorized Use Only

6. Type the parameter Name.
7. Click Save.

   For example, the x-wtg-info type groups CA APM TG users in a group called Synthetic. This helps to distinguish actual traffic from CA APM TG traffic.



Automatically assigning new users to subnet user groups

Users can be grouped via the user's IP subnet. CA CEM automatically compares a new user's IP address with the subnets of existing user groups. The subnet mask that is most restrictive is examined first. For example, the following user groups are defined:

```
SubGroup – 255.255.255.0 mask / 192.168.2.0 IP start
BigGroup – 255.255.0.0 mask / 192.168.0.0 IP start
```

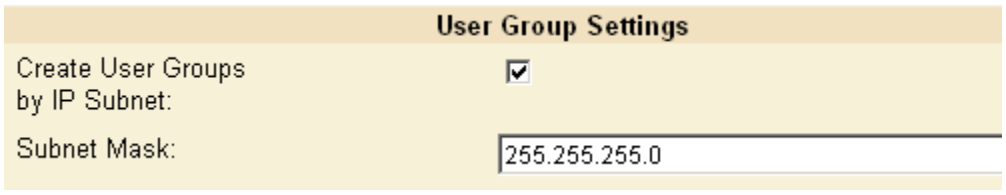
Incoming IP traffic is compared to the SubGroup user group first because its subnet mask is more restrictive. For example, the IP address 192.168.2.4 would be assigned to the SubGroup group and the IP address 192.168.9.9 would be assigned to the BigGroup group.

If an IP address is detected that is not in either the SubGroup or BigGroup groups, the user is assigned according to your configuration choices. For example, an IP address of 192.1.9.9 has a subnet mask of 255.255.255.0, but would not fit into either group's IP address range.

Note: Grouping by user group parameter takes precedence over automatic user group classification by subnet. See Automatically assigning new users by request attributes (see page 294) for more information.

Creating subnet user groups for new users

Subnet user groups are determined by the Create User Groups by IP Subnet option on the Setup > Domain tab:



If you selected the Create User Groups by IP Subnet option on the Setup > Domain page, a new IP address triggers the creation of a new user group rather than being assigned to the group named New Users.

Important! If your installation is a large-scale e-commerce site (hundreds of thousands or millions of users), do not use this feature. The number of user groups created can become unmanageable.

By default, the maximum number of automatically created user groups is 5000. If you want to adjust the default value, see The Default tess-default.properties File (see page 523).

The new user group name is the starting number in the IP address detected. In our example, the new user group would be named:

Usergroup-192.1.9.0 — 255.255.255.0 mask/192.1.9.0 address

To view the newly created user group:

1. Select Administration > User Groups.
2. Click on the Usergroup-192.1.9.0 to see new users listed.

Note: Users that have been placed in a subnet user group using one IP address remain in that user group even if their IP address changes later. You can manually move the original user entry if necessary.

Identifying users behind proxy servers

In the scenario described in Automatically assigning new users to subnet user groups (see page 295), users who are behind a proxy server are grouped according to subnet of their proxy servers and not of their client machines. Depending on your business need, this might not be what you want. You might want to have the user groups created according the client subnets.

Similarly, for the defect list (Incident Management > Defects page) you might want to have the client IP address show the originating machine IP address rather than the IP address of the proxy server.

CA Technologies Preferred Documentation
Unauthorized Use Only.

If the request traffic includes an HTTP header that contains the client IP address, you can configure CA CEM *so that the client IP address is reported for logins and defects instead of the connection IP address.* X-Forwarded-For, Forwarded-For, and iv-remote-address are examples of headers used to identify client IP addresses. Other headers can be used if they contain an IP address in dotted-decimal notation (or an IPv6 address in a standard notation). For example:

```
X-Forwarded-For: 172.24.192.4
```

The TIM setting called ProxyForwardHeader is used to specify the name of the header.

Note: In cases where the header is missing or its value is not in the form of an IP address, then the connection IP address is used. That is, the behavior is the same as the default, where the ProxyForwardHeader setting is unspecified.

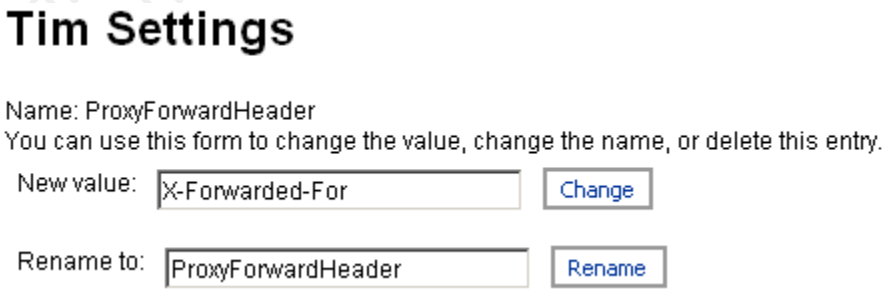
For each TIM, set the value of the ProxyForwardHeader to a header that contains the IP address of the client.

To identify clients behind proxy servers based on the value of a header:

1. Access the TIM System Setup page. (Accessing the TIM from the CEM console (see page 234).)
2. Click Configure TIM Settings.

   The TIM Settings page appears.
3. Click ProxyForwardHeader.
4. In the New value field, enter the header name (without the colon), and click Change.

   This field is not case sensitive.



5. If you have multiple TIMs, repeat the above steps for each TIM. The setting affects all servers and applications monitored by each TIM for which it is defined.

CA Technologies Prerelease Documentation
Authorized Use Only

To create user groups based on the value of a header:

1. Set the ProxyForwardHeader as described in "To identify clients behind proxy servers based on the value of a header" above.
2. On the CEM console, make sure that the option to create user groups by IP subnet is selected:
   a. Select Setup > Domain.
   b. Select Create User Groups by IP Subnet.

   This means new user groups is created based on the client subnet, provided in the header, instead of the connection subnet. Also, defects are reported as coming from the client IP address, instead of the connection IP address.
3. (Optional) Test the configuration by creating a defect from a client behind the proxy server. If the IP address associated with the defect is the client's and not the proxy server's, then the configuration is correct.

Managing the new users group

If you did not set up one of the automatic user group creation and assignment methods, then any new users that do not match your existing subnet user groups is automatically assigned to the New Users group.

To view new users:

1. Select Administration > User Groups.
2. Click on the New Users group to see new users listed.

   New users that have been detected and have not been assigned to a user group appear.

Note: You do not have any new users detected until you start monitoring transactions.

Manually assigning new users to user groups

New users are automatically placed in a subnet user group or in the New Users group.

To move users out of their current group:

1. Select Administration > User Groups.
2. Click New to define a new user group.
   a. Give a meaningful name to the user group (for example, London). The user group name is not case sensitive.
   b. Type in a description of the group.

CA Technologies Preferred Use Documentation. Authorized Use Only.

c. Leave the other parameters as is for now, or see Configuring user groups (see page 299) for more information.

d. Click Save to save the new group.

3. Click on the user group you just created. You do not have any new users detected until you start monitoring transactions, or until you move users into the group.
4. Go back to the User Groups page and click on the New Users group.
5. To move a user (or users) to the new user group, select the check box (or boxes) to the left of the Login Name(s).
6. Select your new user group from the Move to list (for example, London).
7. Click Move to move the user(s) to your new user group.
8. If you want to move every user from the current user group to the selected user group in the list, click Move All. (You do not need to click the check boxes to the left of each Login Name.)
9. When you go back to the User Groups page, click on your new user group (for example, London) to see the moved user(s).

Configuring user groups

To edit user groups on your network and define SLAs:

1. Select Administration > User Groups.

   The current user groups appear.
2. Click the name of a User Group to modify the group's settings.
3. Click the General tab.
4. Change the user group name if needed. The user group name is not case sensitive.
5. Type in or modify a description for the group.
6. Select the Impact Level, which determines the business priority of the defect type for the group's transactions. In most cases, the default value of Medium is sufficient.

   The default is to inherit the Impact Level from the CA CEM domain. If you want to change this value, clear the check box and select the Impact Level from the list.

   Note: To change the CA CEM domain default values, select Setup > Domain as described in Defining your CA CEM domain (see page 262).

CA Technologies Prerelease Documentation
Authorized Use Only

7. Select the Service Level Agreement (SLA) values that define the SLA at the business service level.

   The default is to inherit the SLA values from the CA CEM domain. If you want to change a value, clear the check box and select a new value from the list.

   These values are also used to define default SLA values that can be inherited by the user group's transactions within the business service.

   Success Rate SLA

   Success Rate Service Level Agreement as a percentage. The default is 95 percent.

   Sigma SLA

   Sigma Service Level Agreement as a number. The Sigma score is a measure calculated from defects per million opportunities (DPMO). The default is 4.00, where near-perfection is 6.00 and perfection (zero transaction defects) is infinity.

   Transaction Time SLA

   Transaction Time Service Level Agreement in seconds. There is no default value.

8. Select the Group New Users by IP Subnet check box if you want to automatically add new users detected by login ID to a new user group that is identified by the subnet name.

   a. Type the IP address of the user group (for example, 192.168.1.0).

   b. Type the default subnet mask (for example, 255.255.255.0).

   In this example, any new user that logs in using an IP address from 192.168.1.1 to 192.168.1.254 is automatically assigned to this user group rather than the New Users group.

   Tip: If you are going to use the IP address method for assigning users to groups, it is a good idea to include IP information in the group name.

9. Click Save to save your new user group. It appears in the User Group list.

CA Technologies Pre-Release Documentation. Authorized Use Only.

Configuring user settings with bulk editing

In some cases, you might need to make widespread changes to your user settings. For example, perhaps you want to seek out several users, but not an entire user group, and modify the Sigma SLA setting from 4.00 to 5.00.

Bulk editing allows you to make widespread changes for:

- business service and transaction settings
- transaction and component defect conditions
- user settings

If you want to modify transaction-related settings on a large scale, see the *CA APM Transaction Definition Guide*.

The User Search page allows you to modify user settings on a large scale.

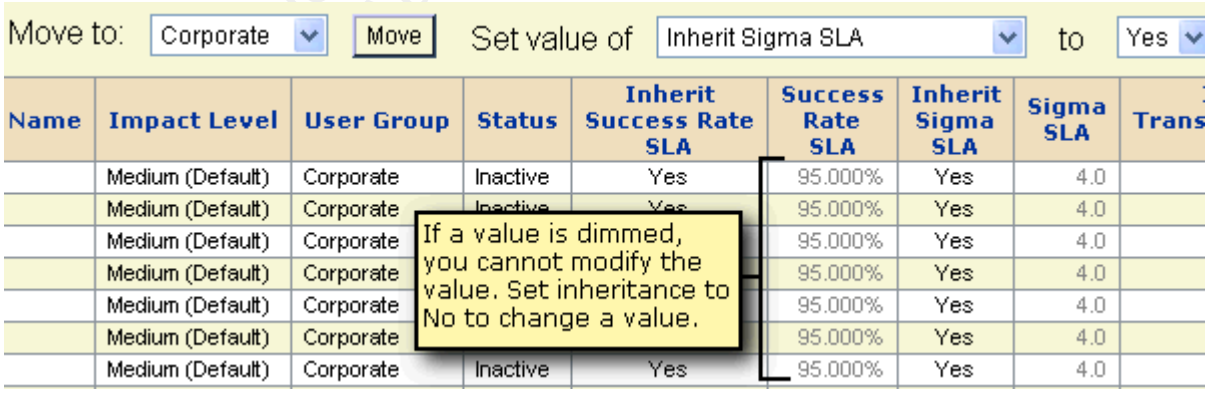
To configure user settings with bulk editing:

1. Select Administration > User Groups > User Search.

   The search options appear.

2. Select the name of the user group (for example, Corporate).
3. Select the Impact Level (for example, All).
4. If you are searching for a particular user, you can optionally type the Login name, First Name, or Last name.
5. Click Search.

   The results of your search appear.



Name	Impact Level	User Group	Status	Inherit Success Rate SLA	Success Rate SLA	Inherit Sigma SLA	Sigma SLA	Trans
	Medium (Default)	Corporate	Inactive	Yes	95.000%	Yes	4.0	
	Medium (Default)	Corporate			95.000%	Yes	4.0	
	Medium (Default)	Corporate			95.000%	Yes	4.0	
	Medium (Default)	Corporate			95.000%	Yes	4.0	
	Medium (Default)	Corporate			95.000%	Yes	4.0	
	Medium (Default)	Corporate			95.000%	Yes	4.0	
	Medium (Default)	Corporate	Inactive	Yes	95.000%	Yes	4.0	

6. Select the check box to the left of the user(s) that you want to change. Or, if you want to change the entire list, click the top check box.

Broadcom Prerelease Documentation. Authorized Use Only.

7. Select an item from the Set value of list:
   - Inherit Sigma SLA
   - Sigma SLA
   - Inherit Success Rate SLA
   - Success Rate SLA
   - Inherit Transaction Time SLA
   - Transaction Time SLA

   A value text box or Yes|No list appears, depending on your choice.
8. Type a value, or select Yes or No, and click Set.

Configuring correlational SLAs

In some cases, you might want to set a user group's SLAs by business service or by business transaction, rather than using one SLA for the whole user group or an SLA for a transaction (independent of user group).

The table below is an example of varying acceptable response times, based on the transaction and the user location.

User Group / Business Transaction Acceptable Response times	Local	Remote	Overseas
Login	3 seconds	6 seconds	9 seconds
Buy	5 seconds	7 seconds	10 seconds

Correlational SLAs allow you to set SLAs by user group, by business service, and by business transaction combined. A correlational SLA can be viewed as a multi-dimensional alternative to the single-dimensional SLAs set for the business transaction or for the user group.

Note: Correlational SLAs do not replace domain, business transaction, or user-group SLAs. Correlational SLAs are not inherited from any of these—they must be set specifically by the CEM administrator.

These are the procedures for configuring correlational SLAs:

- Viewing correlational SLAs (see page 303)
- Adding correlational SLAs (see page 303)
- Updating existing correlational SLAs (see page 304)
- Deleting correlational SLAs (see page 305)

CA Technologies Pre-release Documentation. Authorized Use Only

Viewing correlational SLAs

Follow this procedure to review which correlational SLAs have already been created or to find and select SLAs for deletion.

You cannot view all correlational SLAs in one table. This is because, in a typical production system, the matrix of all user groups, business services, and business transactions is prohibitively large. However, you *can* view correlational SLAs for a given business service or for a given user group.

To view correlational SLAs by user group and business transaction:

1. Select Administration > Correlational SLAs.
2. In the Show field, select By User Group or By Business Service.
3. In the second field, select a group or service.
4. In the third field, select All or select a group or service.
5. Click Refresh.

   The correlational SLAs for the specified user group or business service appear.

Adding correlational SLAs

This procedure describes how to create new correlational SLAs for a user group/ business service/business transaction set.

To add a correlational SLA:

1. Select Administration > Correlational SLAs.
2. Ensure that the correlational SLA you want to create does not already exist. If it does exist and you want to change it, then see Updating existing correlational SLAs (see page 304).
3. Click New.
4. Define the correlational settings per user group/business service/business transaction:
   a. User Group — select the name of the user group or All.
   b. Business Service — select the name of the business service or All.

CA Technologies Prerelease Documentation. Authorized Use Only.

c. Business Transaction — select the name of the business transaction or All.

d. Set the SLA values.

Success Rate SLA

Success Rate Service Level Agreement as a percentage. There is no default value.

Sigma SLA

Sigma Service Level Agreement as a number. The Sigma score is a measure calculated from defects per million opportunities (DPMO). An acceptable Sigma score might be 4.00, where near-perfection is 6.00 and perfection (zero transaction defects) is infinity. There is no default value.

Transaction Time SLA

Transaction Time Service Level Agreement in seconds. There is no default value.

5. Click Save.

   The new correlational SLA appears in the list.

Note: Existing correlational SLAs are not changed when new ones are created. For example, if you already have a correlational SLA for user group A (UG-A), business service X (BP-x), and business transaction 3 (BT-3), then creating a new business service for UG-A, PB-x, and all business transactions does not change the previously defined SLA (for UG-A/BT-x/BT-3).

Note: If you create correlational SLAs for all business transactions in a business service and later add another business transaction, remember to create a new SLA for the added business transaction. (An SLA for this business transaction is not added automatically.)

Updating existing correlational SLAs

To change existing correlational SLAs use the Update button.

To change an existing correlational SLA:

1. Select Administration > Correlational SLAs.
2. Select the correlational SLA that you want to modify. (Use the Show, Business Service, and User Group lists and Refresh to display the correlational SLA that you want to change.)
3. Click Update.
4. Set the Business Transaction fields to specify the correlational SLA that you want to change.

CA Technologies Proprietary/Trade Secret Document. Authorized Use Only.

5. Set the SLA values.
6. Click Update.

   The updated correlational SLA appears in the list.

Deleting correlational SLAs

If you no longer need a correlational SLA, you can delete it.

To delete correlational SLAs:

WARNING! When you click Delete, the SLA is immediately deleted. There is no "are you sure" dialog associated with this deletion.

1. Select Administration > Correlational SLAs.
2. Display and select the correlational SLAs that you want to delete.
3. Click Delete.

Deleting a user group

You can delete user groups and all their related transaction statistics. You might need to do this:

- When an office closes and you no longer need transaction information for that group
- When you want to merge groups: move the users from one group to another and then delete the emptied group
- When automatic creation of user groups has lead to more user groups than you need

Important! Delete user groups that you do not need, including user groups that have no users. Too many user groups can have a negative impact on performance.

To delete a user group:

1. Select Administration > User Groups.
2. On the User Groups tab, select the check box for the user group that you want to delete.
3. Click Delete.

CA Technologies Prerelease Documentation. Authorized Use Only.

Deactivating or deleting a user

You can temporarily deactivate users if you know they will not be monitored by CA CEM for an extended period of time. In this case, transaction history and the user records are preserved. If CA CEM observes a new transaction from any deactivated user, then the user is reactivated automatically.

When you delete users, only the users' transaction history is preserved. You do not see deleted users on the administration pages. The users' history appears dimmed in reports. If CA CEM observes a new transaction from any deleted user, then a new user is created. You should delete users only when you are sure they will not use the application you are monitoring ever again.

To deactivate or delete a user:

1. Select Administration > User Groups.
2. Click User Search and search for the specific user you need to deactivate or delete.
3. Check the user that you want to deactivate or delete.
4. Click Deactivate or Delete.

Reactivating a user

If you deactivated a user in the past, and now want to reactivate that user:

To activate a user:

1. Select Administration > User Groups.
2. Click User Search and search for the specific user you need to activate.
3. Check the user that you want to activate.
4. Click Activate.

CA Technologies Prerelease Documentation. Authorized Use Only

Chapter 20: Managing system operations on CA CEM

This chapter describes the tasks the administrator performs to ensure the CA CEM is performing correctly, such as reviewing error reports and checking disk space and memory.

This list shows the things you need to know about CA CEM monitoring events and viewing audit trails for CA CEM:

1. Learn about CA CEM events (see page 307).
2. Review CA CEM events (see page 311).
3. Maintain the required CA CEM events (see page 311).
4. Configure CA CEM to send email notification events such as errors and warnings (see page 312).
5. Review information such as memory use and disk space for TIM machines (see page 314).
6. Review CA CEM audit trails (see page 316).

CA Technologies Prerelease Documentation
Authorized Use Only

About CA CEM events

CA CEM captures and logs events so you can stay informed about events that can affect system operations and performance. It is a good idea to review the events log on a regular basis.

Severity levels

There are three severity levels for events:

- informational
- warning
- error

List of possible events

CA CEM captures and logs these system events:

Event	Severity	Category	Details
Business application name is not valid	Error	CA APM TG	The business application name in CA APM TG script contains an invalid character such as < or >. In CA APM TG, correct the CA CEM Application Name and migrate the script again.
Database Cleanup Service has started	Informational	Processing	Database Cleanup Service has started. The IP address of the Enterprise Manager running the service is given in the event list.
Database Cleanup Service has stopped	Informational	Processing	Database Cleanup Service has stopped. The IP address of the Enterprise Manager that was running the service is given in the event list.
Defects dropped	Warning	Processing	One or more defects were dropped due to capacity processing overload, reported by incident.
Duplicate session ids detected	Warning	Monitoring	Possible causes: Two or more TIMs are monitoring the same traffic. More than one login has the same session ID. (Might indicate a problem with the parameter used for session identification.)
Enterprise Manager cannot reach TIM	Error	Communications	The Enterprise Manager was unable to reach the TIM to retrieve data.
Enterprise Manager restarted	Informational	Administration	Restart date and time for the Enterprise Manager.
HTTP analyzer plug-in error	Error	Processing	An HTTP analyzer plug-in has failed.

CA Technologies Prerelease Documentation. Authorized Use Only.

Event	Severity	Category	Details
Business application name is not valid	Error	CA APM TG	The business application name in CA APM TG script contains an invalid character such as < or >. In CA APM TG, correct the CA CEM Application Name and migrate the script again.
Limit exceeded adding new components	Warning	Processing	This event can only occur when Discover new non-identifying components is selected. When the maximum number of non-identifying components is exceeded, the new component replaces the oldest. Each time this happens, the Limit exceeded adding new components event is logged. See *CA APM Transaction Definition Guide* for information about automatically discovering components.
Logins dropped	Warning	Processing	One or more logins were dropped due to capacity processing overload.
No network data	Warning	Communications	TIM has not received any data over the network for a period of time. For example, perhaps a cable was moved, or the network switch was not configured properly, or the network topology has changed and no longer matches the CA CEM configuration.
Stats Aggregation Service has started	Informational	Processing	Stats Aggregation Service has started. The IP address of the Enterprise Manager running the service is given in the event list.
Stats Aggregation Service has stopped	Informational	Processing	Stats Aggregation Service has stopped. The IP address of the Enterprise Manager that was running the service is given in the event list.
Enterprise Manager deleted some data files because the lag-time threshold was met	Informational	Administration	The Enterprise Manager deleted some old data files (defect, automatic transaction discovery, and customer experience metrics). This happened because the lag-time threshold (downloadfiles.minutes.collect) was met.
Enterprise Manager deleted some TIM data files because the lag-time threshold was met	Informational	Administration	Some old data files (defect, automatic transaction discovery, and customer experience metrics) were deleted from one or more TIMs. This happened because the lag-time threshold (downloadfiles.minutes.collect) was met.

CA Technologies Prerelease Documentation Authorized Use Only.

Event	Severity	Category	Details
Business application name is not valid	Error	CA APM TG	The business application name in CA APM TG script contains an invalid character such as < or >. In CA APM TG, correct the CA CEM Application Name and migrate the script again.
TIM Collection Service has started	Informational	Processing	TIM Collection Service has started. The IP address of the Enterprise Manager running the service is given in the event list.
TIM Collection Service has stopped	Informational	Processing	TIM Collection Service has stopped. The IP address of the Enterprise Manager that was running the service is given in the event list.
TIM deleted some of its data files because disk space was low	Informational	Administration	A TIM deleted some of its data files to protect disk space. By default, files are deleted when the disk is 75% full. This behavior is determined by DiskSpace/MinFreeDataSpaceInPercent on the TIM Settings page.
TIM restarted	Informational	Administration	Restart date and time for the TIM.
Transaction definitions dropped	Informational	Processing	This event can only occur when automatic transaction discovery is on. It means that some automatically discovered transaction definitions were dropped due to capacity processing overload. If you see this event, consider turning off automatic transaction discovery and running it at another time.
User group limit reached	Warning	Processing	The maximum number of automatically created user groups by IP subnet has been reached, reported by user group.
WTG agent is down	Warning	CA APM TG	A CA APM TG agent is not available.
WTG agent is up	Informational	CA APM TG	A CA APM TG agent is available.

CA Technologies Prerelease Documentation
Authorized Use Only

Viewing CA CEM events

You can monitor CA CEM events periodically, looking for signs of network problems, system or backup failures, and HTTP analyzer plug-in errors.

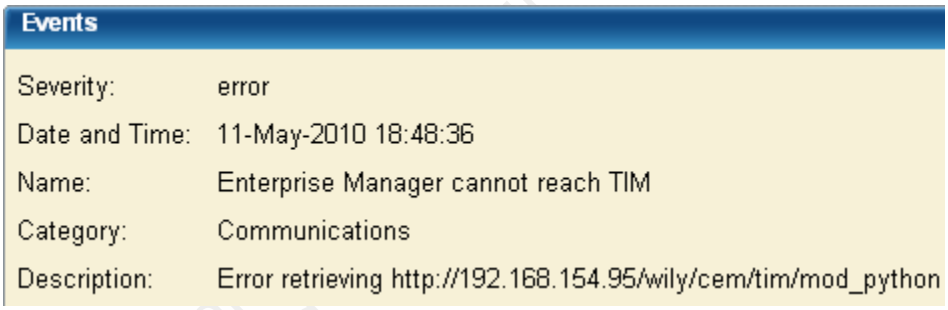
Follow these steps:

1. Select System > Events.

   The list of CA CEM events appears.

2. Click on the name of an event to see more information.

   The event details page appears.



3. Select System > Events to return to the list of events.

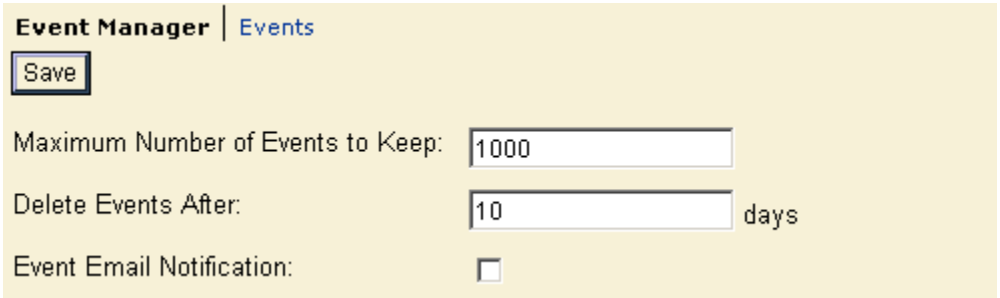
Managing CA CEM events

You can specify the number of events to retain in the CA CEM event log and the number of days to keep events.

Follow these steps:

1. Select System > Events, and then click on the Event Manager tab.

   The current settings for CA CEM events appear.



2. Define the Event settings:

   Maximum Number of Events to Keep

   The CA CEM Event Manager keeps events in the log up to this maximum. The recommended maximum is 1000.

   Delete Events After

   Number of days after which CA CEM events can be deleted from the event log.

3. Click Save.

Receiving email notification of events

You can configure CA CEM to send an email notification when events occur.

At a minimum, you should set up event email notification so that a CA CEM administrator receives email notification of error events.

You can specify:

- which email addresses to send the notifications to.
- which events to be notified about.
- the subject and body text of the message.

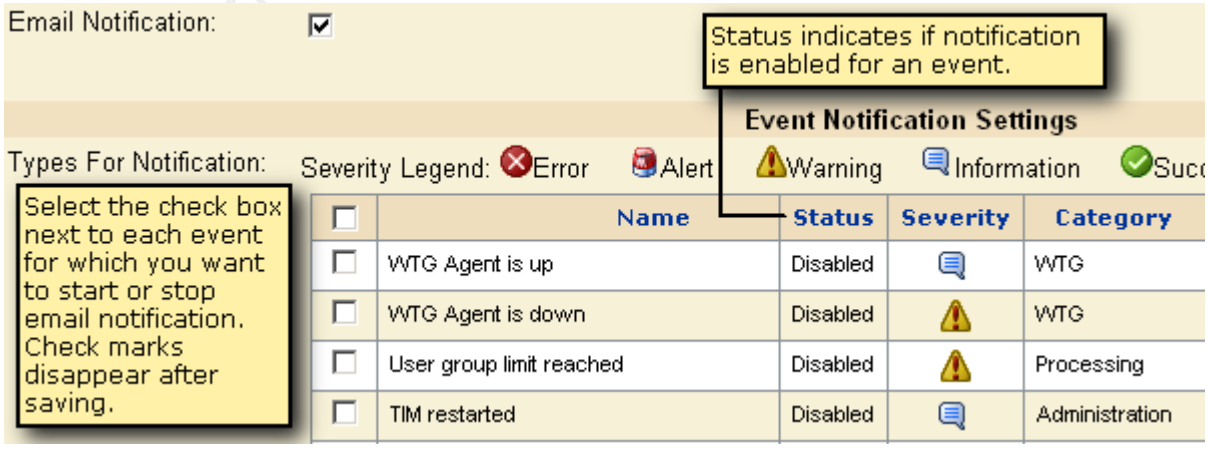
To configure CA CEM to send email notification of events:

1. Select System > Events, and then click on the Event Manager tab.

   The current settings for CA CEM events appear.

2. Select the Event Email Notification check box.

   The Event Notification Settings appear.



Knowledge Preparation Documentation. Authorized Use Only.

3. Select the check box next to each event you want to enable (or disable) for email notification.

   The check box is a toggle option. If the status is currently Disabled, selecting the check box and clicking Save starts email notification for that event (and the status becomes Enabled). If the status is currently Enabled, selecting the check box and clicking Save stops email notification for that event (and the status becomes Disabled).

   Note: The status field indicates whether *email* notification for the event is enabled or disabled. All events, regardless of their email notification status, appear on the Events page.

4. (Optional) Specify an IP address range if you want to receive email notifications from only a subset of the Enterprise Managers and TIMs.

   For example, if you want to see email notifications about the Enterprise Managers but not the TIMs, specify an IP address range that includes the Enterprise Manager address but excludes the TIM addresses.

5. Edit the From and To addresses for the email if necessary.

6. (Optional) Edit the Message field. Use the variables to create a customized description:

   - <id>
   - <name>
   - <severity>
   - <category>
   - <source>
   - <ipaddress>
   - <desc>
   - <time>

   You can rearrange the variables to create your own message. Each variable is replaced with actual event values when the email notification is sent.

   Tip: You can copy and paste the variables and text from the Examples area to the message field.

7. Click Save.

8. Ensure that SMTP server settings are correct. See Configuring SMTP server settings (see page 281).

CA Technologies Prerelease Documentation
Authorized Use Only

Reviewing TIM machine health

For routine maintenance and when troubleshooting, you should check the health of the TIM machines. This includes reviewing disk and memory usage, the processes running on the machine.

Note: You need access to the Wily System Setup pages to see these health reports.

To see a snapshot of the health of a TIM machine:

1. Select Setup > Monitors.
2. Click Get Report (in the TIM Health column on the far right).
3. If you are presented with a login dialog, log in with the user name and password for the Wily System Setup pages; the default user name for this is admin.

   The Appliance Health page appears.



4. Click the Complete Appliance Health link.

   Alternatively, click one of the other links to see a single section of the health report.

   See the Appliance Health pages reference table (see page 315) for information about each link.
5. Repeat the above steps for each TIM machine.

CA Technologies - Authorized Documentation. Only.

Appliance Health pages reference table

This table describes the information presented for each of the links on the Appliance Health page.

Link	Information	Linux command equivalent
Disk Information	Displays the file system disk space usage in MB. If disk usage gets too high, above 80%, reduce disk space usage by deleting TIM log files.	`df -m`
Memory Information	Displays the amount of free and used memory in the system in MB. If the amount of free memory gets low, then you need to beware of excessive memory swapping (thrashing). On a TIM, the most likely cause of thrashing is a plug-in. You might need to modify the plug-in code to improve its memory usage.	`free –m`
Logged in Users	Displays any users who are logged in to the machine itself. Typically, you should not see any users. However, you might see a root user who has logged in to the machine, for example, through a secured shell (ssh). If you see any users who you do not recognize, investigate immediately.	`who`
Currently Running Processes	Displays every process running. (Displays using standard syntax and using some ASCII art in the CMD column to indicate process paths.) In this list, look for unusually high CPU usage and for unexpected processes.	`ps –efl --forest`
Complete Appliance Health	Displays all of the above.	see above

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

Viewing CA CEM Audit Trails

CEM administrators can view an audit trail of CA CEM activity.

You can view Wily CEM activity on administration detail pages. Audit information is not available on list pages.

To view audit trail information:

1. Select one of the administration pages to audit.
   - Administration > Business Services
   - Setup > Domain
   - System > Email Settings
2. In the address, add **&***showHistory=true* or ?*showHistory=true* to the end of the URL. (Use either an ampersand (&) or a question mark (?).)
   - If the URL ends in something like *.html?pId=600001*, use an ampersand:

   `http://<IP_Address>:8081/wily/cem/tess/app/admin/tranDefGroupList.html?pId=1&showHistory=true`
   `http://<IP_Address>:8081/wily/cem/tess/app/admin/domainDetail.html?pId=1&showHistory=true`
   - If the URL ends in *.html*, use a question mark:

   `http://<IP_Address>:8081/wily/cem/tess/app/system/emailSettings.html?showHistory=true`
3. Click Go or press Enter to see the audit trail.

   The audit trail appears at the bottom of the page.

Date and Time	User Name	Action	Property Changes

Name	Old Value	New Value
Calculate Incident Cost	false	true
Capture Comprehensive Defect Details	false	true
Defect Filter Subnet Mask		255.255.255.0
Filter Defects by IP Subnet Enabled	false	true
IT Value	0.00	0.01
Ignore Apps In User Recognition	false	true

Technology Preview Documentation. Authorized Use Only

4. Scroll down on the page to see Wily CEM audit information. The last 200 changes appear.
5. If you want to override the default of displaying 200 records, you can add *&maxResults=nnn* to the end of the query string.

   For example:

   ```
   http://<IP_Address>:8081/wily/cem/tess/app/system/emailSettings.html?showHistory=true&maxResults=500
   ```

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 21: Reporting with CA CEM

This list shows the things you need to know about CA CEM reporting:

1. Learn about CA CEM reports and how to capture them as you navigate the system (see page 319).
2. Schedule the CA CEMM reports you need on a regular basis (see page 325).

Capturing CA CEM reports

You can analyze your CA CEM data in several ways. You can navigate through the CA CEM menus and perform queries based on business service, business transaction, user group, as well as by incident and by defect.

The report options vary, depending on the report type you choose. For example, if you choose a trend report, then you need to specify the primary and secondary time frames to compare in the trend analysis.

See Report delivery (see page 320) for information about report output alternatives.

See Report formats (see page 321) for information about available reports.

Your CA CEM system data retention choices and statistics aggregation schedule can affect report output. See Report time frames (see page 322) for more information.

About CA CEM Report Aggregation

CA CEM aggregation affects the data that appears on your reports. You should understand how CA CEM aggregates daily, weekly, monthly, and yearly statistics.

CA CEM reports are based on aggregated statistics of your business transactions. Once per hour, CA CEM runs the statistics processing which includes hourly aggregation.

Once per day—at midnight, *after* the hourly statistics aggregation process finishes—CA CEM runs:

- daily aggregation process
- weekly aggregation process
- monthly aggregation process

CA Technologies Prerelease Documentation
Authorized Use Only

The once-per-day aggregation is the best method for performance reasons, however you need to be aware of the impact on reports. For example:

- If you run a weekly or monthly report at 5 p.m. today, the report statistics include data as of midnight; the report does not include today's statistics.
- If your Enterprise Manager that runs the Stats Aggregation Service is not running at midnight, the aggregation processes do not run (until the next time the Enterprise Manager is running at midnight). So if you are looking for yesterday's statistics, they are not be included in the daily (Yesterday time frame), weekly, monthly, or yearly reports.
- Report statistics from the Today time frame take slightly longer than the aggregated daily data.
- Yearly reports are based on monthly aggregation data.

Report delivery

CA CEM allows you to capture and manage CA CEM reports in a variety of ways:

- View a report online as you query CA CEM data
- View a predefined report from the My Reports list
- Create a PDF file, which you can save or print
- Email a PDF file to a user or to a group of users
- Save a report, which you can run and view later
- Schedule reports to run and to be delivered via email on a daily, weekly, or monthly basis

Securing report data

If you have security groups to limit CA CEM users' access to business services and associated data, then verify that you have only the Embedded Entitlements Manager security realm.

If there are multiple realms, for example, a local realm and an Embedded Entitlements Manager realm, then it is possible for CA CEM users to see data in scheduled reports that they do not have access rights to.

For information about configuring realms, see the *CA APM Security Guide*.

CA Technologies Prerelease Documentation
Authorized Use Only

Report formats

You can capture CA CEM reports as you navigate these pages on the CEM tab:

- Service Level Management
  - transaction SLAs — performance and quality
  - user SLAs — performance and quality
  - correlational SLAs
  - transaction trends — performance and quality
  - user trends — performance and quality
  - business value
- Incident Management
  - incidents
  - impact leaders
  - defects
  - transaction trace sessions
- Performance Reports
  - business applications, business services, business transactions
  - user groups
- Quality Reports
  - business applications, business services, business transactions, transactions
  - user groups
- Analysis Graphs
  - success rate
  - time
  - throughput
  - size
  - volume
  - count
  - yield
  - defects — frequency
  - session usage

CA Technologies Prerelease Documentation. Authorized Use Only.

Report time frames

CA CEM allows you to specify the report view time frame in various ways:

Time Frame	Current	Previous	Custom
Hour	N/A	Previous hour	Custom hour
Day	Today	Yesterday	Custom day
Week	Current week	Previous week	Custom week
Month	Current month	Previous month	Custom month
Year	Current year	Previous year	N/A

Note: When you select a report time frame and when you specify a scheduled report, it is important to consider CA CEM system aggregation and data retention settings. If the statistics have not yet been aggregated, if the data has been purged, or if incident settings have been changed, your report will not be what you expect.

The CA CEM administrator configures the system to retain the required data for reporting purposes. However, to achieve the best system performance, the CA CEM statistics must be purged periodically. For more information on data retention settings, see Defining your CA CEM domain (see page 262).

The CA CEM system aggregates data to keep hourly, daily, weekly, and monthly statistics. Aggregation also is a tool used for performance and scalability. For more information on CA CEM aggregation, see About CA CEM report aggregation (see page 319).

CA Technologies Product Documentation. Authorized use only.

Exploring reports

On CEM pages where reports are available, you have various options.

PDF

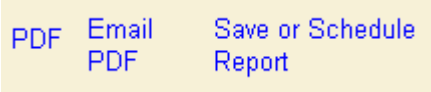
The PDF link presents the report in PDF format. You can save the report to a file or print the PDF.

Email PDF

The Email PDF link opens a new page from which you can send an email with the PDF report attached.

Save or Schedule Report

The Save or Schedule Report link takes you to the My Reports page, where you can save or schedule the report.



You must first query and refresh the report page with valid data before you can click one of these reporting options.

Note: If you click the PDF link on a CEM reports page after a second CEM reports page has been opened (for example, using your browser's Open in New Window option) you get an error. Go back to the first CEM reports page, click Refresh, and click the PDF link again.

PDF reports

The PDF link opens a report in Adobe Acrobat PDF format, which you can save to a file or print.

To print the report:

1. In Internet Explorer, click File > Print to print the PDF-formatted report.
2. Click the Back button to return to the previous page.

To save the report to a file:

1. In Internet Explorer, click File > Save As.
2. Browse to the location where you want to save the report. Type an appropriate File Name for the report. Click Save to save the report to a file.

   Note: If you do not change the name, it is *WilyCEMReport.pdf*.
3. Click the Back button to return to the previous page.

CA Technologies Proprietary Documentation. Authorized Use Only.

To use your own logo on PDF reports:

- Replace the reports.userImage in the ts_settings table in the APM database.

  Note: You cannot use your logo on reports sent via email.

Email reports

The email link allows you to send the report in PDF format. A new window opens with an email form.

To send the report in an email:

1. Complete the email form.

   From

   This is the email address for the logged in user. If an email address has not been provided, then the address might be populated from the email.defaultFromAddress in the tess-default.properties file is used (The Default tess-default.properties file (see page 534)).

   To

   Enter a comma-separated list: email addresses or distribution lists.

   Note: Do not use semicolons to separate email addresses.

   Subject

   Enter a meaningful subject line for the attached report.

   Message

   Add the appropriate information to the standard message for the report.

2. Click Send.

   Note: If you intend to send reports via email, the CEM system administrator needs to configure the mail SMTP server.

Saved or scheduled reports

The Save or Schedule Report and My Reports features allow you to set up reports to run on a regularly scheduled—daily, weekly, or monthly—basis. TheCA CEM system generates and sends the reports via email at the specified time.

For more information, see Scheduling reports (see page 325).

CA Technologies Prerelease Documentation Authorized Use Only

Scheduling reports

There are two ways to schedule reports in CA CEM:

- Define and schedule your report in My Reports (see page 325).
- Generate the report you want and then schedule it (see page 328).

Schedule a report using My Reports

This section describes how to schedule a report manually, via My Reports.

Follow these steps:

1. Select CEM > My Reports and click New.
2. Select the Report Format from the list. Each report format requires additional information, which varies as you make your choices.

   See Report formats (see page 321) for more information.

   A Business Value report requires:

   Business Application

   Select a specific business application or All.

   Business Service

   Select a specific business service or All.

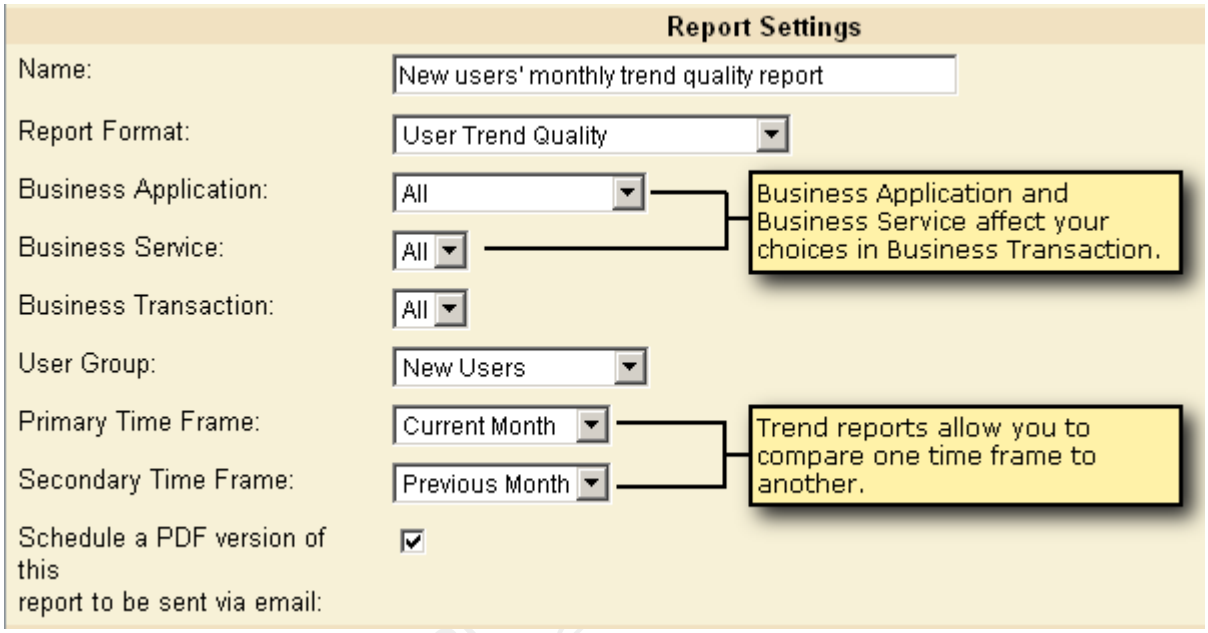
CA Technologies Prerelease Documentation. Authorized Use Only

Time Frame

Specify hour, day, week, month, or year.

Schedule a PDF

Optionally, you can select the option to schedule this report to be delivered via email.



See Report time frames (see page 322) for more information.

A User Trend Quality report requires:

Business Application

Select a specific business application or All.

Business Service

Select a specific business service or All.

Business Transaction

Your choices vary, depending on the business service you selected. Select a specific business transaction or All.

User Group

Select a specific user group or All.

Primary Time Frame

Specify hour, day, week, month, or year. This is usually a current time frame, for example, this week.

CA Technologies Proprietary

Secondary Time Frame

Specify hour, day, week, month, or year. This is usually a previous time frame. For example, you can compare today to yesterday; or this week to last week, or last month, or last year.

Schedule a PDF

Optionally, you can select the option to schedule this report to be delivered via email.

3. Decide how often you want the report to run.

Daily

Select the time-of-day to run the report.

Weekly

Select the time and day-of-week to run the report.

Monthly

Select the time and day-of-month to run the report.

Important! If you select a day of the month that does not occur in a given month, the report is not sent that month. For example, if you selected day-of-month 31, the report is not sent in the months with fewer than 31 days.

4. Decide what time you want the report to run; set the recurrence time.

Tip: If you want a report to include the last hour's data, it is a good idea to schedule the report for late in the following hour—this allows time for the statistics process to finish.
For more information on CA CEM aggregation, see About CA CEM report aggregation (see page 319).

5. Define email settings.

Tip: It is a good practice to use the subject line and message to let your report recipients know the type of report, how often it's run, and a contact name in case they have questions.

From Address

Enter the email address from which to send the report. You can change this, for example if you want report recipients to reply to a different email address.

From Name

Enter the name associated with the email address; this is usually your name, or the name of the responsible CA CEM user.

To

Enter a comma-separated list: email addresses or distribution lists.

CA Technologies Proprietary & Documentation. Unauthorized Use Only

Subject

Enter a meaningful subject line for this report.

Message

Enter a useful message for this report.

6. Click Save. Your new report appears in the list of reports.

All CA CEM users see the reports they created; only the CA CEM system administrator can see reports created by other CA CEM users.

Tip: If you name all your reports in a consistent way, it is easy to find the report you need. For example, all the report names might include:
— the business applications, services and transactions included
— user groups included
— frequency

Generate and schedule a report

This section provides an example of how to schedule a report after you have generated it with the options you want.

To generate and schedule a report:

1. Go to the page you want and generate the report view the way you want it. For example:
   a. Select CEM > Service Level Management > User Trend.
   b. Select the User Group, Business Application, Business Service, and Business Transaction that you want to view in a report.
   c. Select the Primary and Secondary time frames.
   d. Click Quality to view a quality report on the user group.
   e. Click Refresh to update the page to the quality report.
   f. Click Save or Schedule Report.

   The My Reports page appears, with all the Report Settings pre-filled for the User Trend Quality report you just viewed.
2. Modify the pre-filled report settings as required.
3. Select the check box, Schedule a PDF version of this report to be sent via email.

   The Recurrence and Email Settings options appear.
4. Define the recurrence settings as required.

CA Technologies Prerelease Documentation
Authorized Use Only.

5. Define the email settings as required.
6. Click Save to schedule this report.

   See Schedule a report using My Reports (see page 325) for more information on report scheduling options.

Export CA CEM data

You can export CA CEM data from the Export Data tab. This tab allows you to retrieve the business statistics, defects, and incidents data.

You can export CA CEM data for the following:

- Aggregated or nonaggregated business statistics in a specified time range.
- Defects in a specified time range, last *N* number of defects, and defects in last *X* minutes.
- Incidents in a specified time range.

Note: In an e-commerce mode, you cannot export user information in statistics data.

Broadcom Prerelease Documentation. Authorized Use Only.

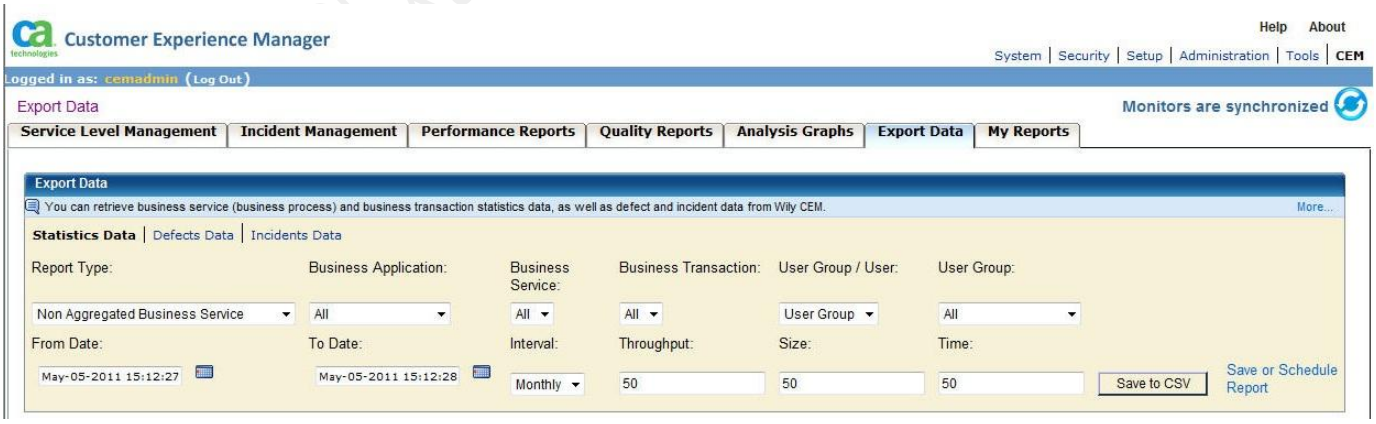
Export data

This section gives an example of how to export CA CEM data.

Follow these steps:

1. Select CEM > Export Data and click Statistics Data.

   The Statistics Data view appears.



2. Specify a combination of the following options:
   a. Select the Report Type, Business Application, Business Service, and Business Transaction that you want to view in data.
   b. Select a User Group or a User.

      Note: If you select a Non Aggregated report type, then only a User Group can be selected.

   c. Specify the From Date and the To Date.

      Note: By default the current system date and time is displayed.

   d. Select an Interval such as Hourly, Daily, Weekly, or Monthly for a Non Aggregated report type.
   e. (Optional) Specify the transaction Throughput, Size, and Time percent value.

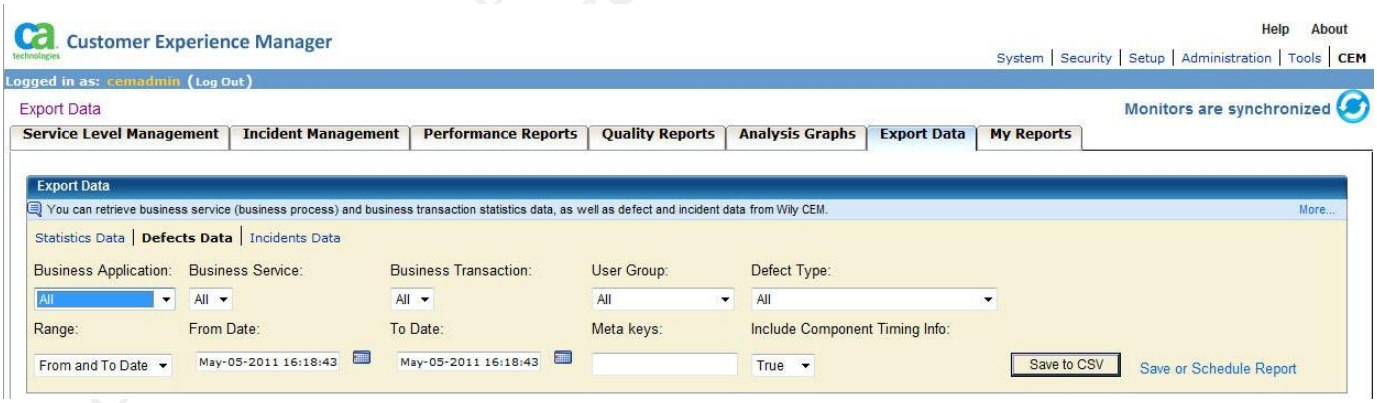
      The Throughput, Size, and Time fields represent the percentile value of a transaction. A semicolon (;) is used to separate multiple percentiles. For example, throughput = 25; 50; 75. You can also specify "all" in these fields to include the percentile value from 5 through 95 incremented by 5.

      Note: By default the transaction Throughput, Size, and Time is 50 percent.

   f. Click Save to CSV.

      The File Download dialog opens a CSV file with the statistics data.
3. Click Defects Data.

   The Defects Data view appears.



4. Specify a combination of following options:
   a. Select the Business Application, Business Service, and Business Transaction that you want to view in the exported data.
   b. Select a User Group and a Defect Type.

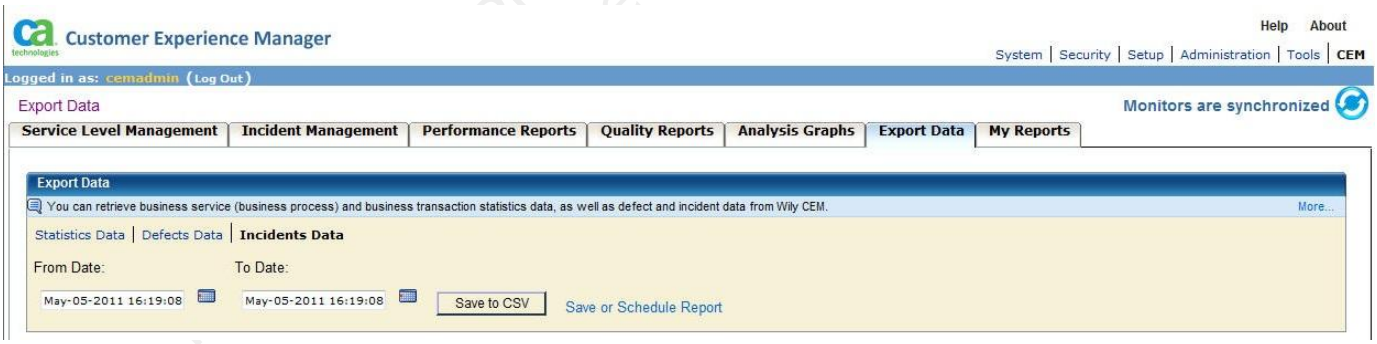
      The Content Error defect type lets you specify content string value that you want to export in the defects data.
   c. Specify the From Date and the To Date, or Last *N* Defects from 0-1000 or Last *X* minutes from 0-1440.
   d. (Optional) Specify the Meta keys that you want to retrieve as part of the defects.

      A Meta key is the value specified in the HTTP Response and Request Headers of the defect.

      Note: A semicolon (;) is used to separate the multiple Meta keys. The Meta keys are case-sensitive.
   e. (Optional) Select Yes to Include Component Timing Information for a defective business transaction.
   f. Click Save to CSV.

      The File Download dialog opens a CSV file with the defects data.
5. Click Incidents Data.

   The Incidents Data view appears.



6. Specify the From Date and the To Date.
7. Click Save to CSV.

   The File Download dialog opens a CSV file with the incident data.

Save or schedule a report

You can save or schedule CA CEM data for reports by selecting the Defects Data, the Incidents Data, or the Statistics Data report format in the My Reports tab.

Note: The scheduled report of CA CEM data is emailed to you as a .zip file attachment. The CSV report of more than 800-KB size is split into equal parts and sent as .zip file email attachments.

See Schedule a report using My Reports (see page 325) for more information about report scheduling options.

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 22: Monitoring Secure Web Applications With CA CEM

This list shows the main tasks involved in setting up CA CEM to monitor web applications that use SSL-based security:

1. Obtain your web server's SSL private keys (see page 333).
2. Make sure that SSL private keys are in a PEM file (see page 333). (Convert if required).
3. Configure CA CEM for SSL and upload the PEM file into CA CEM (see page 338).
4. Verify CA CEM with SSL (see page 339).

Importing and managing SSL private keys

CA CEM supports multiple private keys, which makes it possible to analyze the SSL traffic of various HTTPS servers using different private keys. CA CEM supports the PEM private key format (used by Apache and OpenSSL-based servers). Microsoft IIS, SunONE, iPlanet, and Netscape use a private format, but you can convert those keys to PEM files to be used by CA CEM. CA CEM automatically identifies the key needed for decrypting the data it captures.

A version of OpenSSL must be available on the system used to perform the conversion. OpenSSL is available from http://www.openssl.org.

For details about how to obtain SSL private keys for supported web servers, see:

- Apache or OpenSSL-based web servers (see page 334).
- Microsoft Internet Information Service (IIS) (see page 334).
- SunONE, iPlanet, or Netscape enterprise server (see page 335).

CA Technologies Proprietary Documentation. Authorized Use Only.

Apache or OpenSSL-based web servers

You can save an Apache or OpenSSL-based web server private key file so that it can be used with CA CEM.

To save an Apache or OpenSSL-based web server private key file:

1. Locate the server's configuration file. For example:

   `/etc/httpd/conf.d/ssl.conf`

2. Look for the SSL certificate key file. It points to the private key file. For example:

   `SSLCertificateKeyFile=/etc/httpd/conf/ssl.key/server.key`

3. Copy and rename that file to output.pem.

For more information, see: http://httpd.apache.org/ and http://www.openssl.org/

Microsoft Internet Information Service (IIS)

Microsoft Internet Information Service (IIS) private keys must be exported to a format that CA CEM can read. This section describes how to do this. A version of OpenSSL must be available on the system used to perform the conversion. You can get OpenSSL from http://www.openssl.org/.

First you must export the private key from the IIS server, and then convert the output file to the PEM format.

To export the key in PFX format:

1. On the Windows-based system running IIS, click Start, and then click Run.
2. Type in MMC.exe and click OK.
3. Click the Console menu and then click Add/Remove Snap-in.
4. Click Add. Select the Certificates snap-in and click Add.
5. Select Computer Account and click Next.
6. Select Local Computer and click OK.
7. Click Close and then click OK.
8. Expand the menu for Certificates and click on the Personal folder. Right-click on the certificate that you want to export and select All tasks > Export.

   A wizard appears.

9. Check the box to include the private key and continue through the wizard until you have a PFX file.

CA Technologies Pre-release Documentation. Authorized Use Only.

To convert the PFX file to a PEM key file:

- Execute this command on a system running the OpenSSL software:

  ```
  openssl pkcs12 –in filename.pfx –nocerts –out output.pem
  ```

  where *filename.pfx* is the name of the file you just created in the previous section and *output.pem* is the name of the output file you are creating. The password specified when creating this PEM file is the pass phrase you need when uploading this file to the Enterprise Manager. See Configuring CA CEM for SSL (see page 338) for more information.

SunONE, iPlanet, or Netscape enterprise server

Sun, SunONE, iPlanet, and Netscape web servers are all based on the same core technology. Their private keys must be exported to a format that CA CEM can read. This section describes how to do this.

The following must be available on the system used to perform the conversion:

- OpenSSL
- Network Security Services (NSS)
- Netscape Portable Runtime (NSPR)

These steps assume that a Windows system will be used for the conversion.

To create the private key files:

1. Download and extract the key management utilities:
   - OpenSSL — available from http://www.openssl.org/
   - Network Security Services (NSS) — available from http://www.mozilla.org/projects/security/pki/nss/
   - Netscape Portable Runtime (NSPR) — available from http://www.mozilla.org/projects/nspr/
2. Extract the file.
   - Extract the ZIP files to a temporary directory.
   - Extract nss-3.9.zip to c:\.
   - Extract nspr-4.4.1 to c:\.
3. Run Win32-OpenSSL-v0.9.7c and install in C:\OpenSSL.
4. Modify the PATH variable to include:

   ```
   c:\nss-3.9\lib;c:\nspr-4.4.1\lib
   ```

   The SunONE/iPlanet/Netscape certificates and keys are stored in files with a .db extension.

CA Technologies Prerelease Documentation: Authorized Use Only

5. Copy the *.db files to a writable directory (in this example, c:\dbfiles).
6. Open a command prompt and type:

   ```
   cd c:\nss-3.9\bin
   ```

   By default, the NSS Tools look for files by the name of key3.db and cert7.db. If your database files are named something other than key3.db and cert7.db, do not rename the files. Instead, to have certutil and pk12util find unconventionally named files, provide the –P (prefix) switch.

   For example, if your files are named prodSite-cert.db and prodSite-key.db, use the certutil command to create the files prodSite-cert7.db and prodSite-key3.db. The output of this command displays the contained certificates' names:

   ```
   certutil -P prodSite- -K –d C:\dbfiles\
   ```

7. Enter the password or PIN for the NSS Certificate DB: (enter the database password).

   The files prodSite-cert7.db and prodSite-key3.db are created.

To perform the export and conversion:

1. At the prompt, type:

   ```
   pk12util -P prodSite- -d C:\dbfiles\ -o C:\dbfiles\output.p12 –n Server-Cert
   ```

2. Enter the password or PIN for the NSS Certificate DB: (enter the database password).
3. Enter the password for PKCS12 file: (enter the output password).
4. Re-enter the password: (re-enter the output password).

   You now have a file name output.p12 in the C:\dbfiles directory.

   The output password is the pass phrase you need when uploading this file to the Enterprise Manager. See Configuring CA CEM for SSL (see page 338) for details.

CA Technologies Prerelease Documentation. Authorized Use Only.

5. From the OpenSSL\bin directory, run the command:

   ```
   openssl pkcs12 –in C:\dbfiles\output.p12 -nodes –out C:\dbfiles\pkcs12out.txt
   ```

   This command converts the pkcs12 structure into a text file containing MAC and bag information. It also creates PEM files of the private key and certificate. You now need to copy and paste the key structure into a separate file.

   The contents of the pkcs12out.txt file will look similar to:

   ```
   Enter Import Password:
   MAC verified OK
   Bag Attributes
   localKeyID: 01 00 00 00
   1.3.6.1.4.1.311.17.1: Microsoft Enhanced Cryptographic Provider v1.0
   friendlyName: 98849c683ad0e90810a77235bd728b12_668b93c6-0795-
   4ba7-9da8-78737a299d3f
   Key Attributes
   X509v3 Key Usage: 10
   -----BEGIN RSA PRIVATE KEY-----
   MIICXQIBAAKBgQDLIwhdY7ngYk/AlomX9rqnaZfb8PdJ+Mc6msuRWBUDHeIH2eV1
   yiL7ID/0vkUN6qGnnqVuMQ3kuy7mnN/qSPz37/FKTbwF10QOP2LCZVuhU6SI/90Q
   1rTINqQEPi1hyPV10lDv8Fjevhv80jMD1gvjaCmP84FzcrifyDgHonh4nQIDAQAB
   AoGAUXaJIjqeDjge15MHNuQBUqXr+o4V/ZpA+2WcgXsbYhaX0KI3fRdFFASI4XN0
   ZdpolP7oY/cdF6w5EFsVc64t7dIpD+AineQpGeYJ4WhmBstyVlBShNb9drcCQDFg
   vvcyfmi3tTl3MfecV3WhENU6889iGogbglMMQvqmtOCFDPY7fPM3I38FbEwO3Yu+
   gWEZ/f7DLDdsT7GDKf0CQQDSAqXwTMabqGB6bWxqy6M5lpWBW8YAQAXpKhLILRqb
   C/JxfD63KTN90pegkhEOBn5OdqtqfS4kodv45bHeHYxn
   -----END RSA PRIVATE KEY-----
   Bag Attributes
   localKeyID: 01 00 00 00
   subject=/C=US/ST=Utah/L=SLC/O=MooseCo./OU=Moose
   Land/CN=www.megamoose.com
   issuer=/C=US/ST=Utah/L=SLC/O=MooseCo./OU=Moose
   Land/CN=www.megamoose.com
   -----BEGIN CERTIFICATE-----
   MIICgDCCAemgAwIBAgIDBJ3zMA0GCSqGSIb3DQEBBAUAMIGSMQswCQYDVQQGEwJa
   QTEVMBMGA1UECBMMV2VzdGVybiBDYXBlMRIwEAYDVQQHEwlDYXBlIFRvd24xDzAN
   BgNVBAoTBlRoYXd0ZTEdMBsGA1UECxMUQ2VydGlmaWNhdGUgU2VydmljZXMxKDAm
   BgNVBAMTH1BlcnNvbmFsIEZyZWVtYWlsIFJTQSAyMDAwLjguMzAwHhcNMDEwNDE3
   MDAzNzU5WhcNMDIwNDE3MDAzNzU5WjBFMR8wHQYDVQQDExZUaGF3dGUgRnJlZW1h
   aWwgTWVtYmVyMSIwIAYJKoZIhvcNAQkBFhNqbGV2eUBzb25pY3dhbGwuY29tMIGf
   MA0GCSqGSIb3DQEBAQUAA4Ger3DBiQKBgQDLIwhdY7ngYk/AlomX9rqnaZfb8PdJ
   Y29tMAwGA1UdEwEB/wQCMAAwDQYJKoZIhvcNAQEEBQADgYEAcHF7eiQDoCQ08snA
   izV22+7FWiEkSR0PLcOFTytyOUJoQ0RLt7SPqKPJ2NswyC0A2nHgnOFL7ImFGiAi
   5Xg9lEAtscgC7ceo0C3GBKorPE8hiqGCuVcClYKLws7yGRWRPTzQp7TemCoAMOcv
   iO/6K6qicSzUpEj7eYqCbogmEYU=
   -----END CERTIFICATE-----
   ```

CA Technologies - Unauthorized Documentation.

6. Cut and paste the key into a file with a PEM extension (for example, output.pem).

```
-----BEGIN RSA PRIVATE KEY-----
MIICXQIBAAKBgQDLIwhdY7ngYk/AlomX9rqnaZfb8PdJ+Mc6msuRWBUDHeIH2eV1
yiL7ID/0vkUN6qGnnqVuMQ3kuy7mnN/qSPz37/FKTbwF10QOP2LCZVuhU6SI/90Q
1rTINqQEPi1hyPV10lDv8Fjevhv80jMD1gvjaCmP84FzcrifyDgHonh4nQIDAQAB
AoGAUXaJIjqeDjge15MHNuQBUqXr+o4V/ZpA+2WcgXsbYhaX0KI3fRdFFASI4XN0
ZdpolP7oY/cdF6w5EFsVc64t7dIpD+AineQpGeYJ4WhmBstyVlBShNb9drcCQDFg
vvcyfmi3tTl3MfecV3WhENU6889iGogbglMMQvqmtOCFDPY7fPM3I38FbEwO3Yu+
gWEZ/f7DLDdsT7GDKf0CQQDSAqXwTMabqGB6bWxqy6M5lpWBW8YAQAXpKhLILRqb
C/JxfD63KTN90pegkhEOBn5OdqtqfS4kodv45bHeHYxn
-----END RSA PRIVATE KEY-----
```

7. Upload the output.pem file to the Enterprise Manager. See Configuring CA CEM for SSL (see page 338) for more information.

Configuring CA CEM for SSL

Follow these steps:

1. Select Setup > HTTPS Settings.
2. Select the address type.
3. Specify the web server IP address range.

   For example, 172.16.10.2-172.16.10.15

   Separate the two addresses by a single hyphen without spaces.
4. Specify the web server port number.
5. Browse to the location of the output.pem file (on your client computer) and select it.

   If you have a private key file that needs to be converted to PEM format, see Importing and managing SSL private keys (see page 333).
6. Enter the Key File Pass Phrase, which is the password that was used to created the PEM file.

   If a password was not used to create the file, leave this field blank.
7. Click Save to upload the PEM file.

   If there are different output.pem files for each web server, repeat the steps for each web server.

CA Technologies Prerelease Documentation. Authorized Use Only.

About multiple keys for HTTPS servers

An HTTPS server or a range of HTTPS servers can support multiple private keys if there is a different port or address for each key.

For example, if you have one HTTPS server with one key for internal users and another key for one external users, you could use one port for the internal key and another port for the external key.

If you enter a second key for the same address (or address range)/port pair, the first key is overwritten. If you have overlapping address ranges, the keys are kept separately, but a TIM uses only one of them—the first one it finds that includes the actual IP address and port.

Verifying the functionality of CA CEM with SSL

Follow these steps:

1. Access the TIM System Setup page. (Accessing the TIM from the CEM console (see page 234).)
2. Click View TIM Transaction Inspection.
3. Enter the IP address of the computer to monitor.
4. Click Start to begin monitoring transactions defined as SSL.
5. Execute a simple HTTPS transaction (that you defined before now).

   The transaction inspection data appears on the page after a few seconds.
6. Click Stop when you are finished with your transaction inspection.

Note: If you do not click Stop, the inspection data continues to be collected.

CA Technologies Prerelease Documentation. Authorized Use Only.

Chapter 23: Using a Napatech Adapter with CA CEM

Installing a Napatech adapter can improve the TIM performance in some high traffic environments.

This section contains the following topics:

About using the Napatech adapter with CA CEM (see page 341)
Obtaining the Napatech adapter and software (see page 342)
Installing the Napatech adapter and Napatech software (see page 342)
Enabling the Napatech adapter and monitoring (see page 343)
(Optional) Creating and enabling Napatech filters (see page 344)
Other Napatech adapter tasks (see page 345)
Troubleshooting a Napatech adapter (see page 347)

About using the Napatech adapter with CA CEM

Installing the adapter in a TIM machine can prevent the TIM from dropping packets. This can work in the following ways:

- The adapter buffer stores packets until the TIM can process them.
- The adapter can filter traffic so that there is less traffic for the TIM to analyze.
- The CPU takes less time to read data from a Napatech adapter than from most other Ethernet adapters, leaving more CPU time for the TIM.

The adapter can remove duplicate packets — a good thing if your span port sees the same packets more than once.

The process for using the Napatech adapter with a TIM:

1. Obtain the correct adapter and software. (see page 342)
2. Install the adapter and the software. (see page 342)
3. Configure the TIM to use the adapter. (see page 343)
4. (Optional) Create and enable filters for the adapter. (see page 344)

CA Technologies Pre-release Documentation. Unauthorized Use Only.

Obtaining the Napatech adapter and software

The Napatech NT4E-4T PCIe adapter, with four electrical gigabit interfaces, is the only adapter supported for the TIM. This adapter cannot be used with fiber cable.

Network Allies, LLC is the only reseller who can supply both the correct adapter and Napatech software that you need. This Napatech software is designed for only the NT4E-4T PCIe adapter.

Note: TIM running on the Multi-Port Collector appliance supports different Napatech adapters. See the *CA APM Compatibility Guide* (available from CA Support) for more information about supported adapters.

To obtain the Napatech adapter and Napatech software:

- Purchase the NT4E-4T PCIe adapter from Network Allies, part number 801-0075-0800CA.

  The package includes both a DVD and USB flash drive containing software and installation instructions.

Installing the Napatech adapter and Napatech software

The DVD and USB flash drive contain:

- The Napatech hardware installation instructions
- The software image that you install on the TIM
- The sample filter file

To install the adapter

1. Follow the instructions in the *NT4E-4T Hardware Installation Guide* provided on the DVD and USB flash drive. Install the adapter into the PCI Express slot of the TIM machine.
2. Move the traffic monitoring cable from the monitoring port (typically, eth1) to any port in the adapter.

   Note: You can also do this step later, after Enabling the Napatech adapter and monitoring (see page 343). Doing it later reduces data loss. However, doing it now, means one fewer trip to the server room.

CA Technologies Prerelease Documentation. Authorized Use Only

To install the software

1. Make the Napatech software image (napatech-<version>.image) available on your network.
2. Install the napatech-<version>.image from the TIM Wily System Setup page. For general information about installing TIM software, see the *CA APM Installation and Upgrade Guide*.
3. (Optional) Confirm that the Napatech software is installed by navigating to the Wily System Setup page.

   If the software is installed, a bulleted line “Napatech <installation timestamp>” appears on the page. (It is not a link.)
4. (Optional) Review the Napatech driver log.
   a. On the TIM System Setup page, click Configure Napatech.
   b. Click View the Napatech Driver Log.
   c. Review the log.

      For example, if you see lines like the following in the log on this page:

      ```
      4B85587F.00030BC4 | #WARN | NOTE: The adapter will have limited throughput.
      4B85587F.00030D15 | #WARN | Link negotiated to 4 lanes. Adapter supports 8 lanes.
      ```

      then you installed the Napatech adapter in a four-lane slot instead of an eight-lane slot. This can affect the performance of your Napatech board but does not stop it from working.

Enabling the Napatech adapter and monitoring

After you have installed the Napatech adapter and the Napatech software, follow these procedures to enable the Napatech adapter and configure the TIM to monitor the ntxc0:0 port.

To enable the adapter

1. On the TIM System Setup page, click Configure Napatech.

   The TIM Napatech Configuration page appears.
2. Click Enable Napatech and restart TIM.

   If the TIM is not running, click Enable Napatech. Start the TIM to begin monitoring.

CA Technologies Pre-release Documentation. Authorized Use Only.

To configure the TIM network monitor interface to use the adapter

1. On the TIM System Setup page, click Configure TIM Network Interfaces.

   For general information about configuring the TIM network monitor interface, see the *CA APM Installation and Upgrade Guide*.

2. Select the ntxc0:0 interface and click Set.

   If you moved, or plan to move, the traffic monitoring cable from another interface to ntxc0:0, then clear the check box for the interface that you are not using. TIM can monitor from the ntxc0:0 interface and another interface, such as eth1; or from only the ntxc0:0 interface.

3. If you have not already connected the traffic monitoring cable to one of the adapter ports, do it now: Move the traffic monitoring cable from the monitoring port (typically, eth1) to any port in the adapter.

(Optional) Creating and enabling Napatech filters

The Napatech adapter provides a way to filter packets in the hardware.

Filters are written in the Napatech programming language (NTPL) and then uploaded to the TIM. Use the sample filter file. The sample filter file is on the USB flash drive

The TIM itself can filter traffic by web server IP address. (For information about configuring TIM monitoring this way, see the *CA APM Installation and Upgrade Guide*.) However, best practice is to filter exclusively with the Napatech adapter. Hardware filtering is faster.

If you have set web server filters for the TIM and a Napatech filter, these filters are ANDed. That is, you get the intersection of the filters.

To create and enable a filter for the Napatech adapter

1. Copy the sample filter file (sample_filter.ntpl) from the DVD or USB flash drive and edit it for your environment (using a simple text editor such as Notepad or vi).

   Note: Verify that the DeleteFilter=all command is at the beginning of the file.

2. On the TIM System Setup page, click Configure Napatech.

   The TIM Napatech Configuration page appears.

3. Click Upload a Napatech Hardware Filter Configuration File.

   The TIM Napatech Hardware Filter Configuration page appears.

4. Browse for the file, and then upload and install it.

CA Technologies Proprietary Use Document
Authorized User Only.

5. (Optional) Confirm that the filter is installed.
   a. On the TIM System Setup page, click Configure Napatech.
   b. Click View the Napatech Hardware Filter Configuration File (on the TIM Napatech Configuration page). Or, for more detailed information, click View the Napatech Hardware Filter Configuration Log.

Important! When you upload a new file, the information from the old file is no longer visible on the Napatech hardware filter configuration log or file pages. However, the previous filters themselves are deleted only if you include the DeleteFilter=all command at the beginning of the new filter file.

Other Napatech adapter tasks

Here are a few additional Napatech adapter tasks that you might need to perform:

- Removing the filter from the Napatech adapter
- Removing the Napatech adapter
- Viewing the Napatech adapter-specific logs

Removing the filter from the Napatech adapter

You can remove the filters from the Napatech adapter when troubleshooting or if you no longer want to filter traffic.

To remove the filter from the Napatech adapter

1. On the TIM System Setup page, click Configure Napatech.

   The TIM Napatech Configuration page appears.
2. Click Upload a Napatech Hardware Filter Configuration File.

   The TIM Napatech Hardware Filter Configuration page appears.
3. Click Delete all Napatech hardware filters.

   The deletion applies the capture all filter to the Napatech adapter so that all traffic is sent to the TIM.
4. (Optional) Confirm that the filter has been removed.
   a. On the TIM System Setup page, click Configure Napatech.
   b. Click View the Napatech Hardware Filter Configuration File (on the TIM Napatech Configuration page).

      Verify that the page shows:

      ```
      Capture=All
      ```

CA Technologies Pre-release Documentation
Authorized Use Only.

Removing the Napatech adapter

You might want to remove the Napatech adapter from a TIM; for example, to try the adapter on another TIM machine.

To remove the Napatech adapter

1. Clear the ntxc0:0 interface check box. (To configure the TIM network monitor interface to use the adapter (see page 343).)
2. Disable the Napatech adapter:
   a. On the TIM System Setup page, click Configure Napatech.

      The TIM Napatech Configuration page appears.

   b. Click Disable Napatech and restart TIM. Or, if TIM is not running, click Disable Napatech.
3. Remove the Napatech adapter hardware from the TIM machine.

Viewing the Napatech adapter-specific logs

There are two log files specific to the Napatech adapter:

- The TIM Napatech hardware filter configuration log is updated when a filter is applied or removed.
- The Napatech driver log is updated each time the Napatech driver is loaded. (The driver is loaded whenever the TIM restarts.)

To view the TIM Napatech hardware filter configuration log

1. On the TIM System Setup page, click Configure Napatech.

   The TIM Napatech Configuration page appears.
2. Click View the Napatech Hardware Filter Configuration Log.

   The TIM Napatech hardware filter configuration log appears.

To view the Napatech driver log

1. On the TIM System Setup page, click Configure Napatech.

   The TIM Napatech Configuration page appears.
2. Click View the Napatech Driver Log.

   The Napatech driver log appears.

CA Technologies Pre-release Documentation. Authorized Use Only.

Troubleshooting a Napatech adapter

About Napatech adapter and the TIM log

In the TIM log, all messages related to network interfaces are prefaced with AdapterManager. For example, messages related to the Napatech adapter refer to the ntxc0:0 interface. For example:

```
AdapterManager: "ntxc0:0": thread starting
```

Troubleshooting topics

This page contains the following troubleshooting topics:

- The Napatech adapter is not receiving traffic (see page 347)
- The Napatech adapter as an Ethernet port (see page 348)
- Not seeing the traffic you expect (see page 348)
- Serial number and other Napatech adapter details (see page 348)

The Napatech adapter is not receiving traffic

Symptom:

The TIM is not receiving packets. I review the TIM Packet Statistics page and see that packets are not being captured.

Confirm that there is traffic to see. Move the cable from the ntxc0:0 port to the eth1 port and turn on the eth1 port. If you see packets now, then the problem is the adapter.

Solution:

Verify that the interface is enabled. See To configure the TIM network monitor interface to use the adapter (see page 343).

Verify that the software is configured. See Installing the software (see page 342).

Review the logs. See Viewing the Napatech adapter-specific logs (see page 345).

Look for problems in the logs. See Viewing the Napatech adapter-specific logs (see page 345) and About Napatech adapter and the TIM log (see page 347).

Remove filters. See Removing the filter from the Napatech adapter (see page 345).

If you think that the problem is the Napatech adapter itself, call Network Allies; otherwise, call CA Support.

CA Technologies Prerelease Documentation. Authorized Use Only.

The Napatech adapter as an Ethernet port

Symptom:

When I type the ifconfig command at a Linux prompt on the TIM, the ntxc0:0 port does not appear. The operating system does not recognize the Napatech adapter as an Ethernet port.

Solution:

No action required. The Napatech adapter is behaving normally in this application.

Not seeing the traffic you expect

Symptom:

I'm seeing some traffic but not the traffic I expect.

Symptom:

If you have enabled more than one interface for monitoring, turn off the other interfaces to confirm that you are getting some traffic from the ntxc0:0 interface. If you are not getting any traffic from the ntxc0:0 interface, see The Napatech adapter is not receiving traffic (see page 347).

Review the filters. Remove the filter. Do you see the traffic you expect? If so, review and edit the filter and reapply it. Verify that you included the DeleteFilter=all command at the beginning of the filter file.

Serial number and other Napatech adapter details

Symptom:

I want to find the serial number for my Napatech adapter.

(You might need the serial number if you are contacting Network Allies support.)

Solution:

You can view information about the Napatech adapter on the Napatech status page.

To open the Napatech status page

1. On the TIM System Setup page, click Configure Napatech.

   The TIM Napatech Configuration page appears.

2. View Napatech status.

   The Napatech AdapterInfo program runs, and the output displays.

CA Technologies Proprietary Documentation. Authorized Use Only.

Chapter 24: Configuring CA Introscope to work with CA CEM

CA CEM meets the needs of two distinct users: those responsible for web businesses obtain trend, service level, and success rate reports; and those responsible for responding to incidents get real-time visibility into transaction defect details. Introscope provides the industry-leading solution for problem isolation at the application server level. When customers integrate CA CEM and CA Introscope, they obtain a complete end-to-end view of their web application in operation. This solution is CA APM.

Real-time knowledge of the entire end-to-end customer experience provides new insights that can significantly improve service delivery, increase customer loyalty, and enhance the overall online experience for your customers. Transaction triage information, available in CA APM reports, provides a common language that meets the needs of both IT and business stakeholders.

CA CEM enhances the IT role as a *partner* with business management—with the team *together* ensuring the delivery of high-quality online services to customers.

This chapter describes how to install and configure Introscope for integration with CA CEM. For more information on how to use the products together, see Using Introscope with CA CEM. For information on the Introscope agent recording feature, see the *CA APM Transaction Definition Guide*.

This section contains the following topics:

Getting started with CA APM (see page 349)
CA APM Customer Experience Metrics Overview (see page 350)
CA APM problem resolution triage overview (see page 350)
Checklist for configuring CA CEM (see page 352)
Checklist for configuring Introscope (see page 362)
Troubleshooting the CA Introscope to CA CEM Integration (see page 371)
Upgrade considerations (see page 382)
Setting up the business application (see page 385)

CA Technologies Proprietary. Please Do Not Reproduce. Authorized Use Only.

Getting started with CA APM

CA APM requires the integration of CA Introscope and CA CEM.

This list shows how to get started with CA APM:

1. Learn how CA CEM with CA Introscope provides the business with the customer experience information they need. See CA APM customer experience metrics overview (see page 350).

2. Learn how CA CEM with CA Introscope assists with IT triage and problem resolution. See CA APM problem resolution triage overview (see page 350).
3. Configure the CA CEM side of the integration. See Checklist for configuring CA CEM (see page 352)
4. Configure the CA Introscope side of the integration. See Checklist for configuring CA Introscope (see page 362).
5. Verify that the integration features work. See Troubleshooting the CA Introscope to CA CEM integration (see page 371).
6. If you are upgrading your CA CEM / CA Introscope environment to CA APM, review the upgrade considerations. See Upgrade considerations (see page 382).

CA APM Customer Experience Metrics Overview

CA CEM and CA Introscope customer experience metric integration is controlled by the configuration parameters on the Enterprise Manager and on each TIM.

1. The TIM writes all customer experience metrics to an XML file, every 7 seconds or as defined, until they are deleted or they have reached or exceeded the age-out period.
2. The Enterprise Manager TIM Collection Service retrieves the customer experience metrics files and instructs the TIM to delete them.
3. The Enterprise Manager TIM Collection Service parses the customer experience metrics data, but does not store the data in the APM database, nor does it perform any aggregation on the data.
4. The customer experience metrics are grouped by the originating TIM; the Enterprise Manager stores customer experience metrics in the SmartStor database.
5. The CA Introscope Workstation displays default business dashboards, using typeviews and calculators to provide aggregation and defect percentages.
6. You can create your own custom business dashboards as needed, using the same customer experience metrics data. See Using CA Introscope with CA CEM.

For more information on the TIM Collection Service, see Distributing Enterprise Manager services *(see page 252)*.

CA APM problem resolution triage overview

CA CEM and CA Introscope problem resolution triage integration is controlled by the CA Introscope configuration settings. Configuration information is contained in domain configuration information, which includes transaction definitions, defect specifications, and transaction trace parameters.

CA Technologies Prerelease Documentation Authorized Use Only

1. When you click Setup > Monitors > Synchronize All Monitors, CA CEM notifies CA Introscope of changes to domain configuration information. (CA CEM also sends the domain configuration information to the TIM when you click Synchronize All Monitors.)
2. CA Introscope agents listen for any updates to the domain configuration information. (For pre-9.0 agents, CEMDefinitionHandler sends parsed domain configuration information to all registered agents.)
3. Each registered agent then creates a rule set that enables transaction traces later. The rule set consists of regular expression patterns for parameter names/values that identify CA CEM transactions.
4. Java agents with ServletHeaderDecorator installed add extra information to each HTTP response header. The header information includes the GUID, which is used later to correlate CA Introscope transaction traces with CA CEM defects. The .NET agents need HTTPHeaderDecorator installed to provide the same functionality.
5. The CA CEM TIM parses HTTP response headers and includes the information in any reported defect, which it sends to the Enterprise Manager.
6. CA CEM saves defect and GUID information in the APM database.
7. CA CEM generates a slow time incident and sends a transaction trace session start request to the Enterprise Manager. (CA CEM generates a slow time incident—changing the incident status from pending to open—based on Setup > Incident Settings.)

   The transaction trace session start request includes the business transaction name and the transaction trace time threshold, which is a percentage of the slow time defect specification. The transaction trace continues to run until the incident closes or the configured transaction trace session duration is reached.
8. Agents with BizTrxHttpTracer installed look up the transaction definition by the business transaction name as specified in the agent's rule set. The agents use parameter matching to identify CA CEM transactions. (Pre-9.0 agents have CEMTracer installed to provide the equivalent functionality.)
9. When an agent identifies a transaction and the transaction time is greater than the transaction trace time threshold, the agent sends a transaction trace to the Enterprise Manager.
10. CA CEM receives the transaction traces and matches them against the GUIDs stored in the APM database. When a GUID match is found, CA CEM saves the transaction trace information in the APM database.

Enterprise Manager transaction trace information and other CA CEM defect information are used to help dispatch the incident to the proper team for resolution. See Using CA Introscope with CA CEM.

CA Technologies Internal Use Only

Checklist for configuring CA CEM

This list shows the things you need to know about configuring CA CEM for integration with CA Introscope.

To configure CA CEM with CA Introscope:

1. Decide which integration features to use. See Deciding which integration features you need (see page 363).
2. Configure the CA Introscope workstations. See Configuring the CA Introscope workstations (see page 353).
3. Configure CA Introscope Transaction Traces. See Configuring transaction traces (see page 354).
4. Configure real-time transaction metrics. See Configuring real-time transaction metrics integration (see page 356).
5. Verify that the integration features work. See Verifying the CA CEM configuration (see page 362).

Deciding which integration features you need

There are two types of integration for CA APM:

- Problem resolution triage — allows you to analyze the root cause of problems by using correlated transaction information.
- Customer experience metrics — provides regular updates on volume, errors, and average response times for business transactions.

If you want problem resolution triage, you need to configure the Introscope Settings. This allows you to configure the workstations and transaction traces.

If you want customer experience metrics, you need to configure the TIM Collection Service. This allows you to configure the Collector and to enable metrics collection on each TIM.

You can choose to use both integration features, or only one, depending on what you need in your business environment.

Note: If you want to use the CA Introscope agent recording feature, see the *CA APM Transaction Definition Guide*.

CA Technologies Prerelease Documentation. Authorized Use Only

More information:

Configuring Introscope Settings (see page 353)
Configuring Customer Experience Metrics Integration (see page 356)
Verifying the CA CEM Configuration (see page 362)

Configuring Introscope Settings

To enable CA APM integration, you need to gather basic information to set up CA Introscope.

This section contains the following topics:

Configuring the CA Introscope workstations (see page 353)
Configuring transaction traces (see page 354)

Configuring the CA Introscope workstations

You need to set up the connection parameters for Workstation Web Start and WebView applications.

To configure CA CEM to communicate with the CA Introscope workstations:

1. Select Setup > Introscope Settings.

   The Introscope Settings page appears.

2. Enable Workstation Web Start link — select this check box if you want to enable CA CEM to display the Workstation Web Start link on incident and defect pages. If you do not check this box, the link will not appear on incident and defect pages.

3. Enter the Introscope WebView information:

   a. Enable WebView link — select this check box if you want to enable CA CEM to display the WebView link on incident and defect pages. If you do not check this box, the link will not appear on incident and defect pages.

   b. Hostname — type the host name or IP address of the CA Introscope system that hosts WebView for CA CEM.

   c. Port — type the port where the Enterprise Manager listens for connections to WebView. The default port number for WebView is 8080.

   d. Context path — type the *introscope.webview.context.path* property value, which specifies the application context path for CA Introscope.

   Note: The default value of *introscope.webview.context.path* is "/" (without quotes in the text box). The property is defined in the *IntroscopeWebView.properties* file on the computer that hosts the Enterprise Manager, in the *<EM_Home>**config*\ directory.

CA Technologies Prerelease Documentation. Authorized Use Only.

4. Click Save Settings.

   CA CEM saves the Introscope workstation configuration information.

5. Continue to Configuring transaction traces (see page 354).

Configuring transaction traces

The CA APM integration starts a transaction trace when an incident is generated and based on these rules. You can also manually start a transaction trace from the CEM > Incident Overview page.

You can configure the transaction trace settings or leave them set to the default values.

Important! Transaction traces 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.) Note that these settings are global, not per business transaction.

To configure CA Introscope transaction traces:

1. Enter the Maximum Transaction Trace Session Duration in minutes. The default is for the transaction trace to stop after 30 minutes.

   Note: The default setting changed in the CA CEM 5.0 release; the previous default was 120 minutes.

   CA CEM starts a transaction trace session with CA Introscope only when there is an incident open for the business transaction. CA CEM has access to a special transaction trace method where it traces by business transaction *and* time threshold. For example, trace all Update Profile web requests that take longer than 1000 ms (that is, 1 second).

   This allows the agent to optimize performance because once it identifies that a request is not Update Profile, no further action is required. Although there might be several transaction trace sessions running, overhead is kept to a minimum.

2. Enter the Transaction Trace Time Threshold as a percentage of the slow time defect specification. The default is 50 percent.

   Note: The default setting changed in the CA CEM 5.0 release; the previous default was 25 percent.

   The transaction trace time threshold is the execution time limit for the transaction. When a transaction trace is running, all transactions that do not complete within the threshold are traced.

   The transaction trace time threshold is a percentage of the slow time defect specification. For example, if the slow time defect specification is set to 5.00 seconds and the transaction trace time threshold is set to 50 percent, then all transactions that take longer than 2.50 seconds to complete (in CA Introscope, that is, on the application server) will be traced.

CA Technologies Reference Documentation
Authorized Use Only.

Using the example above with the default values of 30 minutes and 50 percent, transactions will be traced when these are all true statements:

- there is an open incident for the business transaction.
- the transaction takes longer than 2.50 seconds to complete.
- the transaction trace has been running for less than 30 minutes.

Important! In some cases, you might have a need to set the CA CEM slow time defect specification to a very short time such as 1 ms (that is 1/1000 of a second). For example, you might set it this way for a transaction that happens only in the event of an error. However, if this same transaction is involved in an CA Introscope transaction trace, then the 1 ms setting might cause performance problems on the Enterprise Manager MOM, Enterprise Manager Collectors, and agents.

3. Enter the Backend Components. In CA Introscope, any component that matches the given pattern is considered a backend component.

   The asterisk ("*") is a wildcard and is always used at the end of the pattern.

   For example: *Backends|*,JDBC|**

   Each pattern is separated by a comma. So the above example has two patterns: *Backends|** and *JDBC|**

   If any component in a transaction trace starts with *Backends|** or *JDBC|**, then it is recognized as a backend component.

4. Verify the CEMTracer 4.0 / 4.1 Support check box setting.

   Important! This option should be used only as long as you have not yet upgraded all the agents to CEMTracer 4.2 or higher. When you upgrade to CA CEM 4.2 / 4.5 with the CA Introscope plug-in enabled, the default setting is on, which means you expect to have both CEMTracer 4.0 / 4.1 deployed on some agents, and CEMTracer 4.2 / 4.5 deployed on other agents. For more information, see Upgrade considerations (see page 382).

5. Click Save Settings.

   CA CEM saves the Transaction Trace configuration information.

   Note: If you want to learn about monitoring transaction trace sessions, see Viewing transaction trace session information.

CA Technologies Internal Documentation
Authorized Use Only.

Configuring Customer Experience Metrics Integration

This section provides information about upgrading with customer experience metrics, how to start/stop the service, and optional troubleshooting.

This section contains the following topics:

Upgrading with Customer Experience Metrics (see page 356)
Starting the Customer Experience Metrics Collection Service (see page 356)
Stopping the RTTM Collection Service (see page 357)
Troubleshooting Customer Experience Metrics on the TIMs (see page 357)
TIM Settings for Customer Experience Metrics (see page 359)

Upgrading with Customer Experience Metrics

Note: If you used customer experience metrics in a previous release of CA CEM, then the terminology change from "business process" to "business service" can affect you.

If you do nothing after upgrade, then the new term "Business Service" appears in your dashboard metrics. If you prefer to keep the metric name of "Business Process" in the metrics, then you need to set your *tess-customer.properties* and continue to use the CA CEM 4.5 Javascript calculators.

To keep the Business Process metric name for customer experience metrics:

1. Set the *tess-customer.properties*, changing the property: *BTStatsMetricName.backwardCompatible*

   For information about how to edit the tess-customer properties file, see Modifying the tess-customer.properties File (see page 523).

2. Replace the Javascript calculators on the Enterprise Manager (from CA CEM 5.x) with the ones from CA CEM 4.5. You need to replace *CEM-AggregatedTree.js* and *CEM-DefectPercent.js*.

   a. Get the *CEM_BT_Stats_package.zip* file.

   b. Extract *CEM-AggregatedTree.js* into the *<EM_Home>\scripts* directory.

   c. Extract into the *<EM_Home>\scripts* directory.

   d. Restart the Enterprise Manager to reload the Javascript calculators.

Starting the Customer Experience Metrics Collection Service

The TIM Collection Service enables customer experience metrics collection. Usually, as a part of your original installation and set up, you will have distributed the Enterprise Manager services, including the TIM Collection Service.

CA Technologies Prerelease Documentation
Authorized Use Only

See Distributing Enterprise Manager services (see page 252) for more information about Enterprise Manager services.

You must configure the TIM Collection Service and start the RTTM Collection Service before customer experience metrics data will appear.

To start the customer experience metrics collection service:

1. Select Setup > Services > Service Configuration.
2. Select the TIM Collection Service.
3. In the TIM Collection Services list, select RTTM Collection Service, if it is not already selected.
4. Click Start to start the customer experience metrics collection service.
5. Verify that the Service Status changes from Stopped to Running.

Stopping the RTTM Collection Service

You might occasionally want to stop the customer experience metrics feature.

To stop the RTTM Collection Service:

1. Select Setup > Services > Service Configuration.
2. Select the TIM Collection Service.
3. In the TIM Collection Services list, select RTTM Collection Service, if it is not already selected.
4. Click Stop to stop the customer experience metrics collection service.
5. Verify that the Service Status changes from Running to Stopped.

Troubleshooting Customer Experience Metrics on the TIMs

Note: The following are *optional* configuration procedures for the TIMs. In most cases, you should configure the TIM Collection Service via Setup > Services > Service Configuration. See Distributing Enterprise Manager services (see page 252) for more information about the TIM Collection Services.

You can temporarily or permanently disable customer experience metrics collection on a TIM. You can change the collection frequency and file aging and modify other settings to troubleshoot in your environment.

To temporarily disable customer experience metrics on the TIMs:

1. From the Wily System Setup page, select TIM > Configure TIM Settings. (For information on how to get to the Wily System Setup page, see Accessing the Wily System Setup page for the TIM (see page 233).)

   The TIM Settings page appears.

CA Technologies Prerelease Documentation. Authorized Use Only.

2. Click Define a new value.
3. Type the Name: *BtStats/Enabled*
4. Type the New Value: *0*
5. Click Change.

   This temporarily disables customer experience metrics collection on this TIM. The next time you restart the Enterprise Manager, collection resumes.

   Note: The above procedure assumes that customer experience metrics collection was already enabled.

To permanently disable customer experience metrics on the TIMs:

1. From the Wily System Setup page, select TIM > Configure TIM Settings.

   The TIM Settings page appears.
2. Click Define a new value.
3. Type the Name: *BtStats/ForceDisabled*
4. Type the New Value: *1*

   This disallows CA CEM from controlling customer experience metrics on this TIM.
5. Click Change.

   This disables customer experience metrics collection on this TIM. The next time you restart the Enterprise Manager, customer experience metrics collection will *not* resume on this TIM.

   Note: You can re-enable CA CEM control of customer experience metrics collection on this TIM by changing the value back to the default value of *0*.

   More TIM settings can optionally be changed in your environment and are defined in TIM Settings for customer experience metrics. If these TIM settings do not appear on the TIM Settings page, then they are set to the default value.

To modify customer experience metrics processing on the TIMs:

1. From the Wily System Setup page, select TIM > Configure TIM Settings.

   The TIM Settings page appears.
2. If the TIM setting already exists:
   a. Click the name.
   b. Type the new value.
   c. Click Change.
3. If the TIM setting does not exist:
   a. Click Define a new value.
   b. Type a New Name.

CA Technologies Prerelease Documentation. Authorized Use Only.

c. Type a New Value.

d. Click Add.

TIM Settings for Customer Experience Metrics

These TIM settings can help you troubleshoot customer experience metrics.

This section contains the following topics:

BtStats/ForceDisabled (see page 359)
BtStats/Enabled (see page 359)
BtStats/IntervalInSeconds (see page 360)
BtStats/LogLevel (see page 360)
BtStats/WriteEmptyFiles (see page 361)
BtStats/MaxFileAgeInSeconds (see page 361)

BtStats/ForceDisabled

This setting allows you to control the customer experience metrics feature on the TIM itself. It also allows you to override CA CEM control.

Settings

Numeric value of:

- 0 (allows CA CEM to control the feature on this TIM)
- 1 (the feature is forced off on this TIM)

Default

0

Example

```
BtStats/ForceDisabled = 0
```

Notes

- The customer experience metrics feature can be disabled on the TIM.
- 1 = The *BtStats/Enabled* TIM setting (next) is ignored.
- 0 = The customer experience metrics feature is controlled by CA CEM and uses the *BtStats/Enabled* TIM setting (next).

BtStats/Enabled

This setting allows you to control the customer experience metrics feature on the TIM itself.

Settings

Numeric value of:

CA Technologies Prerelease Documentation. Authorized Use Only

- 0 (disabled by default)
- 1 (enabled)

Default

0

Example

```
BtStats/Enabled = 0
```

Notes

- CA CEM controls this setting on the TIM when you start up the Enterprise Manager. However, the *BtStats/ForceDisabled* TIM setting (previous) can override this setting.
- You can also modify this setting on the TIM to temporarily disable customer experience metrics collection on the TIM. However, the next Enterprise Manager restart can change it.
- See Troubleshooting customer experience metrics on the TIMs (see page 357).

BtStats/IntervalInSeconds

This setting allows you to control how often the TIM writes the real-time transaction statistics files.

Settings

Numeric value; must be a positive integer.

Default

7

Example

```
BtStats/IntervalInSeconds = 7
```

Notes

The real-time transaction statistics files are stored on the TIM in the directory */etc/wily/cem/tim/data/out/btstats*.

BtStats/LogLevel

This setting allows you to control TIM log file entries for customer experience metrics.

Settings

Numeric value of:

- 0
- 1
- 2

CA Technologies Prerelease Documentation. Authorized Use Only.

Default

0

Example

```
BtStats/LogLevel = 0
```

Notes

- 0 = No messages (default)
  - 1 = Messages are written to the TIM log when BtStats files are written
  - 2 = Same as messages, verbose mode

BtStats/WriteEmptyFiles

This setting allows you to control whether or not the TIM writes empty real-time transaction statistics files.

Settings

Numeric value of:

- 0
- 1

Default

0

Example

```
BtStats/WriteEmptyFiles = 0
```

Notes

- 0 = Empty files not written (default)
- 1 = Empty files (with empty BTlist) are written
- This can be useful for troubleshooting. For example, if no transactions were detected during a TIM interval, the empty file, with an empty BTlist (that is, no transactions detected), would still be written. This can be used to test if the TIM is creating BtStats files with or without the Enterprise Manager running.
- The real-time transaction statistics files are stored on the TIM in the directory */etc/wily/cem/tim/data/out/btstats*.

BtStats/MaxFileAgeInSeconds

This setting allows you to define how long real-time transaction statistics files are to be kept (if they are not picked up by CA CEM). The TIM deletes files older than this value.

CA Technologies Prerelease Documentation. Authorized Use Only.

Settings

Numeric value; Must be a positive integer.

Default

3600

Example

```
BtStats/MaxFileAgeInSeconds = 3600
```

Notes

- 3600 = 1 hour (default)
- Defines how long BtStats files are to be kept. The TIM deletes files older than this value.
- Setting this too high might cause the TIM to run out of disk space.
- Setting this too low might cause files to be deleted before they have been retrieved by CA CEM.
- The real-time transaction statistics files are stored on the TIM in the directory */etc/wily/cem/tim/data/out/btstats*.

Verifying the CA CEM Configuration

If you have not yet configured CA Introscope for CA APM, go to Checklist for configuring CA Introscope (see page 362).

See Troubleshooting the CA Introscope to CA CEM integration (see page 371) for information on verifying and troubleshooting the configuration.

Checklist for configuring Introscope

This list shows what you need to do to configure CA Introscope for integration with CA CEM.

To integrate CA Introscope with CA CEM:

1. Decide which integration features to use. See Deciding which integration features you need (see page 363).
2. Configure the Enterprise Managers for problem resolution triage. See Configuring the Enterprise Managers (see page 363).
3. Integrate each Java agent for problem resolution triage. See Integrating with Java agents (see page 365).
4. Integrate each .NET agent for problem resolution triage. See Integrating with .NET agents (see page 367).

CA Technologies Prerelease Documentation. Authorized Use Only.

5. Verify that the integration features work. See Verifying the CA Introscope configuration (see page 371).

Deciding which integration features you need

There are two types of integration for CA APM:

- Problem resolution triage — allows you to analyze the root cause of problems by using correlated transaction information.
- Customer experience metrics — provides regular updates on volume, errors, and average response times for business transactions.

If you want problem resolution triage, you need to configure the Enterprise Manager and all affected agents in order to view the problem resolution triage information on slow-time incidents.

Note: The CA CEM problem resolution triage integration is supported only for Java and .NET agent applications that have been instrumented with AutoProbe.

If you want customer experience metrics, you need to configure Enterprise Manager dashboards to view the collected metrics; see Creating your own CA CEM customer experience dashboards.

You can choose to use both integration features, or only one, depending on what you need in your business environment.

Note: If you want to use the CA Introscope agent recording feature, see the *CA APM Transaction Definition Guide*.

More information:

Configuring Introscope Settings (see page 353)
Configuring Customer Experience Metrics Integration (see page 356)
Verifying the CA CEM Configuration (see page 362)

CA Technologies Pre-release Documentation
Authorized Use Only

Configuring the Enterprise Managers

The Enterprise Manager must be configured to enable the CA CEM integration with CA Introscope.

Note: If you have pre-9.0 agents, then you must also redeploy the 4.5 Enterprise Manager extension:

- Download and install the Enterprise Manager extension to the MOM and each Enterprise Manager Collector.

- Restart each Enterprise Manager (JVM) to reload the Enterprise Manager extension.

Establishing CA CEM as a Virtual Agent

You need to configure CA CEM as a virtual agent for the Enterprise Manager. The following procedure enables metrics to be aggregated across agents. It also enables the Go to Introscope Workstation Web Start link to link directly to the Investigator tree, rather than linking only to the CA Introscope Console.

To configure CA CEM as a virtual agent:

1. Locate the *<EM_Home>\config\agentclusters.xml* file.

   Important! CA CEM can be configured to connect either to a single Enterprise Manager or to a MOM Enterprise Manager. If you have multiple Enterprise Managers in your CA Introscope environment, you must modify only the MOM's *agentclusters.xml* file.

2. Uncommment the relevant agent lines in the *agentclusters.xml* file:

```
<agent-cluster name="CEM Agent" domain="SuperDomain">
    <agent-specifier>.*\|.*\|.*</agent-specifier>
    <metric-specifier>Customer Experience\|.*</metric-specifier>
    <metric-specifier>Business Segment\|.*</metric-specifier>
</agent-cluster>
```

   Important! Do not change the name or domain attributes, or the integration will not function correctly.

3. Save your changes to the *agentclusters.xml* file.

   There is no need to restart the Enterprise Manager for this change.

Deploying the Enterprise Manager extension for pre-9.0 agents

Important! You only need to deploy this CA CEM 4.5 extension if you have pre-5.0 agents attached to the Enterprise Manager.

Each Enterprise Manager that connects to pre-9.0 agents needs the CA CEM 4.5 CEMDefinitionHandler Enterprise Manager extension installed to enable the CA CEM integration with CA Introscope.

Follow these steps:

1. Contact CA Support to get the files that you need.
2. Use the CA CEM software link to get the CA CEM 4.5 CEMIntroscopeIntegrationKit.
3. Download the *CEMDefinitionHandler.jar* to the *<EM_Home>\ext* directory on the Enterprise Manager.
4. Continue to Reloading the Enterprise Manager (see page 365).

CA Technologies Proprietary Documentation. Authorized Use Only

Reloading the Enterprise Manager

You must restart the Enterprise Manager in order to complete the Enterprise Manager integration for problem resolution triage.

Restarting the JVM enables the Enterprise Manager extension.

To reload the Enterprise Manager:

1. Restart the Enterprise Manager to reload the Enterprise Manager extension.
2. Return to Deploying the Enterprise Manager extension for pre-9.0 agents (see page 364) and repeat the steps for each Enterprise Manager that connects to pre-9.0 agents.

Integrating with Java agents

You need to perform these procedures for each Java agent that is included in the CA CEM to Introscope integration:

1. Configure AutoProbe tracing to Introscope-enable the applications that are involved in CA CEM transactions.
2. Verify agent properties for AutoProbe and ServletHeaderDecorator.
3. Restart the agent (JVM) to reload the agent configuration properties.

If you need to integrate with a .NET agent, see Integrating with .NET agents (see page 367).

Configuring AutoProbe tracing in Java environments

Details in the Introscope Investigator tree are based on BizTrxHttpTracer and ServletHeaderDecorator being activated in the PBL files in the agent folder. With these PBDs enabled, CA Introscope can display metrics for each business transaction.

Once the integration is in place, the business transactions appear in the tree, using the same transaction name as appears in CA CEM.

For more information about AutoProbe, see the *CA APM Java Agent Implementation Guide*.

To configure AutoProbe tracing on a Java agent:

1. Open the *IntroscopeAgent.profile* in the agent's *<Agent_Home>* directory.
2. Locate the *introscope.autoprobe.directivesFile* property, for example:

   introscope.autoprobe.directivesFile=
   default-typical.pbl

CA Technologies Prerelease Documentation. Authorized Use Only

3. Locate the PBL file specified in the property.

   In the example above, it is *default-typical.pbl*.

4. Open the specified PBL file in a text editor.

5. In the PBL file, uncomment the ServletHeaderDecorator.pbd line:

```
#######################
# Servlet Header Decorator
#
# Enable this PBD for CEM integration. It appends additional
# information to the HTTP response that is then captured by CEM.
# ===============
ServletHeaderDecorator.pbd
```

6. Save your changes to the PBL file.

7. Continue to Verifying default agent properties in Java environments (see page 366).

Note: In CA Introscope 9.0, *biz-trx-http.pbd* is enabled by default. If you have pre-9.0 agents, then *CEMTracer.pbd* must also be uncommented.

Verifying default agent properties in Java environments

The agent profile properties file allows you to modify default agent properties for AutoProbe and ServletHeaderDecorator.

All the CA CEM-related properties are pre-configured for the CA CEM to CA Introscope integration. You need to verify that the default values are appropriate for your installation.

Important! You should review at least the security-related properties. For example, the *introscope.agent.decorator.security* property is set to clear by default. This is appropriate for initial testing, but might reveal information in the transaction header that you do not want known beyond the firewall. You can set the property to encrypted for a more secure production environment.

For more information about CA CEM-related agent properties, see CA Introscope agent profile properties.

For more information, see the *CA APM Java Agent Implementation Guide*.

To verify the agent profile properties:

1. Open the *IntroscopeAgent.profile* in the agent's *<Agent_Home>* directory.

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

2. Compare the properties you find in the agent profile to the properties' descriptions in CA Introscope agent profile properties. For example, you might want to change these properties from the default settings:
   - *introscope.agent.decorator.enabled* — change the default Java agent setting from false to true; this enables the problem resolution triage feature.
   - *introscope.agent.decorator.security* — change the default setting from clear to encrypted; this secures confidential information in transaction headers.
3. If you need to modify default values, you might need to add the property to the *IntroscopeAgent.profile* file and then configure it.
4. Save your changes to the *IntroscopeAgent.profile* file.
5. Continue to Reloading agent properties in Java environments (see page 367).

Tip: If a property is not listed at all, it can be confusing, so it's a good idea to explicitly include it, even if the default value is what you need.

Reloading agent properties in Java environments

You must restart the agent for the managed application in order to complete the integration.

Restarting the JVM reloads the agent properties and enables AutoProbe and ServletHeaderDecorator.

To reload the agent profile properties on a Java agent:

1. Restart the agent to reload the agent properties.
2. Return to Integrating with Java agents (see page 365) and repeat the procedures for each agent.

Integrating with .NET agents

You need to perform these procedures for each .NET agent that is included in the CA CEM to CA Introscope integration:

1. Configure the .NET agent version number to match the wily.Agent assembly version number.
2. Configure AutoProbe tracing to CA Introscope-enable the applications that are involved in CA CEM transactions.
3. Verify agent properties for AutoProbe and HTTPHeaderDecorator.
4. Restart the agent (CLR) to reload the agent configuration properties.

If you need to integrate with a Java agent, see Integrating with Java agents (see page 365).

CA Technologies Prerelease Documentation. Authorized Use Only.

Configuring the agent .NET version

The .NET wily.Agent assembly version needs to be added to the .NET *machine.config* file, under the *<runtime>* section.

Important! If you do not configure the wily.Agent version, your .NET agent machine might fail to start.

Tip: The *machine.config* file location is usually:
`C:\WINDOWS\Microsoft.NET\Framework\vn.n.nnnnn\CONFIG\machine.config`

To configure the agent assembly version on a .NET agent:

1. Verify the wily.Agent assembly version.
   a. Select Administrative Tools > Microsoft .NET Framework n.n Configuration.

      The Microsoft Management Console appears, displaying Console Root > My Computer.

   b. Select Manage the Assembly Cache.

   c. Select View List of Assemblies in the Assembly Cache.

   d. Scroll to the bottom of the Assemblies list to see the wily.Agent.

   e. Note the wily.Agent Version in the next column. For example, 8.2.0.1, or 9.0.0.0.

   WARNING: The next steps must be performed exactly. If you do not update the *machine.config* file properly, the .NET agent will not start.

2. In the .NET *machine.config* file, look for an existing *<runtime>* section.
   a. If the *<runtime>* section exists, add this to the section:

```
<assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1">
    <dependentAssembly>
    <assemblyIdentity name="wily.Agent" publicKeyToken="2B41FDFB6CD662A5"/>
    <bindingRedirect oldVersion="0.0.0.0 - 65535.65535.65535.65535"
    newVersion="AGENT.VERSION.NUMBER" />
    </dependentAssembly>
</assemblyBinding>
```

   b. If the *<runtime>* section does *not* exist, add the section to *machine.config*:

```
<runtime>
    <assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1">
    <dependentAssembly>
    <assemblyIdentity name="wily.Agent" publicKeyToken="2B41FDFB6CD662A5"/>
    <bindingRedirect oldVersion="0.0.0.0 - 65535.65535.65535.65535"
    newVersion="AGENT.VERSION.NUMBER" />
    </dependentAssembly>
    </assemblyBinding>
</runtime>
```

CA Technologies Confidential/Restricted Documentation. Authorized Use Only.

Important! If the agent is running on .NET 2.0 Framework, look for the empty section (look for a *<runtime />* section tag); replace it with the above section.

3. Update the AGENT.VERSION.NUMBER with the wily.Agent version in step 1 of Verify the wily.Agent assembly version.

   Note: This change globally configures all .NET applications to accept the new version number. For more information, see the related Symptoms and solutions (see page 374) or the *CA APM .NET Agent Implementation Guide*.

4. Save the *machine.config* file with the new wily.Agent version number.
5. Continue to Configuring AutoProbe tracing in .NET environments (see page 369).

Configuring AutoProbe tracing in .NET environments

Details in the Introscope Investigator tree are based on the BizTrxHttpTracer and HTTPHeaderDecorator being activated in the PBL files in the agent folder. With the BizTrxHttpTracer enabled, CA Introscope can display metrics for each business transaction.

Once the integration is in place, the business transactions appear in the tree, using the same transaction name as appears in CA CEM.

For more information about AutoProbe, see the *CA APM .NET Implementation Agent Guide*.

To configure AutoProbe tracing on a .NET agent:

1. Open the *IntroscopeAgent.profile* in the agent's *<Agent_Home>* directory.
2. Locate the *introscope.autoprobe.directivesFile* property, for example:

   introscope.autoprobe.directivesFile=
   default-typical.pbl

3. Locate the PBL file specified in the property.

   In the example above, it is *default-typical.pbl*.

4. Open the specified PBL file in a text editor.
5. In the PBL file uncomment the httpheaderdecorator.pbd line:

```
#######################
# HTTP Header Decorator
#
# Enable this PBD for CEM integration.
# It appends additional information to the
# HTTP response that is then captured by CEM.
# ===============
httpheaderdecorator.pbd
```

CA Technologies Pre-release Documentation. Authorized Use Only.

6. Save your changes to the PBL file.
7. Continue to Verifying default agent properties in .NET environments (see page 370).

Note: In CA Introscope 9.0, *biz-trx-http.pbd* is enabled by default. If you have pre-9.0 agents, then *CEMTracer.pbd* must also be uncommented.

Verifying default agent properties in .NET environments

The agent profile properties file allows you to modify default agent properties for AutoProbe and HTTPHeaderDecorator.

All the CA CEM-related properties are pre-configured for the CA CEM to CA Introscope integration. You need to verify that the default values are appropriate for your installation.

Important! You should review at least the security-related properties. For example, the *introscope.agent.decorator.security* property is set to clear by default. This is appropriate for initial testing, but might reveal information in the transaction header that you do not want known beyond the firewall. You can set the property to encrypted for a more secure production environment.

For more information about CA CEM-related agent properties, see CA Introscope agent profile properties.

For more information, see the *CA APM .NET Implementation Agent Guide*.

To verify the agent profile properties:

1. Open the *IntroscopeAgent.profile* in the agent's *<Agent_Home>* directory.
2. Compare the properties you find in the agent profile to the properties' descriptions in CA Introscope agent profile properties. For example, you might want to change these properties from the default settings:
   - *introscope.agent.decorator.enabled* — the default setting for a .NET agent is true; this enables the problem resolution triage feature.
   - *introscope.agent.decorator.security* — change the default setting from clear to encrypted; this secures confidential information in transaction headers.
3. If you need to modify default values, you might need to add the property to the *IntroscopeAgent.profile* file and then configure it.
4. Save your changes to the *IntroscopeAgent.profile* file.
5. Continue to Reloading agent properties in .NET environments (see page 371).

Tip: If a property is not listed at all, it can be confusing, so it's a good idea to explicitly include it, even if the default value is what you need.

CA Technologies Prerelease Documentation. Authorized Use Only

Reloading agent properties in .NET environments

You must restart the agent for the managed application in order to complete the integration.

Restarting the CLR reloads the agent properties and enables AutoProbe and HTTPHeaderDecorator.

To reload the agent profile properties on a .NET agent:

1. Restart the agent to reload the agent properties.
2. Return to Integrating with .NET agents (see page 367) and repeat the procedures for each agent.

Verifying the Introscope configuration

If you have not yet configured CA CEM for CA APM, go to Checklist for configuring CA CEM (see page 352).

See Troubleshooting the Introscope to CA CEM integration (see page 371) for information on verifying and troubleshooting the configuration.

Troubleshooting the CA Introscope to CA CEM Integration

You might need to troubleshoot the CA CEM integration with CA Introscope. This section provides high-level verification steps and a troubleshooting section.

This section contains the following topics:

Verifying CA CEM integration on CA Introscope (see page 371)
Verifying CA Introscope integration on CA CEM (see page 373)
Viewing transaction trace session information (see page 374)
Symptoms and solutions (see page 374)

Verifying CA CEM integration on CA Introscope

You can verify the integration on CA Introscope in several ways.

To verify CA CEM functionality on CA Introscope:

1. Verify that ServletHeaderDecorator / HTTPHeaderDecorator is adding GUID information to headers as expected. See Go to Introscope Workstation Web Start.
   a. Navigate to the Introscope Historical Query Viewer.
   b. Look for the CorGUID in the Query field.

CA Technologies Pre-release Documentation. Unauthorized Use Only.

c. Select CEM > Incident Management > Defects.

d. Search for all slow time defects.

e. Choose Introscope View.

f. Look for Yes values in the Transaction Trace column.

g. Click on Yes to see details.

In the HTTP Information section, look for a Response Header *x-wily-info*, displaying the GUID.

2. Verify that BizTrxHttpTracer is adding metrics to the Introscope Investigator tree as expected.

   a. See CA CEM problem resolution triage metrics on CA Introscope.

   b. See Go to Introscope Workstation Web Start.

   c. Verify that you see the expected metrics in the tree.

3. Verify that the domain configuration information was received by the agents.

   a. Check the log in debug or verbose mode.

   b. Look for entries with "BizTrxDef" in the message.

4. Verify that transaction traces can be started manually.

   a. Select CEM > Incident Management.

   b. Find a slow time incident and select it to see the incident overview.

   c. Click Start Transaction Trace. The transaction trace should start immediately.

   d. See Viewing transaction trace session information.

   e. Your transaction trace should appear in the list.

5. Verify that CA Introscope captures the transaction traces that you started in step 4.

   a. See Go to Introscope Workstation Web Start.

   b. Navigate to the Introscope Historical Query Viewer.

   c. Click on Trace View tab to see the expected transaction trace.

6. Verify that CA Introscope displays the customer experience metrics as expected.

   a. See Starting Introscope Workstation.

   b. Open a new Introscope Investigator window.

CA Technologies Proprietary Documentation
Authorized use only

c. Look for these nodes and metrics to appear under *SuperDomain* > *server_name* > CEM > TESS Agent (*SuperDomain*):

- Business Service — aggregated metrics from all the TIM monitors. (Note that you might see Business Process instead, if you have reset the *BTStatsMetricName.backwardCompatible* property. See Upgrading with customer experience metrics (see page 356).)
- TIM — one entry for each active TIM monitor you have.

7. See Navigating the Introscope Workstation Investigator.
8. Verify that CA Introscope displays the real-time transaction dashboards as expected.
   a. See Starting Introscope Workstation.
   b. In the Introscope Console Dashboard box, select CEM Overview.
   c. Verify that real-time transaction data appears in the dashboard.
9. See Navigating the Introscope Workstation Console.

Verifying CA Introscope integration on CA CEM

You can verify the integration on CA CEM in several ways.

To verify CA Introscope functionality on CA CEM:

1. Verify that the Enterprise Manager is configured correctly.

   See Configuring CA Introscope Settings (see page 353).
2. Verify communications with the Enterprise Manager.
   a. Select System > Events.
   b. Ensure that there are no outstanding communications errors between the Enterprise Manager and the agents.
   c. Ensure that there are no outstanding communications errors between the Enterprise Manager and the TIMs.
   d. See Viewing CA CEM events (see page 311).
3. Synchronize all monitors to send the latest domain configuration information, which includes transaction definitions, defect specifications, and transaction trace parameters, to the agents. (It also sends to all enabled TIM monitors.)
   a. Select Setup > Monitors.
   b. Click Synchronize All Monitors.
   c. Verify that Enterprise Manager-to-Agent communications are active. See the agent log file.

CA Technologies Prerelease Documentation. Authorized Use Only.

4. Check the Enterprise Manager log for any related error messages. Search for CA CEM-related messages. For example, search the log for "tess" or "Biz" in the messages.
5. Verify that transaction trace sessions are running when expected. See Viewing transaction trace session information.
6. Verify that the real-time transaction metrics feature is enabled via the TIM Collection Service. See Selecting the TIMs for customer experience metrics (see page 257).
7. Verify that the real-time transaction metrics feature is enabled on each TIM.

   See Troubleshooting customer experience metrics on the TIMs (see page 357).

Viewing transaction trace session information

You can view a list of the transaction trace sessions that are currently running. You can also select a specific incident to get more information about it.

To view all the transaction traces in progress:

1. Select CEM > Incident Management.

   The Incidents page appears.
2. Select the Transaction Trace Sessions link.

   A list of the currently running transaction traces appears.
3. Click on a specific incident number to go directly to the CEM > Incident Overview page for that incident.

Symptoms and solutions

The following topics provides some tips that might help you troubleshoot the CA Introscope integration with CA CEM.

Categories

.NET (see page 375)
Events & logs (see page 375)
Customer Experience Metrics (see page 376)
Transactions & traces (see page 377)
User interface (see page 380)

CA Technologies Pre-release Documentation. Authorized Use Only.

.NET

Symptom

Transactions for .NET are not recognized when they are defined to match on host name of IP address. (Transactions that are monitored via Java agents under the same conditions are recognized and appear in the Investigator tree as expected.)

Possible solution

Change the Host name to match on UPPERCASE host name instead of IP address. (Business Services > Business Transactions > Transactions > Components > Identification).

Symptom

The .NET agent does not restart if the .NET wily.Agent assembly version does not match the version number of the .NET agent extension for CA CEM.

Possible solution

Follow the instructions as defined in Configuring the agent .NET version (see page 368).

Note: The steps must be performed exactly. If you do not update the machine.config file properly, the .NET agent will not start.

The larger issue on CA Introscope is:

Periodically, the Standalone agent installer does not install extensions (such as DLLs) correctly if the version number of the extension differs from the version number of the .NET agent.

There are two workarounds for this situation:

- configure individual applications to accept different version numbers.

OR

- configure globally all applications to accept different version numbers.

Note: Choose only one option; do not perform both options.

For more information, see the *CA APM .NET Agent Implementation Guide*.

Events & logs

Symptom

Enterprise Manager log file message appears when synchronizing monitors:

```
Failed to send domain configuration xml file to the collector Enterprise Manager.
The requisite service was not defined by the Enterprise Manager.
```

Note: This can happen when you have agents from CA CEM 4.x.

CA Technologies Reference Documentation
Authorized Use Only

Possible solution

Verify that the CEMDefinitionHandler Enterprise Manager extension has been installed on the MOM and each Collector, in that order.

You must also restart the JVM for each Enterprise Manager in order to complete the integration.

For more information, see Configuring the Enterprise Managers (see page 363).

Customer Experience Metrics

Symptom

On the Introscope Investigator tree, business transaction metrics are either grayed out or do not appear.

Possible solution

Check the IntroscopeAgent.log for connection errors to the specified Enterprise Manager.

Symptom

On the Introscope Investigator tree, business transaction metrics are either grayed out or do not appear.

Possible solution

If you have changed the default property: *BTStatsMetricName.backwardCompatible* so your metrics will retain the 4.x "Business Process" name, then you need to also use the 4.5 Javascript calculators. See Upgrading with customer experience metrics (see page 356).

Symptom

On the Introscope Investigator tree, business transaction metrics are either grayed out or do not appear.

Possible solution

Verify that the Enterprise Manager TIM Collection Service is started and enabled on the TIMs.

Look for these nodes and metrics to appear under *SuperDomain* > server_name > CEM Agent:

- Business Segment — aggregated metrics from all the TIM monitors.
- TIM — one entry for each active TIM monitor you have.

The Business Segment aggregation node turns gray when there is no new data.

CA Technologies Prerelease Documentation. Authorized Use Only.

The TIM entries should be zero when there is no data (no transactions being monitored), but they should not be gray.

When you restart the Enterprise Manager, metrics go away.

For more information, see Verifying CA CEM integration on CA Introscope (see page 371).

Symptom

On the Introscope Investigator tree, Business Service appears where Business Process is expected.

Possible solution

As of the CA APM 9.0 (CA CEM 5.0) release, "business process" in CA CEM changed to "business service."

If you prefer to see the metrics with the pre-9.0 names, you can change this.

For more information, see Upgrading with customer experience metrics (see page 356).

Transactions & traces

Symptom

Do not see the Business Segment node in the Investigator, nor any transaction trace data in CA CEM.

Possible solution

Restart the Enterprise Manager and Synchronize All Monitors, in that order. (There must be changes in the domain configuration in order for Synchronize All Monitors to send the information.)

Symptom

For some business transactions, the TIM and the Enterprise Manager do not recognize them in the same way.

For example, CA Introscope might recognizes a business transaction as TEST/txLang_confirm.jsp while CEM recognizes the same transaction with TEST/txLang.

Possible solution

The TIM monitors network traffic before the load balancer. The agent monitors traffic behind the load balancer. The load balancer can redirect the URL that comes in, or change it completely.

Try matching the business transaction with another parameter definition (POST parameters or HTTP headers for example).

CA Technologies Prerelease Documentation
Authorized Use Only.

Symptom

Catch-all transactions are not working as expected with transaction traces.

Catch-all transactions are defined in CA CEM to identify transactions during the set-up phase as well as to catch any transactions that haven't been defined.

A generic catch-all transaction definition can potentially appear higher up in the list of transaction definitions rather than at the bottom of the list. This could result in practically un-defining the transaction definitions that appear in the list below the catch-all transaction definition.

Possible solution

As of CA CEM 4.2, there is a sorting algorithm in CEMTracer where transaction definitions are sorted by business application and then by their length. The more specific a transaction definition (more parameters specified), the higher it appears in the list within that business application.

All catch-all transaction definitions should be made application specific by specifying the host name, the port, the application context URL path, or any parameter used by the application to identify a specific business application in the catch-all transaction definition. A combination of the above parameters can also be used to uniquely associate a catch-all transaction with a business application as defined in CA CEM.

Another recommendation: Do not define business service / business transactions with the same exact definition.

Another recommendation: Make business service / business transactions on the same business application even more specific than the catch-all transaction. The more specific a business transaction definition is, the higher it is on the list.

If the catch-all definition is application specific, but it is also more specific (matches host, port, URL path) than other business service / business transactions on that business application, the catch-all would still end up higher on the list.

When defining business service / business transactions other than the catch-all, match the parameters that are matched by the catch-all, plus add more matching parameters. This should make it both application specific, and also more specific than the catch-all.

Symptom

Transaction traces stopped before the maximum session duration elapsed.

Possible solution

If you restart the Enterprise Manager, then the transaction traces will stop.

Go to the Incident > Overview page and click Start Transaction Trace to restart.

CA Technologies Pre-release Documentation. Authorized Use Only

Symptom

The *domainconfig.xml* file does not match the current CA CEM configuration.

Note: If your integration is using mixed mode support, allowing both 4.0 / 4.1 CEMTracer and 4.2 / 4.5 CEMTracer on the agents, then you need to understand which domainconfig.xml file to look for and where to look for it. The *domainconfig.xml* version is v1.9. See CA CEM domain configuration information for 4.0 through 4.5 releases (see page 382).

Possible solution

When you click Synchronize All Monitors, the domain configuration information gets sent automatically through the MOM / Enterprise Manager to the agents.

To correct:

1. Check System > Events to make sure communications are OK with the Enterprise Manager.
2. Check the communications between the Enterprise Manager and the agents.
3. Go to the Setup > Introscope Settings page and verify the CEMTracer Support.
4. Click Save Settings.
5. Go to the Setup > Monitors page and click Synchronize All Monitors. (There must be changes in the domain configuration in order for Synchronize All Monitors to send the information.)
6. Verify that the new *domainconfig.xml* file was received by the agents. Check the:
   - date and time stamp on the files.
   - Enterprise Manager and agent logs.

Symptom

The *domainconfig.xml* file on 7.2.x agents does not match the current CA CEM configuration.

(These agents should be using CEMTracer 4.0 / 4.1.)

Possible solution

CA CEM might be reloading the *domainconfig.xml* file too infrequently (or too often).

The agent property *introscope.agent.cemtracer.domainconfigfile.reloadfrequencyinminutes* is set to once per minute by default.

See the CA Introscope agent profile properties.

Symptom

The domainconfig.xml seems OK on one 7.2.x agent, but not on another.

(These agents should be using CEMTracer 4.0 / 4.1.)

CA Technologies Proprietary Documentation. Authorized Use Only

Possible solution

- Verify that the agent's <*Agent_Home*> directory is not read-only.
- Verify that it is a compatible agent (that is, not a .NET agent prior to 7.2.2, an EPA agent, or other non-Java agent).
- Verify the Enterprise Manager-to-Agent communications with the agent that does not seem to have the correct domainconfig.xml file.

Symptom

The agent is looking for the wrong domain configuration file.

Possible solution

The agent property *introscope.agent.cemtracer.domainconfigfile* should be set to *domainconfig.xml* (not *domainconfig-introscope.xml*) by default.

Change the agent property to *domainconfig.xml*.

See the CA Introscope agent profile properties.

User interface

Symptom

Workstation Web Start link does not appear on the incident and defect pages.

Possible solution

Make sure the Enable Workstation Web Start Link check box has been selected on the Introscope Settings page and the setting saved.

Symptom

WebView link does not appear on the incident and defect pages.

Possible solution

Make sure the Enable WebView Link check box has been selected on the Introscope Settings page and the setting saved.

Symptom

Workstation Web Start is not working.

Possible solution

Make sure the Workstation Web Start port is set to the correct port number (8081 is the default). Verify Java 5 or later is installed on your PC.

CA Technologies Prerelease Documentation. Authorized Use Only.

Symptom

WebView is not working.

Possible solution

Make sure the WebView Port is set to the correct port number (8080 is the default) on the Introscope Settings page and the setting saved. Also verify that the WebView process is running.

Symptom

WebView displays duplicate metrics (domain, business service, business transaction, application, and components).

Possible solution

Agent aggregation is configured via agentclusters.xml. The file should be placed only on the MOM, not the MOM and Enterprise Manager Collectors.

If you modify agentclusters.xml on more than one Enterprise Manager, you might see duplicate metrics. See Establishing CA CEM as a virtual agent (see page 364).

You must configure the *agentclusters.xml* file as specified in the Introscope documentation.

Symptom

After logging in to WebView, get a 404 error.

Possible solution

Context path is not correct. Verify that the Introscope Settings page configuration for Webview is "/" (no quotes) and save the setting.

Symptom

Incident > Overview > Go to Introscope Workstation Web Start link goes to the Introscope Console; it does not link to the Introscope Workstation, displaying the Investigator with the related business services for the CA CEM agent.

Possible solution

If you do not configure agentclusters.xml, the default landing page is the Introscope Console. See Establishing CA CEM as a virtual agent (see page 364).

You must configure the *agentclusters.xml* file as specified in the Introscope documentation.

Symptom

Incident > Overview > Go to Introscope Workstation Web Start link does not link to detailed business transaction metrics.

Possible solution

You need to click to navigate down to the detailed business transaction metrics that correlate to the incident.

CA Technologies Prerelease Documentation. Authorized Use Only.

Symptom

Incident > Overview > Go to Introscope Workstation Web Start link displays a "Page cannot be found" error.

Possible solution

Check the Java Web Start configuration. See Java documentation.

Upgrade considerations

Now that the TESS is part of the Enterprise Manager:

- Functionality previously provided with CEMTracer is now supplied by BizTrxHttpTracer, which is enabled by default. (BizTrxHttpTracer also supports the agent recording feature.)
- There is no longer a need to forward the domain configuration information file (*domainconfig.xml*) with 5.0 agents. When you click Synchronize All Monitors, the CA CEM notifies the Enterprise Manager, and the information is stored in the shared APM database.
- Backward compatibility for CEMTracer is provided. You can use CEMTracer 4.0/4.1 if you select the "mixed mode" integration option. You can use CEMTracer 4.2/4.5 by default.
- If you have CEMTracer 4.2/4.5 agents, you need CEMDefinitionHandler.jar 4.5.
- There is no longer a TESS log file. The messages that were previous written to the TESS log file are now written to the Enterprise Manager log file located in *<EM_Home>/logs*.

Before you upgrade CA CEM or CA Introscope for CA APM integration, you should review the specific product upgrade requirements in the *CA APM Installation and Upgrade Guide*.

Planning your upgrade before you actually begin the individual procedures is highly recommended in the case of CA CEM to CA Introscope integration. Taking into consideration your own unique environment and particular requirements, you can avoid unnecessary confusion along the way.

CA Technologies Pre-release Documentation. Authorized Use Only

CA CEM domain configuration information for 4.0 through 4.5 releases

CA CEM automatically forwards the domain configuration file (*domainconfig.xml*) file to the MOM or Enterprise Manager when you click Synchronize All Monitors on the Setup > Monitors page. CA CEM also sends the same domain configuration file to the TIM when you click Synchronize All Monitors.

If there is an MOM, then the MOM sends the file to all the Enterprise Manager Collectors that are connected to the MOM. The Enterprise Manager sends parsed transaction related information to registered agents.

This file contains all the information that Introscope needs to recognize transactions for CA CEM.

Important! Previous product integrations between CA CEM and CA Introscope involved manual manipulation of the domainconfig.xml files. *Manual manipulation of these files is no longer supported*.

Integration changes as of CA CEM 4.2

CA CEM 4.2 introduced the "mixed mode" option that allows old and new *domainconfig.xml* file handling at the same time. This means you do not need to convert all your CA Introscope agents at once.

Instead of the CA CEM 4.0 / 4.1 "push" method of delivering CA CEM transaction definitions (in domainconfig.xml) to agents, the CA CEM 4.2 / 4.5 method changes to a "pull" method, in which agents register with the Enterprise Manager to indicate that they require the CA CEM transaction definitions. Also in 4.2, the *domainconfig.xml* information is stored on the Enterprise Manager disk; agents store the CA CEM transaction definitions rule set in memory.

Upgrade Support in CA CEM 4.2 / 4.5

The CEMTracer 4.0 / 4.1 Support option allows you to stagger your Introscope agent migration from 4.0 or 4.1 to 4.2 / 4.5 over time. If the Introscope Setting is enabled when upgrading to CA CEM 4.2 / 4.5, then mixed mode integration is enabled by default.

Important! For existing 4.0 / 4.1 integrations, you must ensure that CEMTracer is from the CA CEM 4.0q9 or later release. (In the *CEMTracer.pbd* file on each Introscope agent, you should see: Version 4.0.Q9 or later.)

If you have an older version of CEMTracer, we recommend you upgrade the agent to CEMTracer 4.5. The agent must also be on Introscope 7.2.0 or later in order to use the new CA CEM 4.5 CEMTracer.jar.

If you have selected the CEMTracer 4.0 / 4.1 Support option on the Introscope Settings page, then it is a mixed mode integration, meaning that both the "push" and "pull" methods are in effect at the same time. You should use this feature only as long as it is required. Mixed mode means:

- Setup > Monitors > Synchronize All Monitors sends *domainconfig.xml* (version 1.9) to MOM / Enterprise Manager.
- both "push" and "pull" methods in play, which depends on the agents' CEMTracer version.

CA Technologies Prerelease Documentation
Unauthorized Use Only

- Enterprise Manager stores domainconfig.xml v 1.9 file on disk and pushes it to all agents with the *introscope.agent.remoteagentconfiguration.enabled* property set to true.
- agents with CEMTracer 4.0 / 4.1 enabled store the *domainconfig.xml* file on disk.
- agents with CEMTracer 4.2 / 4.5 enabled store the transaction definition rule set in memory.

Things to Know About Upgrade Support in CA CEM 4.2 / 4.5

You should know this about upgrade support:

- If CEMTracer 4.0 / 4.1 is not enabled on the agent, then no action will be taken on the domain configuration file once it is sent (for example, if it were received by a CEMTracer 4.2 agent, a .NET agent prior to 7.2.2, EPA agent, or other agent).
- If an agent is not connected to the Enterprise Manager, the domain configuration file (*domainconfig.xml* v 1.9) does not get sent.
- If the agent directory is read-only, the domain configuration file cannot be written to disk.
- If you do not want an agent to receive the *domainconfig.xml* file, set the *introscope.agent.remoteagentconfiguration.enabled* property to *false*.
- For non-compatible agents (that is, a .NET agent prior to 7.2.2, an EPA agent, or other non-Java agent), you should set the *introscope.agent.remoteagentconfiguration.enabled* property to *false*.

Creating a custom agent for each TIM on CA SiteMinder

These steps describe how to create a custom agent (described in CA SiteMinder documentation as an agent identity or an agent object).

You must create a custom agent for each TIM.

To create a custom agent for each TIM on CA SiteMinder:

1. Open a browser and enter the CA SiteMinder URL on the policy server. For example:

   `http://<IP_Address>/siteminder/smadmin2.html`
2. Click Administer Policy Server and use the Java applet to define a new agent.
3. On the Object pane (left), click the System tab.
4. Click the Agent icon on the Object pane.
5. From the Edit menu, select Create Agent.
6. Type the Name of the TIM agent: *TIM <TIM_IPaddress>*.
7. Enter a Description of the TIM agent.

CA Technologies Proprietary/Confidential Document. Authorized Use Only.

8. Verify that the Support 4.x agents check box is selected.
9. Select the CA SiteMinder option and select the custom agent—for example, *TIMAgent*—from the list.
10. In the IP Address or Host Name text box, enter the IP address of the TIM.
11. In the Secret and Confirm Secret text boxes, enter a password that will also be used by the custom agent—for example, *quality1*.

    Tip: Make note of the secret; you will need it when you configure the CA SiteMinder plug-in.
12. Click OK.

    It is not necessary to restart the policy server.

Setting up the business application

When you set up the business application as described in the *CA APM Transaction Definition Guide*, ensure that you use the CA SiteMinder settings.

To confirm correct settings for the business application:

1. Ensure that the Authentication Type is CA SiteMinder.
2. Ensure that the User Identification parameter type is CA SiteMinder and the parameter name is UserName or DistinguishedName.

   For some types of CA SiteMinder authentication (anonymous authentication and certificate-based authentication), the SM_AGENTAPI_ATTR_USERNAME is blank. In this case, the value of SM_AGENTAPI_ATTR_USERDN (the distinguished name) must be used to identify the user.
3. Ensure the Session Identification parameter type is CA SiteMinder and the parameter name is SessionId.

CA Technologies Proprietary Documentation
Authorized Use Only

Chapter 25: Integrating CA SiteMinder with CA CEM

CA CEM works with CA SiteMinder security software. CA SiteMinder provides security features such as single sign-on and centralized control of user access to web applications.

In simple authentication cases, you may be able to use cookies to identify users and sessions without the CA SiteMinder plug-in (and without following the procedures in this chapter).

If the user name and session ID are always observed in the header and if the session ID does not change or expire, then consider using cookie parameters to identify users and sessions, for example, SMUSER and SMSESSION or HTTP_SM_USER and HTTP_SM_SESSION.

If single sign-on (SSO) or expiring cookies are used, then you do need to use the CA SiteMinder plug-in as described in this chapter. In these cases, the simpler cookie approach does not work because the request headers do not consistently contain the user name and a static cookie. Then, the TIM needs to communicate with the CA SiteMinder server to determine user and session.

This section contains the following topics:

Checklist for CA SiteMinder (see page 387)
Configuring CA SiteMinder on CA CEM (see page 388)
Creating custom CA SiteMinder agents (see page 389)
Defining policy servers on CA CEM (see page 389)
Defining agents on CA CEM (see page 390)
Enabling CA SiteMinder (see page 390)
Troubleshooting the CA SiteMinder to CA CEM integration (see page 391)

CA Technologies Pre-Release Documentation
Authorized Use Only.

Checklist for CA SiteMinder

This list shows the things you need to know about CA CEM integration with CA SiteMinder.

To integrate with CA SiteMinder:

1. In CA SiteMinder, create a custom agent type for CA CEM. See Creating a custom agent type for CA CEM on CA SiteMinder (see page 389).
2. In CA SiteMinder, create a custom agent for each TIM. See Creating a custom agent for each TIM on CA SiteMinder (see page 384).

3. Define the shared secret. See Configuring CA SiteMinder on CA CEM (see page 388).
4. Define the policy servers. See Defining policy servers on CA CEM (see page 389).
5. Define the agent names. See Defining agents on CA CEM (see page 390).
6. Enable CA SiteMinder. See Enabling CA SiteMinder (see page 390).
7. Set up the business application to specify CA SiteMinder authentication. See Setting up the business application (see page 385).
8. Verify that CA CEM is working with CA SiteMinder. See Troubleshooting the CA SiteMinder to CA CEM integration (see page 391).

Configuring CA SiteMinder on CA CEM

In order for CA CEM to work with CA SiteMinder you must configure the CA SiteMinder plug-in.

To configure CA SiteMinder on CA CEM:

1. Select Setup > Plug-ins.

   The Plug-ins page appears.

2. Select CA SiteMinder.

   The CA SiteMinder Plug-in > General page appears.

3. Select the Enable CA SiteMinder Plug-in check box, which activates CA CEM and CA SiteMinder integration.
4. Type the Session Cookie Name, which stores CA SiteMinder user session information in encrypted form.

   Note: The default session cookie name is SMSESSION.

5. Type the Shared Secret, which is used together with the custom agent name for encrypting traffic between the custom agent and the CA SiteMinder policy server.

   The shared secret is a text value that must match the value given when configuring the CA CEM custom agent(s) in CA SiteMinder.

   Note: You must use the same shared secret for all CA CEM agents when configuring CA SiteMinder.

6. Click Save.

CA Technologies Prerelease Documentation. Authorized Use Only.

Creating custom CA SiteMinder agents

These steps describe how to integrate CA CEM with CA SiteMinder policy servers. You perform these steps on the CA SiteMinder policy server. For more information, see the CA SiteMinder documentation.

Creating a custom agent type for CA CEM on CA SiteMinder

You need to perform the following steps once. If you have multiple TIMs, you do not need multiple custom agent types.

Note: The following procedure describes the steps in CA SiteMinder r6. If you have a later release, the exact steps might be different, but the same functions must be performed.

To create a custom agent type for CA CEM on CA SiteMinder:

1. Open a browser and enter the CA SiteMinder URL on the policy server. For example: `http://<IP_Address>/siteminder/smadmin2.html`
2. Click Administer Policy Server and use the Java applet to define a new agent type.
3. On the Object pane (left), click the System tab.
4. From the View menu, verify that Agent Types is checked.
5. Click the Agent Types icon on the Object pane.
6. From the Edit menu, select Create Agent Type.
7. Type the Name the custom agent uses to identify itself—for example, *TIMAgent*.
8. Enter a Description of the custom agent type.
9. On the Agent Type Definition tab, leave the RADIUS Device check box cleared and leave the IETF Vendor ID text box empty.
10. If no actions are defined, then define a new action and set the Type to Authenticated.

CA Technologies Pre-release Documentation. Authorized Use Only.

Defining policy servers on CA CEM

A policy server is a CA SiteMinder component providing authorization and authentication for web applications. A CA SiteMinder configuration may include more than one policy server.

To define a policy server:

1. Select Setup > Plug-ins. The Plug-ins page appears.

2. Select CA SiteMinder.
3. On the CA SiteMinder Plug-in page, select Policy Servers.
4. Click New.

   The Policy Server page appears.
5. Enter the IP address and authentication port of the computer on which the policy server is installed.
6. Click Save.

Defining agents on CA CEM

A CA SiteMinder agent is a software module that authorizes requests for a secured resource and forwards those requests to the CA SiteMinder policy server.

On the Plug-ins Agents page, agents and associated monitors (TIMs) are listed. Initially, agents and monitors have the same names. You must change the agent name so that it exactly matches the name used by CA SiteMinder for that agent.

To define an agent:

1. Select Setup > Plug-ins.

   The Plug-ins page appears.
2. Select CA SiteMinder.
3. On the CA SiteMinder Plug-in page, select Agents.

   The Agents page appears.
4. On the Agents page, click on the IP Address of an agent.

   The Monitor page appears.
5. Update the CA SiteMinder Agent Name so that it exactly matches the name used by CA SiteMinder.
6. Click Save.

Enabling CA SiteMinder

In order for CA CEM to work with CA SiteMinder you must enable the CA SiteMinder plug-in.

To enable CA SiteMinder on CA CEM:

1. Select Setup > Plug-ins.

   The Plug-ins page appears.

CA Technologies Prerelease Documentation. Authorized Use Only

2. Select CA SiteMinder.
3. Select the Enable CA SiteMinder Plug-in check box.
4. Click Save.

Troubleshooting the CA SiteMinder to CA CEM integration

To verify that CA CEM is working with CA SiteMinder:

1. Verify that CA SiteMinder is enabled.

   See Enabling CA SiteMinder (see page 390).

2. Verify that the policy servers are using the correct IP addresses and ports.

   See Defining policy servers on CA CEM (see page 389).

3. Verify that the shared secret used by both CA CEM and CA SiteMinder is the same.

   See Creating custom CA SiteMinder agents.

4. Verify that the agent names used by both CA CEM and CA SiteMinder are the same.

   See Defining agents on CA CEM (see page 390).

5. Verify that the business application has been correctly set up.

   See Setting up the business application (see page 385).

6. In CA CEM, select Setup > Monitors > Synchronize All Monitors to synchronize monitors.
7. In CA CEM, select Setup > Monitors > IP Address to check the TIM log for a confirmation message or for any CA SiteMinder error messages.
8. Check the CA SiteMinder agent log for error messages.
9. In CA SiteMinder, check the policy server log for errors.

CA Technologies Pre-release Documentation. Authorized Use Only.

Chapter 26: Integrating CA Service Desk with CA CEM

You can send CA CEM incident and defect information to CA Service Desk, so tickets can be generated and tracked as appropriate.

CA Service Desk employs the Information Technology Infrastructure Library (ITIL) best practice processes by automating service support incident management.

This section contains the following topics:

About CA Service Desk with CA CEM (see page 393)
CA Service Desk integration overview (see page 393)
Configuring CA Service Desk for CA CEM (see page 394)
Using CA Service Desk with CA CEM (see page 397)
Troubleshooting the CA Service Desk to CA CEM integration (see page 398)
CA Service Desk to CA CEM integration terminology (see page 400)

About CA Service Desk with CA CEM

CA Service Desk supports help desk analysts perform their incident management responsibilities including detecting and recording, classifying and supporting, tracking and resolving incidents.

You can configure CA CEM and CA Service Desk to create service requests (or incidents) in CA Service Desk based on incidents detected by CA CEM.

CA Service Desk integration overview

Integration with CA Service Desk provides incident management support by automatically generating service requests from CA CEM incidents.

To enable CA CEM incidents to be created in CA Service Desk, you need to configure and enable the CA Service Desk plug-in. This establishes communication and defines the way incidents appear in CA Service Desk.

If you want to separate CA CEM generated incidents in CA Service Desk, you need to create a request area in CA Service Desk.

For integration-related terms, see CA Service Desk to CA CEM integration terminology (see page 400). See the *CA APM Glossary* for all CA APM glossary terms and definitions.

CA Technologies Pre-release Documentation. Authorized Use Only.

Configuring CA Service Desk for CA CEM

To enable CA Service Desk, you first must configure the CA Service Desk plug-in.

Optionally, you can create CA Service Desk request areas to separate CA CEM incidents from other service requests in CA Service Desk.

Configuring the CA Service Desk plug-in

You need to gather connection and request information to implement the CA Service Desk plug-in.

To configure CA CEM to communicate with CA Service Desk:

1. Select Setup > Plug-ins.

   The Plug-ins page appears.

2. Select CA Service Desk.

   The CA Service Desk Plug-in page appears.

3. Select the Enable CA Service Desk Plug-In check box.

   This allows CA CEM to send incident information to CA Service Desk.

   Note: You must enable the CA Service Desk plug-in or the integration will not function.

4. Externally Resolvable Hostname — in most cases, this field already contains the CA CEM host name. If it does not, type the IP address or host name of the Enterprise Manager.

   This information is used for links in CA Service Desk-generated email. The CA Service Desk email includes a link back to the incident in CA CEM.

5. Port — in most cases, this field already contains the CA CEM port.

   Note: The default port is 8081. It is defined in the IntroscopeEnterpriseManager.properties file as *introscope.enterprisemanager.webserver.port=8081* and can be changed.

6. Enter the CA Service Desk connection information:

   a. Hostname — type the host name or IP address of the CA Service Desk system that hosts CA CEM.

   b. Port — type the port where CA CEM connects to CA Service Desk. The default port number for CA Service Desk is 8080.

CA Technologies Prerelease Documentation. Authorized Use Only.

c. User Name — type the CA CEM user name for CA Service Desk access. This user should have the appropriate permissions and access levels in to create requests in CA Service Desk.

d. Password — type the password associated with the user name.

7. If you want CA CEM incidents to be created in CA Service Desk as incidents (instead of requests), select the Create ITIL Incidents in Service Desk check box.

   If you do not want incidents created in CA Service Desk, then the default behavior is to create requests.

   Note: CA Service Desk incidents will be created only if CA Service Desk has been configured for ITIL (Information Technology Infrastructure Library). If not, CA CEM incidents will be created as requests.

8. Type the Request Area — type the request area created in CA Service Desk to classify CA CEM related issues.

   This request area should be the same name as you create in CA Service Desk.

9. Type a Request Filter if required.

   Tip: You can use the request filter to test the integration with a single business service and business transaction first, rather than allow all incidents to be forwarded to CA Service Desk immediately.

Syntax Example	**Description**
Request Filter= BusinessServiceName/ BusinessTransactionName	Creates a request in CA Service Desk only if the incident occurs for the specified business service and business transaction.
*Request Filter=BusinessServiceName/**	Creates a request in CA Service Desk for all incidents and any business transaction triggered for the specified business service.
*Request Filter= BusinessServiceName/Tra**	Creates a request in CA Service Desk for all incidents and any business transaction that begins with "Tra" triggered for the specified business service.
Request Filter=/**	Creates a request in CA Service Desk for incidents associated with all business services and all business transactions.

   If the business service or business transaction name contains a forward slash ("/") character, you must replace it in the filter with a double-slash ("//").

   If the business service or business transaction name contains a comma (",") character, you must replace it in the filter with a slash-comma ("/,").

CA Technologies Prerelease Documentation. Authorized Use Only.

For example, for:

- business service name = Trade
- business transaction name = Sell Big/Medium Stocks

The filter should be:

Request Filter=Trade/Sell Big//Medium Stocks

For example, for:

- business service name = Trade
- business transaction name = Sell Big, Medium Stocks

The filter should be:

Request Filter=Trade/Sell Big/, Medium Stocks

10. Type a Request Summary, which will appear in the CA Service Desk summary. Use the variables to create a customized summary:
    - *<defect_name>*
    - *<business_service>*
    - *<business_transaction>*

    You can rearrange the variables to create your own request summary. Each variable will be replaced with actual values when the request summary appears in CA Service Desk. Examples:

    <defect_name> for <business_service>,<business_transaction>

    Appears like this: Slow Time for Support Desk, KB Lookup

    Business Transaction: <business_transaction> Defect: <defect_name>

    Appears like this: Business Transaction: CheckOut Defect: Slow Time
11. Type a Request Description, which will appear in the CA Service Desk description. Use the variables to create a customized description:
    - *<incident_id>*
    - *<defect_name>*
    - *<business_service>*
    - *<business_transaction>*
    - *<user_count>*
    - *<defect_count>*
    - *<incident_link>*

CA Technologies Pre-release Documentation. Authorized Use Only.

You can rearrange the variables to create your own request description. Each variable will be replaced with actual values when the request description appears in CA Service Desk.

The <incident_link> variable will provide a link in CA Service Desk. You can use the link to go directly to the incident overview information (CEM > Incident Management > Incident Overview).

12. Create a Service Desk request if incident Severity is — enable the selected incident levels to create service requests in CA Service Desk. Select the appropriate check boxes:
    - Low
    - Moderate
    - Severe
    - Critical
13. Click Save.

    Future CA CEM incidents matching your saved configuration will trigger service requests in CA Service Desk, effective immediately.

    Note: If there are any existing open incidents in CA CEM, they will not trigger CA Service Desk requests.

Configuring a CA Service Desk request area

If you want a separate request area for the incidents reported by CA CEM, you need to create a request area in CA Service Desk.

To configure a request area in CA Service Desk for CA CEM:

1. Create a request area, for example *CEM_reported*.
2. Attach three properties to the request area:
    - Business Service
    - Business Transaction
    - Incident

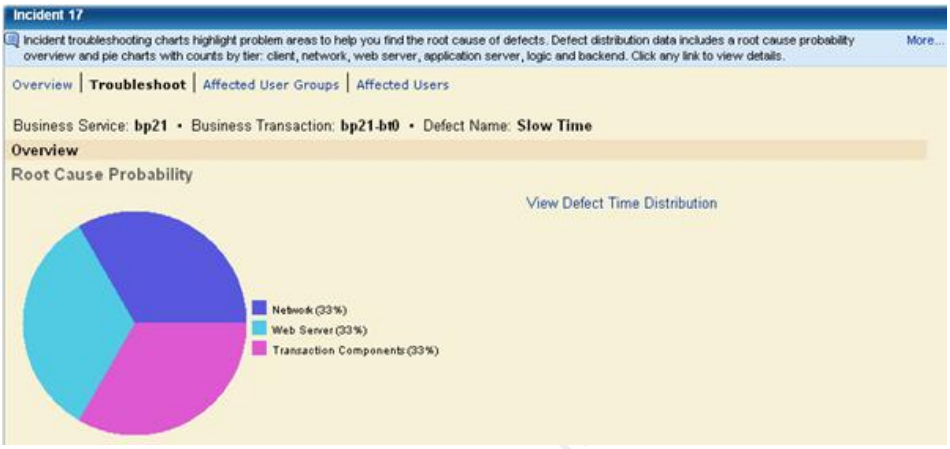
See *Defining Request Areas* in CA Service Desk Online Help for information about creating and configuring request areas.

Using CA Service Desk with CA CEM

In CA Service Desk, you can see CA CEM incident information. You can click a link from CA Service Desk to CA CEM to learn more about the incident.

CA Technologies Prerelease Documentation Authorized Use Only

You can navigate around CA CEM to troubleshoot and discover the root cause of the incident.



Troubleshooting the CA Service Desk to CA CEM integration

You might need to troubleshoot the integration with CA Service Desk. This section provides a troubleshooting table.

CA Technologies Prerelease Documentation
Authorized Use Only

Symptoms and solutions

This table provides some tips that might help you troubleshoot the CA Service Desk integration with CA CEM.

Symptom:

The CA Service Desk request is not created for any or some CA CEM incidents.

Solution:

Verify that the CA Service Desk plug-in is properly configured in CA CEM:

The correct CA Service Desk credentials are configured.

The filter is set for the correct business services and business transactions.

The appropriate severity levels are enabled.

Check the log file for connection or other errors:

http://<IP_Address>/cgi-bin/wily/packages/cem/service-desk/index

Symptom:

The CA Service Desk requests are created without a request area.

Solution:

Confirm that the request area was created by the CA Service Desk Administrator.

Verify that the CA CEM CA Service Desk plug-in configuration includes the correct request area information.

Symptom:

The link to the CA CEM incident detail from the CA Service Desk Incident or Request does not work.

Solution:

The incident might have been deleted from the CA CEM database.

Check the Setup > Incident Settings > Incident Age-out Rules.

Symptom:

The link to the CA CEM incident detail from the CA Service Desk email does not work. The link in the email might look something like this:

http://localhost.localdomain/wily/cem/tess/...

Solution:

Verify that the CA CEM CA Service Desk plug-in configuration includes the Externally Resolvable Hostname for CA CEM.

CA Technologies Prerelease Documentation. Authorized Use Only

The incident might have been deleted from the CA CEM database. Check the Setup > Incident Settings > Incident Age-out Rules.

Symptom:

When I close an incident in CA CEM, it does not close in CA Service Desk.

Solution:

You must close the incident / service request separately in CA Service Desk.

Symptom:

When I close an incident / service request in CA Service Desk, it does not close in CA CEM.

Solution:

You must close the incident separately in CA CEM.

CA Service Desk to CA CEM integration terminology

An important aspect of product integration is terminology. You need to learn the language of the products that are unfamiliar to you. You also need to be aware of how terms and definitions in CA CEM map to the terms in CA Service Desk.

CA Service Desk terms

These CA Service Desk terms are related to the product integration.

incident

Any event that causes a reduction in the quality of the service of your IT infrastructure.

request area

The CA Service Desk system administrator sets up request areas to categorize service requests by type. For example, hardware support or software install.

service request

A request from a user for support, information, or assistance and does not represent a failure in the IT infrastructure.

ticket

Ticket is a general term in CA Service Desk, which can mean a request, incident, problem, or an issue. CA CEM incidents in CA Service Desk are classified either as a request or an incident.

CA Technologies Pre-Release Documentation. Authorized Use Only.

CA Service Desk to CA CEM term mapping

Products can have different evolution cycles and different core audiences. These differences often mean that the product terms and definitions also differ.

Sometimes there is a one-to-one mapping (for example an incident in one product might be a ticket in another). And sometimes the same term can have a slightly different meaning in a different product, which can be confusing.

These term mappings will help you navigate the product integration.

CA CEM Term	CA Service Desk Term Mapping
incident	A CA CEM *incident* triggers either a service request or an incident, based on the CA Service Desk configuration, and the CA CEM configuration (Setup > Plug-ins).

CA Technologies Prerelease Documentation. Authorized Use Only.

Appendix A: Properties Files Reference

The following are the properties files located in these directories allow you to control the Enterprise Manager, Workstation, and WebView settings.

Important! On Windows machines, when adding a path in the properties files you must escape a backslash (\) with another backslash (each one doubled)—for example: *C:\\Introscope\\lib\\Agent.jar*.

This section contains the following topics:

IntroscopeEnterpriseManager.properties (see page 403)
APMEnterpriseManager.properties (see page 443)
apm-events-thresholds.xml (see page 447)
ResourceMetricMap.properties (see page 461)
Catalyst.properties (see page 464)
EMservice.conf (see page 472)
Introscope Enterprise Manager.lax (see page 477)
IntroscopeWorkstation.properties (see page 483)
Introscope Workstation.lax (see page 487)
IntroscopeWebView.properties (see page 492)
Introscope WebView.lax (see page 500)

IntroscopeEnterpriseManager.properties

A variety of Enterprise Manager behaviors are controlled by properties defined in one of three Enterprise Manager properties files, *IntroscopeEnterpriseManager.properties*. (The others are APMEnterpriseManager.properties (see page 443) and apm-events-thresholds.xml (see page 447).)

After configuring most Enterprise Manager properties, you must restart the Enterprise Manager before any changes you make will take effect.

Some Enterprise Manager properties, however, use the Introscope hot configuration service. It is not necessary to restart the Enterprise Manager after making changes to these so-called "hot config" properties. For more information, see Hot Configuration Service (see page 410).

CA Technologies Proprietary and Confidential Documentation
Authorized Use Only.

You can configure these properties in this file:

Available Processors (see page 404)
Alert Behaviors (see page 405)
Communications (see page 405)
Hot Configuration Service (see page 410)
Application Heuristics (see page 410)
Logging (see page 411)
Persistent collections—recording data to a database (see page 412)
Flat File Properties (see page 413)
Enterprise Manager Metric Clamping (see page 414)
SmartStor (see page 415)
Transaction Event Database (see page 419)
Transaction Trace Sampling (see page 420)
Baseline Data Storage (see page 421)
SNMP Collections (see page 421)
WebView (see page 422)
Management Module Bootstrapping (see page 422)
Management Module Deployment (see page 423)
Clustering Settings (see page 423)
Agent interactions (see page 427)
Enterprise Manager hot failover (see page 431)
Configuration directory (see page 432)
Application triage map data (see page 432)
Location Map data (see page 437)
Web Services (see page 437)
Outgoing message queue (see page 438)
Thread dumps (see page 439)
SNMP Triage Map Alert Actions (see page 441)
Application Map Objects (see page 443)

CA Technologies Confidential Documentation. Authorized Use Only.

Available Processors

The Available Processors section of *IntroscopeEnterpriseManager.properties* contains the following property.

introscope.enterprisemanager.availableprocessors

Use this property to set the number of processors available to this Enterprise Manager. It should be used if there is more than one Enterprise Manager on a host.

Default

<blank>

Alert Behaviors

The Alert State Metrics section of *IntroscopeEnterpriseManager.properties* contains the following property:

introscope.enterprisemanager.alertstatemetric.prefix

Name of the Investigator node under which Alert state metrics are shown.

Default

Alerts.

Communications

The Port Settings section of *IntroscopeEnterpriseManager.properties* contains the following properties:

introscope.enterprisemanager.enabled.channels

Specifies the communication channels that the Enterprise Manager can use. To enable a communication channel, add the channel type to the comma separated list of enabled channels. To disable a communication channel, remove it from the list.

The enabled.channels property works in conjunction with the introscope.enterprisemanager.port.<ChannelName> (see page 406) property.

Example

If you have two communications channels called channel1 and channel2, you would set these properties:

```
introscope.enterprisemanager.enabled.channels=channel1,channel2
introscope.enterprisemanager.port.channel1=5001
introscope.enterprisemanager.port.channel2=5002
```

Default

channel1

CA Technologies Pre-Release Documentation. Authorized Use Only.

introscope.enterprisemanager.port.<ChannelName>

This listener port that the Enterprise Manager uses for encrypted communication with Workstations and agents.

To create a custom communication channel, add these properties to the properties file:

- introscope.enterprisemanager.port.ChannelName
- introscope.enterprisemanager.serversocketfactory.ChannelName
- where *ChannelName* is the name of the new channel.

To enable the custom channel, add its name to the introscope.enterprisemanager.enabled.channels (see page 405) property.

Default

5001

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.serversocketfactory.<ChannelName>

This property identifies the classname of the Server Socket Factory that the Enterprise Manager uses for encrypted communication with Workstations and agents.

If you create a custom communication channel, add a new property to the properties file, for instance:

introscope.enterprisemanager.serversocketfactory.ChannelName

where ChannelName is the name of the new channel, and set the property value to the fully qualified classname of the custom Server Socket Factory that implements the *com.wily.isengard.postofficehub.link.net.IServerSocketFactory* interface.

You must also specify the listener port for the channel, and enable it, as described in introscope.enterprisemanager.port.<ChannelName> (see page 406).

Example

If you have two communications channels called channel1 and channel2 set like this:

```
introscope.enterprisemanager.enabled.channels=channel1,channel2
introscope.enterprisemanager.port.channel1=5001
introscope.enterprisemanager.port.channel2=5002
```

you would set the serversocketfactory property like this:

```
introscope.enterprisemanager.serversocketfactory.channel1=<classname>
introscope.enterprisemanager.serversocketfactory.channel2=<classname>
```

Default

com.wily.isengard.postofficehub.link.net.server.DefaultServerSocketFactory

introscope.enterprisemanager.serversockets.reuseaddr

Set this property to true to allow a communication channel server socket to rebind to a local port that is stuck in a *FIN_WAIT* or *TIME_WAIT* state. Once the Enterprise Manager starts up again, this property should be disabled.

Do not enable this property by default, since there may be a valid reason that the port is already in use.

Default

False

CA Technologies Proprietary/Confidential Documentation. Authorized Use Only.

introscope.enterprisemanager.disableInteractiveMode

Disables interactive console (whether or not the console accepts input).

Notes

On many UNIX platforms, nohup will not work if the application is trying to read from standard input. To run in nohup mode, set this property to true, and change the lax.stdin.redirect property in the Application triage map data file to <blank>. If both of these properties are not set exactly as stated above, it can either render the application unable to start up in nohup, or cause maximum CPU utilization.

Default

True

introscope.enterprisemanager.webserver.port

The port on which the Enterprise Manager's embedded Web server listens.

Default

8081

introscope.enterprisemanager.webserver.max.threads

The maximum number of listener threads in the Web server—that is, the number of concurrent HTTP clients that can be served by the Enterprise Manager's embedded Web server.

Notes

Configure this property only if more than 80 HTTP tunneling agents are connecting to this Enterprise Manager. In this situation, CA Technologies recommends setting the property value to 20 more than the maximum number of connecting HTTP tunneling agents. If you expect fewer than 80 concurrent HTTP agents, do not modify this property value.

Default

100

CA Technologies Prerelease Documentation
Authorized Use Only

introscope.enterprisemanager.ipaddress

Binds all Enterprise Manager communication channels, including the embedded Web server, to a specific local IP address.

When this property is not configured, the Enterprise Manager accepts connections on all local addresses.

Rules to consider when configuring this property:

- If this property is configured to a IPv4 address, the Enterprise Manager will only accept connections from agents or clients on running on an IPv4 system.
- If the property is configured to an IPv6 address, the Enterprise Manager will only accept connections from agents or clients on machines running an IPv6 compliant JVM. Agents or clients on machines running an IPv4 JVM cannot use an IPv6 address to connect.
- If you do not bind the Enterprise Manager or WebView instance running on IPV6 enabled system to a specific IP address, the instance is bound to all local addresses. It can accept connections from both IPv4 and IPv6 environments.
- Workstations, agents, Browsers, and other clients should use the DNS name, not the IP address, to connect to the Enterprise Manager or WebView instance.
- IPv6 support on the agent is dependent upon IPv6 support in the environment (JVM, operating system, hardware stack) in which the agent is running.

Default

<commented out, no value>

introscope.enterprisemanager.workstation.connection.channel

Specifies the channel for connecting to the Web Start Workstation.

Default

channel1

CA Technologies Proprietary & Confidential Documentation. Authorized Use Only.

Hot Configuration Service

The Enterprise Manager includes a hot configuration that is enabled by default for these properties:

- log4j logging properties—see Logging (see page 411)
- Persistent collections—recording data to a database properties—see Persistent collections—recording data to a database (see page 412)
- Enterprise Manager Metric Clamping properties—see Enterprise Manager Metric Clamping (see page 414)
- clustering properties—see Clustering Settings (see page 423)

introscope.enterprisemanager.hotconfig.enable

Enables the hot configuration service. Set to false to disable the service.

Default

true

introscope.enterprisemanager.hotconfig.pollingInterval

Specifies how often to check this file for changes, in seconds.

Default

60 seconds

Application Heuristics

The Application Heuristics section of *IntroscopeEnterpriseManager.properties* contains the following property:

introscope.enterprisemanager.application.overview.baselines

To improve Enterprise Manager performance, you can turn off heuristics by setting the property to false.

If you do this, the Enterprise Manager turns off the Enterprise Manager's assessment of JVM, application, and backend health. It also turns off the baseliner, which means that metrics in the Workstation will be inactive and report no values. The metrics will appear gray unless the configuration is set back to true, at which point they will be active again.

Default

True

CA Technologies Prerelease Documentation. Authorized Use Only.

Logging

Logging properties define the level of detail in the Enterprise Manager logs and the name and location of Enterprise Manager log files. See Log4J documentation at http://jakarta.apache.org/log4j/docs/documentation.html for information on advanced Log4J functions.

These Enterprise Manager properties use the Introscope hot configuration service. It is not necessary to restart the Enterprise Manager after making changes to these hot configuration properties. For more information, see Hot Configuration Service *(see page 410).*

log4j.logger.Manager

Controls the level of detail, and destination for log messages.

Level of detail value can be: *INFO* or *VERBOSE#com.wily.util.feedback.Log4JSeverityLevel*.

Destination value can be console or logfile. The default setting writes log information to the console and the log file.

Default

INFO, console, logfile

log4j.appender.logfile.File

Name and location of Enterprise Manager logfile, relative to the Enterprise Manager home directory.

Default

logs/IntroscopeEnterpriseManager.log

CA Technologies Prereleased Documentation
Authorized Use Only.

log4j.appender.logfile.layout.ConversionPattern

Allows you to customize your log files.

The property can be set the following values:
d{M/dd/yy hh:mm:ss a z} [%-3p] [%c] %m%n

The pattern is implemented in the log4j library and the format is described here:
http://logging.apache.org/log4j/docs/api/org/apache/log4j/PatternLayout.html

The date/time format appears in the *%d{..}* and follows a pattern described here:
http://java.sun.com/j2se/1.4.2/docs/api/java/text/SimpleDateFormat.html

Default

<blank>

Persistent collections—recording data to a database

You can configure the properties described in this section to record traced data to an external, relational database. For more information, see the SQL Database Settings section of the *IntroscopeEnterpriseManager.properties* file. These settings are deployed by hot configuration, so they do not require a manual restart of the Enterprise Manager.

Introscope includes an API that enables you to extract metric and event data for historical ranges, then use a data analysis tool such as Excel or SAS to view Introscope data. CA Technologies recommends that you use this API record data—see JDBC API (see page 507). You can also use the Command-Line Workstation commands for extracting persisted data—see Using the Command-Line Workstation (see page 141).

These Enterprise Manager properties use the Introscope hot configuration service. You do not need to restart the Enterprise Manager after changing these so-called "hot config" properties. For more information, see Hot Configuration Service (see page 410).

introscope.enterprisemanager.database.collection.agentExpression

This is the agent regular expression property. Add this setting to the *IntroscopeEnterpriseManager.properties* file and is deployed via hot configuration.

For example:

Introscope.enterprisemanager.database.collection.agentExpression=rhart-dt1\\|EPAgentProcess1\\|EPAgent5

Default

<blank>

CA Technologies Proprietary Documentation
Authorized Use Only.

introscope.enterprisemanager.database.collection.metricExpression

This is the metric regular expression property. Add this setting to the *IntroscopeEnterpriseManager.properties* file and is deployed via hot configuration.

Example

introscope.enterprisemanager.database.collection.metricExpression=Servlets\\|MyServlet:Average Response Time \\(ms\\)

Default

<blank>

introscope.enterprisemanager.database.collection.frequencyinseconds

Where agentExpression is the agent regular expression, metricExpression is the metric regular expression, and frequencyinseconds is the frequency in seconds to aggregate the data and write it to the database. Add this setting to the *IntroscopeEnterpriseManager.properties* file and is deployed via hot configuration.

There is no default setting as each persistent collection can have a different value. This property should be set for each persistent collection.

Default

none

introscope.enterprisemanager.database.recordingQueue.upperLimit.value

The maximum number of metric data points on the queue. A metric data point is one time slice of data for one metric. This setting is deployed via hot configuration.

Default

150k (thousand)

Flat File Properties

Flat files are plain-text comma-separated value (*.csv*) files. Introscope can write to flat files, but cannot read from them. You can analyze flat files with numerical analysis tools.

Flatfile recording is enabled by default and directed to the *<EM_Home>/flatfiles* directory. You can also specify that these files reside in a directory outside of the Introscope directory. Note that for Windows paths, you must escape backslashes (\) with another backslash (doubled), for example: *C:\\Introscope\\FlatFileDir*.

CA Technologies Prerelease Documentation. Authorized Use Only

Enterprise Manager Metric Clamping

These properties define the limits set for the number of results you get for metrics from a standalone Enterprise Manager or Collector. They are hot configurable meaning the values are applied when the Enterprise Manager checks the property file during a hot deploy. A hot deploy occurs about once every 60 seconds.

To disable clamps:

1. Open the IntroscopeEnterpriseManager.properties file for editing and navigate to the EM Metric Clamping section.
2. To disable a clamp, set the value of the clamp to 0.
3. Save and close IntroscopeEnterpriseManager.properties.

There are additional Enterprise Manager clamps in the apm-events-thresholds.xml file in the <*EM_Home*>/config directory. See apm-events-thresholds.xml (see page 447).

introscope.enterprisemanager.query.datapointlimit.

This property defines the maximum number of metric data points an Enterprise Manager retrieves for any single CLW or JDBC query. The effect of this clamp is to reduce the number of metric a clustered Collector or standalone Enterprise Manager retrieves from SmartStor. However, the Enterprise Manager always returns the full data for at least one metric. The clamp is for each CLW or JDBC query, and not for multiple concurrent queries.

Example

If:

- The limit is set to 10000.
- A CA APM user makes a CLW or JDBC query that returns 11000 data points.
- Another user makes a query that returns 8000 data points.

Then, the Enterprise Manager limits the query for the first user to 10000 data points, and the second user can continue performing their query.

Notes

You cannot set this property to clamp a query on a MOM.

Default

0 (where 0 means that there is no limit.)

CA Technologies Pre-release Documentation. Authorized use only.

introscope.enterprisemanager.query.returneddatapointlimit

This property limits the number of data points that SmartStor returns for a CLW or JDBC query. The effect of the clamp is to reduce the number of metrics for which data is returned to a clustered Collector or standalone Enterprise Manager. The clamp is for each CLW or JDBC query, and not for multiple concurrent queries.

Example

If:

- The limit is set to 10000.
- A CA APM user makes a CLW or JDBC query that returns 11000 data points.
- Another user makes a query that returns 8000 data points.

Then, the Enterprise Manager limits the query for the first user to 10000 data points, and the second user can continue performing their query.

Notes

You cannot set this property to clamp a query on a MOM.

Default

0 (where 0 means that there is no limit.)

SmartStor

The properties in this section define the directory where SmartStor data is stored, and configure data collection and aging behavior for three tiers of storage. It also allows you to disable SmartStor if desired.

Note: All three tiers must be defined. You cannot define more than three tiers.

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.smartstor.enable

Allows you to enable or disable SmartStor. While disabling SmartStor will save disk space, it will not store statistics and you will be prevented from monitoring them.

Set to False to disable SmartStor.

Notes

If you do not see the property in the SmartStor section, then SmartStor is enabled. To disable SmartStor add the following line to the properties file:
introscope.enterprisemanager.smartstor.enable=false

Default

True

introscope.enterprisemanager.smartstor.directory

Specifies the directory, relative to the Enterprise Manager home directory, that SmartStor will use.

Default

data

introscope.enterprisemanager.smartstor.directory.archive

Specifies the directory that SmartStor will use to store archived data.

Default

data/archive.

introscope.enterprisemanager.smartstor.dedicatedcontroller

Specifies whether SmartStor data is written to a dedicated disk.

Set to true if SmartStor data is being written to a dedicated disk.

For more information about using this property, see the CA APM S*izing and Performance Guide*.

Default

false

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.smartstor.tier1.frequency

Specifies the frequency (in seconds) of data collection for the first data tier. The frequency:

- must be a multiple of 15 seconds.
- must not be greater than 1800 seconds (30 minutes).
- must be an even multiple of the previous tier's frequency—this multiple can be one; each tier can have the same frequency).

Default

15

introscope.enterprisemanager.smartstor.tier1.age

Specifies the number of days after which data in first tier ages out of SmartStor. The value cannot be zero.

Default

7

introscope.enterprisemanager.smartstor.tier2.frequency

Specifies the frequency (in seconds) of data collection for the second data tier. The frequency:

- must be a multiple of 15 seconds.
- must not be greater than 1800 seconds (30 minutes).
- must be an even multiple of the previous tier's frequency—this multiple can be one; each tier can have the same frequency).

Default

60

introscope.enterprisemanager.smartstor.tier2.age

Specifies the number of days after which data in second tier ages out of SmartStor. The value cannot be zero.

Default

23

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.smartstor.tier3.frequency

Specifies the frequency (in seconds) of data collection for the third data tier. The frequency:

- must be a multiple of 15 seconds.
- must not be greater than 1800 seconds (30 minutes).
- must be an even multiple of the previous tier's frequency—this multiple can be one; each tier can have the same frequency).

Default

900

introscope.enterprisemanager.smartstor.tier3.age

Specifies the number of days after which data in third tier ages out of SmartStor. The value cannot be zero.

Default

335

introscope.enterprisemanager.smartstor.reperiodizationOffsetHour

This property allows you to configure the time offset for data tiering.

Constraints

Must be an integer from 0 to 23, representing hours of the day, where 0 signifies midnight and 13 signifies 1:00 p.m.

Default

0 (0 means that data tiering starts at midnight).

Example

If you set this property to 2, the data tiering starts at 2:00 a.m.

introscope.enterprisemanager.smartstor.conversionOffsetMinute

This property allows you to configure a time offset for spool converting to data.

Constraints

Must be an integer from 0 to 59, representing the number of minutes after the hour.

CA Technologies Prerelease Documentation. Authorized Use Only.

Default

0 (for example, setting the value to xx:00 means that conversion starts on the hour).

Example

If you set the value to 10 (xx:10), the spool conversion starts at 10 minutes after the hour.

Transaction Event Database

The Enterprise Manager stores captured transaction events locally on its file system for a fixed length of time or until a maximum disk usage is exceeded. You can control where to store these events and when they should be deleted by setting the transaction event database properties.

introscope.enterprisemanager.transactionevents.storage.dir

Specifies the location, relative to the Enterprise Manager home directory, where Introscope stores transaction event data.

Default

traces

introscope.enterprisemanager.transactionevents.storage.max.data.age

Specifies, in days, how long to store transaction event data.

Default

14

introscope.enterprisemanager.transactionevents.storage.optimize.frequency

Specifies how often, in hours, the trace storage should be optimized by cleaning up aged data.

Enables aging of transaction trace data when the Enterprise Manager starts.

Default

24 hours

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.transactionevents.storage.optimize.timeoffsethour

Specifies the time offset for storage optimization, which allows you to schedule purges to start on a specific hour of every day against the time offset.

For example, if the time offset is set to 02 and optimize frequency (see introscope.enterprisemanager.transactionevents.storage.optimize.frequency (see page 419)) is set to 24 hours, optimization will run every 24 hours offset from time 02:00AM.

Entering a 00 value for this property will display an error message, as this value must be a positive integer.

Default

02 (2AM)

Transaction Trace Sampling

These transaction trace sampling properties are located, by default, in *IntroscopeEnterpriseManager.properties*.

introscope.agent.transactiontracer.sampling.perinterval.count

Configured for 1 sample trace by default. This property works in conjunction with the introscope.agent.transactiontracer.sampling.interval.seconds property.

You can change the number of traces per interval to trace, by modifying its value and uncommenting it in the properties file.

Default

1

CA Technologies Internal/Restricted Documentation. Authorized Use Only.

introscope.agent.transactiontracer.sampling.interval.seconds

Configured to trace every 2 minutes by default. This property works in conjunction with the introscope.agent.transactiontracer.sampling.perinterval.count property.

You can change the sampling interval by modifying its value and uncommenting it in the properties file.

Notes

Before overriding the defaults for the sampling.perinterval and sampling.interval properties, consider the potential for increased load in the Enterprise Manager with higher sampling rates.

Default

120 seconds

Baseline Data Storage

The EM Object Database Settings section of *IntroscopeEnterpriseManager.properties* contains the following property:

introscope.enterprisemanager.dbfile

Specifies the filename, and its location relative to the Enterprise Manager home directory, of the database in which Introscope stores metric baseline information.

Default

data/baselines.db

SNMP Collections

The SNMP Adapter section of *IntroscopeEnterpriseManager.properties* contains the following properties:

introscope.enterprisemanager.snmp.enable

Enables SNMP agent (only relevant if you have installed SNMPServices.jar)

Default

false

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.snmp.agent.port

Enterprise Manager port that communicates with the SNMP agent

Default

161

WebView

The WebView Configuration section of *IntroscopeEnterpriseManager.properties* contains the following properties:

introscope.enterprisemanager.webview.disableLogin

Determines whether WebView can establish a connection with Enterprise Manager.

Default

false

introscope.enterprisemanager.webview.password

Specifies password WebView users to connect to Enterprise Manager.

Default

WebView

Management Module Bootstrapping

The Management Module Bootstrapping section of *IntroscopeEnterpriseManager.properties* contains the following property:

introscope.enterprisemanager.managementModules.invalid.references.as.errors

Logs invalid references as warnings. If set to true, invalid references will prevent the Enterprise Managerfrom starting.

Default

false

CA Technologies Prerelease Documentation. Authorized Use Only.

Management Module Deployment

The Management Module Deployment section of *IntroscopeEnterpriseManager.properties* contains the following properties:

introscope.enterprisemanager.hotdeploy.directory

Defines the directory, relative to the Enterprise Manager home directory, to check for updated applications to deploy.

Default

deploy

introscope.enterprisemanager.hotdeploy.checkFrequencyInSeconds

Specifies how often, in seconds, to check for updated Management Modules to deploy.

Default

60

Clustering Settings

For more information about Enterprise Manager clustering, see Working with Enterprise Manager Clustering (see page 101).

These Enterprise Manager properties use the Introscope hot configuration service. It is not necessary to restart the Enterprise Manager after making changes to these so-called "hot config" properties. For more information, see Hot Configuration Service (see page 410).

introscope.enterprisemanager.clustering.manager.enable

Setting this property to true identifies the Enterprise Manager as the Manager of Managers (MOM) in a cluster of Enterprise Managers.

Default

false

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.clustering.manager.slowcollectorthreshold

Set this property to define the threshold for the ping metric for your environment. The ping metric monitors only the lower boundary of the round-trip response time from the MOM to each Collector.

Default

10000 (in milliseconds)

introscope.enterprisemanager.clustering.manager.slowcollectordisconnectthresholdseconds

Set this property to define the threshold for which the MOM will automatically disconnect from the Collector if exceeded. This will prevent the entire cluster from hanging, which is a side effect of when one Collector in a cluster is greatly under performing.

If this property is set to zero or a negative value, which is not recommended, this will not be enabled.

Notes

The MOM checks Collectors every 7.5 seconds, so small changes will not affect its behavior.

Default

60 (in seconds)

introscope.enterprisemanager.clustering.login.<emName>.host

Hostname or IP address of the Collector Enterprise Manager.

Default

<blank>

introscope.enterprisemanager.clustering.login.<emName>.port

Value of the Collector Enterprise Manager's introscope.enterprisemanager.port.<ChannelName> (see page 406) property.

Default

<blank>

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.clustering.login.<emName>.publickey

Specifies the public key for encrypting the password the MOM uses to connect to a Collector Enterprise Manager.

Default

config/internal/server/EM.public, commented out

introscope.enterprisemanager.clustering.privatekey

Specifies the private key a Collector Enterprise Manager uses to decrypt the password the MOM uses to connect to it.

Default

config/internal/server/EM.private

introscope.enterprisemanager.clustering.collector.enable

Set to true to identify an Enterprise Manager as a Collector, for clustering. If an Enterprise Manager is configured for both the Manager of Managers role and the Collector role, the Enterprise Manager becomes the Collector.

If the Enterprise Manager is designated a Collector during installation, the Enterprise Manager installer configures this property. If the Enterprise Manager is designated a Manager of Managers (MOM), the Enterprise Manager installer can configure up to five Collectors. If you want a MOM to track six or more Collectors, set this property manually.

You can also use this property if you configured the Enterprise Manager standalone during installation, then decide later it will be part of a cluster.

Default

false

introscope.enterprisemanager.clustering.collector.identifier

Optional. You can set this property if you want to give this Collector a descriptive name.

If you don’t specify an identifier, Introscope uses the default identifier of host@port, where host is the Enterprise Manager's host name, and port is the Enterprise Manager's connection port.

Default

false

CA Technologies Pre-release Documentation. Author Use Only.

introscope.enterprisemanager.clustering.manager.staggeredstartdelay

Set this optional property on the MOM to alter the length of time between MOM connection attempts to clustered Collectors. If this property is not set, the default time between connection attempts is 0 seconds, meaning that the MOM recognizes each Collector as soon as the Collector becomes available.

Setting this property helps even out the MOM's startup overhead. It gives the MOM a period of time to process the connection as each Collector comes online before starting the connection process with another waiting Collector.

The minimum value is 0 seconds.

Default

0 (seconds). The minimum value is 0 seconds.

introscope.enterprisemanager.loadbalancing.threshold

Set this property to define the tolerance of the Load Balancer for imbalance. The number of metrics in a Collector must differ from the weight-adjusted cluster average (i.e. the average for all Collectors) by at least the threshold before the Load Balancer rebalances the cluster.

Rebalancing means redistributing agents among Collectors to equalize the weight-adjusted metric load.

Default

20000 (metrics)

introscope.enterprisemanager.loadbalancing.interval

Set this property to determine how often the Load Balancer should rebalance in seconds. The minimum value is 120 seconds.

Default

600 (10 minutes)

CA Technologies Pre-release Documentation
Authorized Use Only.

introscope.enterprisemanager.clustering.login.em1.weight

Set this optional property to control the relative load of the Collector.

All other things being equal, the weight divided by the total weight for the cluster is the percentage of metric load assigned to that Collector. For example, if the weight of a Collector is 50 and the sum of the weights of all Collectors is 250, the Collector should get 20% of the cluster's metric load.

Average value for the cluster, with the weight being the load balancing weight for this Enterprise Manager Collector.

Agent interactions

Some agent interaction properties require Enterprise Manager restart after configuration. Others are automatically deployed using the Enterprise Manager hot deploy configuration service, and hot deploy is noted in the property description.

Some agent interaction properties are not included in the Enterprise Manager properties file. To change their default values, add them to the *IntroscopeEnterpriseManager.properties* file and save the file. If a property is not hot deployed, shut down and restart the Enterprise Manager to enable the new values.

introscope.apm.agentcontrol.agent.allowed

Determines, when an agent does not match any regular expression in loadbalancing.xml, whether the agent is allowed to connect to any of the following Enterprise Managers:

- Specific standalone Enterprise Manager.
- Collectors.

Default

true

Example

introscope.apm.agentcontrol.agent.allowed=true

Notes

This setting is deployed using hot configuration.

If this property is:

- Missing from the IntroscopeEnterpriseManager.properties file, then the value is true.
- Set to true on the MOM, then all agents not matched in loadbalancing.xml are allowed in the cluster.

CA Technologies Proprietary & Documentation. Authorized Use Only.

- Set to false, then all agents not matched in loadbalancing.xml are disallowed and:
  - Pre-9.1 agents do not connect to any Enterprise Manager
  - 9.1 agents connect to the Enterprise Manager in disallowed mode.
  - Agent is added to the APM Status console Denied Agents list.

If a pre-9.1 agent is:

- Connected to an Enterprise Manager.
- Disallowed based on a change of this property setting.

Then the Enterprise Manager disconnects the agent. The Enterprise Manager denies all further agent connection attempts until the agent is allowed either by changing this property or configuring loadbalancing.xml. The agent waits to reconnect based on the time configured in the introscope.enterprisemanager.agent.disallowed.connection.limit (see page 430) property.

introscope.apm.agentcontrol.agent.emlistlookup.enable

Determines whether the following can send to each agent the list of Enterprise Managers to which the agent is allowed to connect:

- MOM.
- Collectors.
- Standalone Enterprise Manager.

When any of the following goes down with the agent connected, this list tells the agent to which Enterprise Manager to connect next:

- MOM.
- Standalone Enterprise Manager.
- Initial Collector to which the agent connects.

Default

true

Example

introscope.apm.agentcontrol.agent.emlistlookup.enable=true

CA Technologies Pre-Release Documentation. Authorized Use Only.

Notes

This setting is deployed using hot configuration.

If an agent restarts, it uses its agent profile introscope.agent.enterprisemanager.connectionorder property for its initial Collector connection information.

introscope.apm.agentcontrol.agent.reconnect.wait

Provides the base duration, in seconds, that an agent waits before trying to reconnect to a standalone Enterprise Manager or Collector.

Default

45 (seconds)

Example

introscope.apm.agentcontrol.agent.reconnect.wait=45

Notes

This setting is deployed using hot configuration.

This property applies to pre-9.1 agents and 9.1 and later agents.

To avoid many simultaneous agent connections, the Enterprise Manager adds to the default value a random number of seconds from this set of values: -30, -15, 0, 15, 30.

If the minimum calculated time is 15 seconds or less, the Enterprise Manager waits 45 seconds before allowing the agent to reconnect.

introscope.enterprisemanager.agentconnection.metrics.enable

Controls whether or not the Enterprise Manager generates agent connection metrics. You can set this to True or False.

This property is not included in the Enterprise Manager properties file. To change the value, add the property to the properties file and restart the Enterprise Manager to enable the property.

Default

True

CA Technologies Internal / Customer Use Only. Unauthorized Use Prohibited. Documentation.

introscope.enterprisemanager.agentconnection.metrics.agentTimeoutInMinutes

Specifies how many minutes after an agent disconnect the Enterprise Manager stops generated connection metrics for the agent.

This property is not included in the Enterprise Manager properties file. To change the value, add the property to the properties file and restart the Enterprise Manager to enable the property.

Default

30 (minutes)

introscope.enterprisemanager.agent.disallowed.connection.limit

Limits the number of disallowed agents that can connect to an Enterprise Manager.

This clamp property applies to the MOM, standalone Enterprise Managers, and Collectors.

When this clamp is hit and a disallowed agent tries to connect, the Enterprise Manager does not accept the new connection. In addition, the Enterprise Manager logs this activity and the agent is added to the APM Status console Denied Agents list.

Default

0 (unlimited)

Example

introscope.enterprisemanager.agent.disallowed.connection.limit=0

Notes

This setting is deployed using hot configuration.

This property applies to 9.1 and later agents.

If this property is set on the MOM, the MOM value does not override the Collector value.

introscope.enterprisemanager.autoUnmountDelayInMinutes

Number of minutes the Enterprise Manager continues to track an agent after the agent is disconnected, in case the disconnection is temporary. After this time runs out, the Enterprise Manager releases agent tracking resources.

Default

60 (minutes)

CA Technologies Prerelease Documentation. Authorized Use Only.

Enterprise Manager hot failover

These properties allow you to enable and apply hot failover settings, which can be used for enterprise managers sharing the same installation directories. When setting these properties, you might want to consider the Enterprise Manager hot failover rules.

introscope.enterprisemanager.failover.enable

This property enables hot failover.

Default

false

introscope.enterprisemanager.failover.primary

This property allows you to set the hostname or IP address of the primary Enterprise Manager. To configure both Enterprise Managers as primary, set to a comma-separated list of hostnames or IP addresses of the primary Enterprise Managers. The following is an example of both Enterprise Managers configured as primary: *introscope.enterprisemanager.failover.primary=myhost1, myhost2*.

If you configure Enterprise Manager hot failover (see Using Enterprise Manager hot failover (see page 104)) for a single host rather than a shared filesystem, the primary Enterprise Manager needs to be set to itself, and the *introscope.enterprisemanager.failover.secondary* property needs to blank.

Default

<blank>

introscope.enterprisemanager.failover.secondary

This property allows you to set the hostname or IP address of the secondary Enterprise Manager by specifying the hostname or IP address of the Enterprise Manager's host. For example, *introscope.enterprisemanager.failover.secondary=myhost2*.

Default

<blank>

CA Technologies Proprietary/Confidential Information. Authorized Use Only.

introscope.enterprisemanager.failover.interval

This property allows you to set in seconds how often a running secondary Enterprise Manager checks to see if a primary Enterprise Manager is waiting to start. This setting is deployed via hot configuration.

Default

120 (2 minutes)

Configuration directory

This property, in the EM Config Directory section, allows you to specify the location of the Enterprise Manager *config* directory.

introscope.enterprisemanager.directory.config

The directory the Enterprise Manager uses for configuration. In Windows, backslashes must be escaped (doubled). For example: *C:\\Introscope\\config*

Default

The config directory in the original Introscope installation.

Application triage map data

These properties allow you to enable and apply application triage map data settings. For configuration instructions, see Configuring the Application Triage Map (see page 219).

Additional properties pertaining to the application triage map are configured in the *APMEnterpriseManager.properties* file. See APMEnterpriseManager.properties (see page 443).

introscope.apm.data.agingTime

This property, which is loaded by agents, controls how often agents report application triage map data to the Enterprise Manager.

Notes

This property is not hot configurable. If you change this property, restart the agent and Enterprise Manager in order for the changes to take effect.

CA Technologies Prerelease Documentation. Authorized Use Only.

WARNING! Before changing this property, delete the map data from the APM database. When this value is changed and old map data is used with new map data, historical queries will show inconsistent results with respect to the new value. For information on deleting map data from the APM database, see Application triage map data pruning (see page 224).

This value cannot be less than 1 minute. If you set it to less than 1 minute, the property defaults to 1 day.

When deciding on this setting, consider how often your organization adds and modifies applications. If you set a lower value, then the map displays application dependencies on a more real-time basis, but this can use more agent and Enterprise Manager overhead. If you set a higher value, then the map displays application dependencies that can be less current, but this can use less agent and Enterprise Manager overhead.

For information about agent and Enterprise Manager overhead, see the CA APM *Sizing and Performance Guide*.

Use any of these time units, which are not case sensitive. For:

- Days use: days, D
- Minutes use: mins, M
- Hours use: hours, H

This agingTime property works in conjunction with the *introscope.apm.data.timeWindow* property, such that agingTime is usually a subset of the timeWindow.

Example

If the *introscope.apm.data.agingTime* value is 3 days, and the *introscope.apm.data.timeWindow* value is 10 days, the application triage map displays application dependencies that have been reported by agents during the previous 10 days. If no data has been reported for a dependency during those 10 days, the dependencies will appear as aged after the agingTime period has expired, or 3 days, and will remain as aged elements on the map until the end of the 10 day timeWindow.

Default

1 day

CA Technologies Prerelease Documentation. Authorized Use Only

introscope.apm.data.timeWindow

This property defines a time period, in addition to the period defined in the *introscope.apm.data.agingTime* property.

Notes

This property is not hot configurable. If you change this property, restart the Enterprise Manager in order for the changes to take effect.

This value cannot be less than 1 minute. If you set it to less than 1 minute, the property defaults to 1 day.

Use any of these time units, which are not case sensitive. For:

- Days use: days, D
- Minutes use: mins, M
- Hours use: hours, H
- Seconds use: secs, S
- Milliseconds use: MS

This property works in conjunction with the *introscope.apm.data.timeWindow* property, such that agingTime is usually a subset of the timeWindow.

Example

If the *introscope.apm.data.agingTime* value is 3 days, and the *introscope.apm.data.timeWindow* value is 10 days, the application triage mapdisplays application dependencies that have been reported by agents during the previous 10 days. If no data has been reported for a dependency during those 10 days, the dependencies will appear as aged after the agingTime period has expired, or 3 days, and will remain as aged elements on the map until the end of the 10 day timeWindow.

When deciding on this setting, consider how often your organization adds and modifies applications. If you set a lower value, then the map displays application dependencies on a more real-time basis, but this can use more agent and Enterprise Manager overhead. If you set a higher value, then the map displays application dependencies that can be less current, but this can use less agent and Enterprise Manager overhead.

For information about agent and Enterprise Manager overhead, see the *CA APM Sizing and Performance Guide*.

Default

3 days

CA Technologies Pre-Release Documentation. Authorized Use Only.

introscope.apm.pruning.enabled

This property enables the pruning of application triage map data from the APM database.

Notes

- This property is not hot configurable. If you change this property, restart the Enterprise Manager in order for the changes to take effect.
- When this property is set to true, application triage map data is pruned from the APM database.
- When this property is set to false, application triage map data is not pruned from the APM database, and is stored indefinitely.

Default

Application triage map data pruning is enabled by default to be done on the APM database in a cluster environment for standalone Enterprise Managers and MOMs.

Important! You cannot configure a Collector for application triage data pruning. If you set the introscope.apm.pruning.enabled property to true for a Collector, errors are returned.

introscope.apm.data.preserving.time

This property controls how long application triage map data is stored in the APM database. Map data older than the preserving time value is considered aged out, and can be pruned.

Notes

This property is not hot configurable. If you change this property, restart the Enterprise Manager in order for the changes to take effect.

During pruning, application triage map data that is older than the sum of the *introscope.apm.data.preserving.time* and *Introscope.apm.data.agingTime* values is deleted.

For example, if the *introscope.apm.data.preserving.time value* is 30 days and the *Introscope.apm.data.agingTime* value is 1 day, map data older than 30 + 1 = 31 days is deleted.

CA Technologies Prerelease Documentation. Authorized Use Only

WARNING! If the preservation time value is set too low (seconds or milliseconds), application triage map data might be deleted from the APM database right after the data arrives from the Enterprise Manager.

Use any of these time units, which are not case sensitive. For:

- Days use: days, D
- Minutes use: mins, M
- Hours use: hours, H
- Seconds use: secs, S
- Milliseconds use: MS

Default

365 days

introscope.apm.pruning.cron.trigger.expression

This property schedules how often aged-out application triage map data is pruned from the APM database.

Notes

This property is not hot configurable. If you change this property, restart the Enterprise Manager in order for the changes to take effect.

The *introscope.apm.pruning.cron.trigger.expression* property is formatted as an Open Symphony Quartz CronTrigger scheduler expression. For information about Quartz ChronTrigger, see www.quartz-scheduler.org (www.quartz-scheduler.org).

This expression is a string of 6 or 7 fields separated by white spaces. The fields in string order from left to right are:

Seconds, Minutes, Hours, Day of month, Month, Day of week, Year

Fields can contain:

- any of the allowed values
- various combinations of the allowed special characters for that field.

For the allowed values, special characters, and example expressions, see www.quartz-scheduler.org (www.quartz-scheduler.org).

Default

*0 0 6 * * ?* — This default expression prunes map data daily at 6 AM.

CA Technologies Prerelease Documentation. Authorized Use Only.

Location Map data

These properties allow you to configure settings on the Location Map.

introscope.workstation.locationmap.clamp

The property controls the maximum number of elements the location map can contain. The effect of the clamp is to control the size of the location map.

Notes

When the number of elements returned exceeds the clamp's value, some map elements are not displayed. Users will see an error message informing that some elements are not displayed.

Default

500

Web Services

Users who access the web services APIs on a regular basis may notice a performance degradation due to a high number of web services calls. The number of threads devoted to web services is set, by default, to 10. You can change this setting by adding a property to the IntroscopeEnterpriseManager.properties file.

introscope.apm.model.webservices.threadpool.numberofthreads

This property specifies the number of threads used in a thread pool. Use of this property is optional.

WARNING! Setting this property to a large number may cause performance degradation. Use caution in setting this property.

Constraint

Must be an integer.

Default value

10

Example

```
introscope.apm.model.webservices.threadpool.numberofthreads=15
```

This property is also documented in the *CA APM Catalyst Connector Guide*.

CA Technologies Pre-release Documentation. Authorized Use Only.

Outgoing message queue

These properties allow you to configure outgoing message queue settings.

transport.outgoingMessageQueueSize

This property specifies the number of messages in the outgoing message queue.

Note

Increasing the outgoing message queue size uses more RAM.

Constraint

Must be an integer.

Default value

3000

Example

```
transport.outgoingMessageQueueSize=3000
```

This property is also documented in the *CA APM Sizing and Performance Guide*.

transport.override.isengard.high.concurrency.pool.max.size

This property specifies the maximum number of threads for simultaneous output socket connections. A higher number does not directly improve response time, but keeps isolated slow connections from greatly slowing the Enterprise Manager.

The default, 5, is suitable for most standalone Enterprise Managers and Collectors. MOMs should have one thread per Collector, plus about one thread for every 10 Workstations.

Because of their role as a network traffic router, MOMs in particular can benefit from having more outgoing threads. Not only because it can allow them to simultaneously service more Workstations, but also because it helps keep the entire cluster from freezing when a single Collector is overloaded.

Important! Do not oversize, as additional thread overhead may cause OutOfMemory exceptions, especially on 32 bit JVMs

Constraint

Must be an integer.

Default value

5

Example

```
transport.override.isengard.high.concurrency.pool.max.size=5
```

CA Technologies Prerelease Documentation. Authorized Use Only.

This property is also documented in the *CA APM Sizing and Performance Guide*.

Thread dumps

These properties allow you to configure thread dumps.

For information about enabling and configuring thread dumps, see the *CA APM Java Agent Implementation Guide*.

You can configure these properties for this feature:

introscope.enterprisemanager.threaddump.enable (see page 439)
introscope.enterprisemanager.threaddump.storage.dir (see page 440)
introscope.enterprisemanager.threaddump.storage.clean.disk.freq.days (see page 440)
introscope.enterprisemanager.threaddump.storage.clean.disk.olderthan.days (see page 440)
introscope.enterprisemanager.threaddump.storage.max.disk.usage (see page 441)

introscope.enterprisemanager.threaddump.enable

This property enables the Thread Dumps tab on the metric browser tree agent node for any agent connected to an Enterprise Manager.

Notes

This property is hot configurable. If you change this property, you do not need to restart the Enterprise Manager for the changes to take effect.

This property works together with the IntroscopeAgentProfile introscope.agent.threaddump.enable property. For more information, see the *CA APM Java Agent Implementation Guide*.

When this property is set to *true* and the IntroscopeAgentProfile introscope.agent.threaddump.enable property is set to *true*, then users can see the Thread Dumps tab.

When this property is set to *false*, then users cannot see the Thread Dumps tab.

Default

true

CA Technologies Pre-release Documentation. Authorized Use Only

introscope.enterprisemanager.threaddump.storage.dir

Specifies the directory, relative to the Enterprise Manager home directory, where thread dumps are stored.

Notes

This property is not hot configurable. If you change this property, restart the Enterprise Manager for the changes to take effect.

<*EM_Home*>/threaddumps is not the directory where users save thread dump text files. Users can save thread dump text files to any location.

Default

threaddumps

introscope.enterprisemanager.threaddump.storage.clean.disk.freq.days

How often, in days, that CA Introscope purges thread dump files.

Notes

This property works with the introscope.enterprisemanager.threaddump.storage.clean.disk.olderthan.days property. Set them both so that the Enterprise Manager knows how long to store and purge thread dump files.

This property is not hot configurable. If you change this property, restart the Enterprise Manager for the changes to take effect.

Default

1 (day)

introscope.enterprisemanager.threaddump.storage.clean.disk.olderthan.days

The age, in days, after which CA Introscope purges thread dump files.

Notes

This property works with the introscope.enterprisemanager.threaddump.storage.clean.disk.freq.days property. Set them both so that the Enterprise Manager knows how long to store and purge thread dump files.

This property is not hot configurable. If you change this property, restart the Enterprise Manager for the changes to take effect.

CA Technologies Pre-release Documentation. Authorized Use Only.

Default

30 (days)

introscope.enterprisemanager.threaddump.storage.max.disk.usage

The number of thread dumps that can be stored on the Enterprise Manager.

Notes

This property is not hot configurable. If you change this property, restart the Enterprise Manager for the changes to take effect.

Constraint

Must be an integer 8000 or less

Default value

5000

SNMP Triage Map Alert Actions

The SNMP triage map alert action plug-ins are installed as part of the Enterprise Manager installer. Set these configurations to specify SNMP triage map alert action properties.

Note: The SNMP triage map alert action plug-in properties are different from the existing SNMP alert action adapter properties. You can distinguish the new SNMP triage map alerts by their prefix, *introscope.apm.catalyst.triagemapalert.snmp*.

introscope.apm.catalyst.triagemapalert.snmp.destination.host.ip

Set this property with the IP address of the APM connector.

introscope.apm.catalyst.triagemapalert.snmp.destination.trap.port

Set this property with the SNMP port where the SNMP traps are sent.

Default

162

Example

introscope.apm.catalyst.triagemapalert.snmp.destination.trap.port=162

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.apm.catalyst.triagemapalert.snmp.community

Set this property with the SNMP community string.

Default

public

Example

introscope.apm.catalyst.triagemapalert.snmp.community=public

introscope.apm.catalyst. triagemapalert.snmp.trigger

Set this property with the trigger type desired for SNMP trap notifications.

Default

3

Options

1

Each Period while Problem Exists

2

When Severity Increases

3

Whenever Severity Changes

4

Report Only Final State Whenever Severity

Example

introscope.apm.catalyst. triagemapalert.snmp.trigger=3

CA Technologies Prerelease Documentation. Authorized Use Only.

Application Map Objects

Set this property to specify the limit of application map objects to be considered as obsolete.

Default

300 (days)

Example

introscope.apm.data.obsolete.time=300

APMEnterpriseManager.properties

A variety of Enterprise Manager behaviors are controlled by properties defined in the Enterprise Manager properties file, *APMEnterpriseManager.properties*. (The others are IntroscopeEnterpriseManager.properties (see page 403) and apm-events-thresholds.xml.) This file mainly supports the application triage map.

After configuring most Enterprise Manager properties, you must restart the Enterprise Manager before any changes take effect. Some Enterprise Manager properties, however, use the Introscope hot configuration service. It is not necessary to restart the Enterprise Manager after making changes to these so-called "hot config" properties. For more information, see Hot Configuration Service *(see page 410).*

c3p0 properties

The c3p0 properties section configures properties for APM database connection pooling. For complete documentation on c3p0, see http://www.mchange.com/projects/c3p0/index.html.

Important! You are strongly encouraged to look at documentation for c3p0 properties in the *CA APM Sizing and Performance Guide* before changing property values.

c3p0.dataSourceName

Name of the data source to connect to. Replace the default name with the name of your APM database.

Default

apmDataSource

CA Technologies Product Reasonably Documentation. Authorized Use Only

c3p0.acquireIncrement

Determines how many connections at a time c3p0 will try to acquire when the pool has run out of connections.

Default

3

c3p0.idleConnectionTestPeriod

In seconds, the period of time c3p0 will attempt to test connections when they fall idle. At the end of this period, C3p0 will test all idle connections in the pool of this interval.

Default

120

c3p0.maxIdleTime

In seconds, the amount of time an idle connection will remain before being removed from the connection pool. Setting this to 0 means connections never time out from the pool.

Default

0

c3p0.initialPoolSize

Number of connections a pool will try to acquire upon startup. Should be between minPoolSize and maxPoolSize.

Default

2

c3p0.minPoolSize

Minimum number of connections a pool will maintain at any time.

Default

2

CA Technologies Prerelease Documentation. Authorized Use Only.

c3p0.maxPoolSize

Maximum number of connections a pool will maintain at any time.

Default

4

Hibernate specific properties

The Hibernate specific properties section configures aspects of the open-source Hibernate API, which enables transparent persistence, thereby simplifying storing of data in the APM database. For more information on Hibernate, see http://www.hibernate.org.

hibernate.jdbc.batch_size

Designates the JDBC batch size.

Default

10

hibernate.order_inserts

Enables batch inserts.

This is critical for enabling hibernate to batch insertions for each entity. Users are strongly cautioned not to edit this property.

Default

true

hibernate.order_updates

Enables batch updates. This is critical for enabling hibernate to batch updates for each entity. Users are strongly cautioned not to edit this property.

Default

true

Hibernate cache settings

The Hibernate cache settings section configures properties for data caching in the APM database.

CA Technologies Prerelease Documentation. Authorized Use Only.

hibernate.cache.provider_class

Designates the type of Hibernate caching to use.

Do not change this property.

Default

org.hibernate.cache.EhCacheProvider

hibernate.cache.use_second_level_cache

Enables Hibernate second level caching.

Default

true

hibernate.statement_cache.size

Designates the statement cache size.

Default

10

ApplicationTriageMap settings

This section allows you to customize settings for the application triage map.

hibernate.cache.provider_configuration_file_resource_path

Designates the Hibernate ehCache configuration file.

Default

apm-ehcache.xml

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.apm.query.max.results

Use this property to a threshold for the amount of data to be returned to the application triage map display by an agent. The value is, roughly, the number of dependencies for the frontend or business transaction, times the number of agents where the frontend or transaction is deployed and which have reported data for the period of time being queried.

Default

20000

apm-events-thresholds.xml

Properties defined in one of three Enterprise Manager properties files control a number of Enterprise Manager behaviors. This section describes the properties in *apm-events-thresholds.xml*. (The others are *IntroscopeEnterpriseMangager.properties* and *APMEnterpriseManager.properties* (see page 443).) The apm-events-thresholds.xml file, located in the *<EM_Home>/config* directory, contains properties for CA Introscope and CA CEM clamps and CA APM Important Events.

- When these clamps meet or exceed the configured threshold values, the Enterprise Manager reports alerts to the APM Status console Clamps tab.
- When the Important Events meet or exceed the configured threshold values, the Enterprise Manager reports alerts to the APM Status console Important Events table view.

Configure *only* the threshold values in the *threshold value* fields. Do *not* change any other fields such as clamp ID, importantEvent ID, and description because CA Introscopeuses these fields for system monitoring.

You do not need to restart the Enterprise Manager or monitored application to put into effect any apm-events-thresholds.xml configuration changes. These properties use the Introscope hot configuration (hot config) service. If a property uses hot config, then you do not need to restart the Enterprise Manager after making configuration changes. For more information, see Hot Configuration Service (see page 410).

Important! Back up the apm-events-thresholds.xml file before making any configuration changes. If configuration introduces an error into the apm-events-thresholds.xml file, the Enterprise Manager cannot read the file and puts default clamp threshold settings into effect. For example, if a close bracket is missing in apm-events-thresholds.xml, the Enterprise Manager cannot read the file. If new configurations or the updated file do not work properly, either troubleshoot the last-changed file or reinstate a previous backup that did work properly.

CA Technologies Product Documentation. Authorized Use Only

Introscope Clamp Properties

These properties define the limits set for Introscope-specific clamp properties. When the property thresholds are met or exceeded, the Enterprise Manager notifies the APM Status console.

These properties are hot configurable, meaning the values are applied when the Enterprise Manager checks the property file during a hot deploy. A hot deploy occurs about once every 60 seconds.

To disable clamps:

1. Open the apm-events-thresholds.xml file in the <*EM_Home*>/config directory.
2. Navigate to the clamps section.
3. Set the property value to 0.
4. Save and close apm-events-thresholds.xml.

Find additional Enterprise Manager clamps in the IntroscopeEnterpriseManager.properties file in the <*EM_Home*>/config directory. See IntroscopeEnterpriseManager.properties (see page 403).

introscope.clw.max.users

The property limits the number of Command LIne Workstations (CLW) that are allowed to be connected to the Enterprise Manager at any given time.

Use this clamp if CLW connections tend to slow down the Enterprise Manager. CLW connections can overload the Enterprise Manager by trying to extract too many metrics.

Notes

You can use this clamp for a standalone Enterprise Manager, MOM and CDV. However, Workstation connections are not typically made to Collectors.

Default

10

introscope.enterprisemanager.agent.connection.limit

The introscope.enterprisemanager.agent.connection.limit property limits the number of agents that can connect to an Enterprise Manager.

This property applies to both standalone Enterprise Managers and Collectors.

CA Technologies Pre-release Documentation. Authorized Use Only.

When the clamp is hit and a new agent tries to connect to the Enterprise Manager, the Enterprise Manager logs this and does not accept the new connection.

Default

400

Example

introscope.enterprisemanager.agent.connection.limit=400

introscope.enterprisemanager.agent.error.limit

The introscope.enterprisemanager.agent.error.limit property clamps error snapshots. The property defines the limit for number of error snapshots that the Enterprise Manager can handle for each agent per interval.

This property applies to both standalone Enterprise Managers and Collectors.

If the number of error snapshots that the agent reports exceeds the specified limit, all the rest of the errors are dropped by the Enterprise Manager.

Default

10

Example

introscope.enterprisemanager.agent.error.limit=10

introscope.enterprisemanager.agent.metrics.limit

Set this property with the value for which agent reporting is clamped when the value is exceeded. For example, if the default number of metrics, 50,000, is met, the agent reports only the 50,000 metrics and does not send new metrics to the Enterprise Manager. An inappropriately configured agent can create thousands of metrics in quick succession, overloading the Enterprise Manager.

Use this property to prevent metric overload. If an agent generates more than the specified number of metrics, the Enterprise Manager will drop them. If this occurs, modify agent configuration to reduce the volume of metrics reported. For more information, see the *CA APM Java Agent Implementation Guide* or the *CA APM .NET Agent Implementation Guide*.

Default

50000

CA Technologies Prerelease Documentation. Authorized Use Only.

Note

This clamp property works in conjunction with the introscope.agent.metricClamp property located in the Introscope.Agent.profile file.

If introscope.enterprisemanager.agent.metrics.limit clamp value is triggered before the introscope.agent.metricClamp value, then the Enterprise Manager reads agent metrics but does not report them in the Investigator metric browser tree.

If the introscope.agent.metricClamp clamp value is triggered before the introscope.enterprisemanager.agent.metrics.limit clamp value, the agent stops sending metrics to the Enterprise Manager.

introscope.enterprisemanager.agent.trace.limit

The introscope.enterprisemanager.agent.trace.limit property clamps Transaction Traces. The property defines the limit for number of transaction events that the Enterprise Manager can handle for each agent per interval.

This property applies to both standalone Enterprise Managers and Collectors.

If the number of Transaction Traces that the agent reports exceeds the specified limit, all the rest of the traces are dropped by the Enterprise Manager.

Default

1000

Example

introscope.enterprisemanager.agent.trace.limit=1000

introscope.cdv.max.users

This property limits the number of Cross-cluster Data Viewers (CDVs) that can connect to a single Collector. The Cross-cluster Data Viewer usually connects to multiple Collectors across two or more clusters.

A CDV connection can add to Collector performance load by extracting metrics. Use this clamp to restrict the number of CDV connections that a Collector can accept.

Notes

Default

5

CA Technologies Pre-release Documentation. Authorized Use Only

introscope.enterprisemanager.disconnected.historical.agent.limit

The property limits the number of disconnected but still mounted historical agents. This clamp minimizes the use of Enterprise Manager run-time memory.

This clamp limits both historical agents that:

- The Enterprise Manager can unmount automatically.
- CA APM users must unmount manually.

For information about mounting and unmounting agents, see the *CA APM Workstation User Guide*.

When the clamp threshold value is hit, the Enterprise Manager tries to unmount the oldest disconnected historical agent that can be unmounted automatically. If there is no historical agent that the Enterprise Manager can automatically unmount, this means that CA APM users mounted manually all the disconnected historical agents. The Enterprise Manager never tries to unmount a disconnected historical agent that a CA APM user mounted manually.

Example

Notes

The Enterprise Manager notifies the APM Status console when the following occur:

- Disconnected mounted agent is unmounted due to the clamp.
- Manual agent mount fails.

This clamp does *not* prune disconnected mounted agents when the following occur:

- An Enterprise Manager is low on memory.
- A disconnected historical agent times out and is automatically unmounted.

A disconnected historical agent stays mounted when the following occur:

- The historical agent becomes disconnected and can be automatically unmounted.
- The disconnected historical agent count is below the introscope.enterprisemanager.disconnected.historical.agent.limit value.

The Enterprise Manager mounts a disconnected historical agent when the following occurs:

- A CA APM user tries to mount manually a disconnected historical agent.
- The disconnected historical agent count is below the introscope.enterprisemanager.disconnected.historical.agent.limit value.

CA Technologies Prerelease Documentation. Authorized Use Only

The Enterprise Manager searches for mounted disconnected agents that can be automatically unmounted when the following occur:

- A historical agent that can be automatically unmounted becomes disconnected.
- A CA APM user tries to mount a historical agent manually.
- The disconnected historical agent count is equal to the introscope.enterprisemanager.disconnected.historical.agent.limit value.

If the Enterprise Manager finds a mounted disconnected agent, then it unmounts the oldest disconnected historical agent and keeps newly disconnected agent mounted. If there is no historical agent that the Enterprise Manager can unmount automatically, then the Enterprise Manager immediately unmounts the newly disconnected historical agent. If a CA APM user then tries to mount manually a historical agent, then the:

- Workstation displays an error message instructing the user to unmount some historical agents to make room to mount new historical agents.
- Enterprise Manager sends notification to the APM Status console.

Default

300

introscope.enterprisemanager.events.limit

The introscope.enterprisemanager.events.limit property clamps transaction traces. The property defines the limit for number of transaction events processed by the Enterprise Manager for each interval.

This property applies to standalone Enterprise Managers, Collectors, and the MOM.

This clamp is applied to the total number of all transaction events generated by all agents. When the clamp is hit, all the rest of the transaction events reported by the agent are dropped for the remainder of that interval.

Default

1250

Example

introscope.enterprisemanager.events.limit=1250

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.metrics.historical.limit

This property limits the number of historical metrics in the Enterprise Manager. The limit is checked every 15 seconds, so the Enterprise Manager may go above this value before the clamp is applied. When the Enterprise Manager detects that it has exceeded this limit, the Enterprise Manager will stop all agents from registering new metrics. Old metrics will continue to report.

Default

1200000

introscope.enterprisemanager.metrics.live.limit

This property limits live metrics currently reporting from agents. When the Enterprise Manager detects that it has exceeded the limit, the Enterprise Manager will stop all agents from registering new metrics. The limit is checked every 15 seconds, so the Enterprise Manager may go above this value before the clamp is applied. Old metrics will continue to report. Metrics for which we have historical data will never be clamped.

Example

If an agent has 60k metrics it has reported to an Enterprise Manager before clamping, and then the clamp of 50k is applied; all 60k old metrics can still register with the Enterprise Manager. If however, any new metrics attempt to register for which we historical data is not already collected, then they will not be allowed to register.

Default

500000

introscope.enterprisemanager.transactionevents.storage.max.disk.usage

Determines the maximum desired disk usage in MB for transaction trace storage. If this maximum is exceeded, then the daily historical trace stores are deleted starting with the oldest first until the total historical trace storage size is below this property value. Even if the size of this property value is exceeded, stored traces may not be deleted because the size of the trace index is not considered when determining what historical trace stores to delete.

A negative or zero value indicates that all historical trace values stored should be deleted.

If the Transaction Events database is larger than this value, historical storage files are deleted until the total database size is less than this value, starting with the oldest storage file.

CA Technologies Pre-release Documentation. Authorized Use Only

This might delete historical data even if it has not reached the fully aged date. The current live storage is not deleted, even if its size exceeds the max database size.

If you determine that the Transaction Events database is too large, you can delete database files manually, when the Enterprise Manager is shut down. When the Enterprise Manager is restarted, the Transaction Trace storage system recognizes that the files have been deleted manually and adjusts the trace index accordingly.

Notes

If you delete files manually, delete only files in the traces directory that have this format: *traces_yyyyMMdd.db*, where *yyyyMMdd* represents the year, month, and day of the traces in that file.
Do not delete any other files in that directory or subdirectory.

Default

1024

introscope.workstation.max.users

The property limits the number of Workstations that are allowed to be connected to the Enterprise Manager at any given time.

Use this clamp if Workstation connections tend to slow down the Enterprise Manager. Workstation connections can overload the Enterprise Manager by trying to extract too many metrics.

Note

You can use this clamp for a standalone Enterprise Manager, MOM and CDV. However, Workstation connections are not typically made to Collectors.

Default

10

CA Technologies Internal/Restricted Use Only. Please Do Not Distribute Outside CA Technologies Without Authorization.

Customer Experience Clamp Properties

These properties define the limits set for CA CEM-specific clamp properties. When these property thresholds are met or exceeded, the Enterprise Manager notifies the APM Status console.

These properties are hot configurable, meaning the values are applied when the Enterprise Manager checks the property file during a hot deploy. A hot deploy occurs about once every 60 seconds.

To disable clamps:

1. Open the apm-events-thresholds.xml file in the <*EM_Home*>/config directory.
2. Navigate to the clamps section.
3. Set the property value to 0.
4. Save and close apm-events-thresholds.xml.

introscope.enterprisemanager.max.transaction.user.groups

The maximum number of application users retrieved.

Limits the Admin.Configuration.UserGroups.UserGroup.Users page to the top N records

Too many application users can negatively affect performance.

Options

numeric value

Maximum allowed value is 10,000.

Default

1000

Example

introscope.enterprisemanager.max.application.user.rows =1000

Notes

When the maximum is reached:

- User groups are no longer created automatically.
- Users are assigned to the New Users group instead of creating a new group.
- The System > Events list displays a warning message.

See Managing users and groups on CA CEM (see page 291).

CA Technologies Prerelease Documentation. Authorized Use Only.

Limiting the number of changes listed since last synchronization

Before synchronizing monitors, you can review configuration changes that have made since the last synchronization. This is described in Viewing configuration changes since last synchronization (see page 289).

The configuration changes appear in the Domain Configuration History table on the Monitors page. The introscope.enterprisemanager.max.number.domain.configuration.changes property determines the maximum number of rows in this table.

introscope.enterprisemanager.max.number.domain.configuration.changes

Sets the maximum number of changes that appear in the Domain Configuration History table when View Changes Since Last Synchronization is clicked (on the Setup > Monitors).

If the value is -1, then all the changes appear.

Options

numeric value

Default

-1

Example

introscope.enterprisemanager.max.number.domain.configuration.changes=1000

Notes

See Synchronizing monitors (see page 288).

introscope.enterprisemanager.max.transaction.user.groups

The maximum number of automatically created user groups.

This property works in conjunction with automatically created user groups via request attribute or by IP subnet.

Too many user groups, even groups without any users, can negatively affect performance.

Options

numeric value

Maximum allowed value is 10,000.

Default

5000

CA Technologies Prerelease Documentation Authorized Use Only.

Example

introscope.enterprisemanager.max.transaction.user.groups=5000

Notes

When the maximum is reached:

- User groups are no longer created automatically.
- Users are assigned to the New Users group instead of creating a new group.
- The System > Events list displays a warning message.

See Managing users and groups on CA CEM (see page 291).

introscope.enterprisemanager.transaction.discovery.max.nonidentifying.components

Maximum number of discovered non-identifying components to keep for a currently defined transaction. (This is in addition to any original non-identifying components.)

Options

numerical value (a positive integer)

Default

50

Example

introscope.enterprisemanager.transaction.discovery.max.nonidentifying.component=60

Notes

This limit is important for transactions where non-identifying components change frequently, for example, if a component is a GIF file that changes daily. Setting this limit prevents saving an excessive number (of probably obsolete) non-identifying components for a transaction.

A CA CEM event called Limit exceeded adding new components is logged each time a new non-identifying component replaces an old one.

For general information about automatically discovering non-identifying components, see *CA APM Transaction Definition Guide*.

CA Technologies Prerelease Documentation. Authorized Use Only.

Setting Important Event Thresholds

These properties define the limits used to notify the APM Status console about events important to Enterprise Manager health monitoring. They are hot configurable, meaning the values are applied when the Enterprise Manager checks the property file during a hot deploy. A hot deploy occurs about once every 60 seconds.

To disable important event properties:

1. Open the apm-events-thresholds.xml file in the <*EM_Home*>/config directory.
2. Navigate to the Important Events section.
3. Set the property value to 0.
4. Save and close apm-events-thresholds.xml.

apm.events.thresholds.config.error

This property monitors the apm-events-thresholds.xml file for syntax errors. When the Enterprise Manager finds an error, it:

- notifies the APM Status console Important Events table view.
- reverts all the apm-events-thresholds.xml file clamp ?and important event? properties values to their default values.

This property applies to all the Collectors in a cluster or to a Standalone Enterprise Manager.

Default

Not applicable. This property is for notification purposes only and cannot be modified.

CA Technologies Preferred Documentation. Authorized Use Only.

introscope.enterprisemanager.cpu.usage

This property defines the limit, in percent (%) at which Enterprise Manager CPU usage is considered an Important Event. When this threshold is met or exceeded, the Enterprise Manager notifies the APM Status console Important Events table view.

This property is based on Enterprise Manager monitoring of the CPU Capacity (%) metric. CPU usage greater than 60% indicates Enterprise Manager resource problems. For information about CPU usage and Enterprise Manager performance, see the *CA APM Sizing and Performance Guide*.

This property applies to all the Collectors in a cluster or to a Standalone Enterprise Manager.

Default

60

Example

introscope.enterprisemanager.cpu.usage=60

introscope.enterprisemanager.database.connectivity

This property monitors the connection between the Enterprise Manager and the APM database. When the connection is lost, the Enterprise Manager notifies the APM Status console Important Events table view.

This property applies to all the Collectors in a cluster or to a Standalone Enterprise Manager.

Default

Not applicable. This property is for notification purposes only and cannot be modified.

introscope.enterprisemanager.harvest.cycle.diff

This property defines the limit, in number of time slices, that the Collector and MOM harvest cycle can be out of sync before notification as an Important Event. When this threshold is met or exceeded, the Enterprise Manager notifies the APM Status console Important Events table view.

This property is based on comparing the Collector metrics harvest cycle to that of the MOM. Ideally within the same 15-second time slice:

- the Collector collects and aggregates agent metrics data

CA Technologies Pre-release Documentation
Authorized Use Only.

- sends metrics data to the MOM
- MOM aggregates data from the Collector
- MOM reports the data to the Workstation.

The Collector does not send data to the MOM until the aggregation is complete. When this completion happens in a later time slice, this results in the Collector and MOM data aggregation being out of sync. The Collector is considered to be running behind the MOM in terms of processing.

If the Collector and MOM are more than four time slices of sync, the Collector considered is slow in providing metrics data to the MOM.

This property applies to all the Collectors in a cluster.

Example

Default

4

Example

introscope.enterprisemanager.harvest.cycle.diff=4

introscope.enterprisemanager.events.limit

This property defines the limit, in milliseconds (ms), at which Enterprise Manager harvest duration is considered an Important Event. When this threshold is met or exceeded, the Enterprise Manager notifies the APM Status console Important Events table view.

This property is based on Enterprise Manager monitoring of the Harvest Duration (ms) metric. A harvest duration longer than 3500 ms indicates Enterprise Manager resource problems. For information about harvest duration and Enterprise Manager performance, see the *CA APM Sizing and Performance Guide*.

This property applies to all the Collectors in a cluster or to a Standalone Enterprise Manager.

Default

3500

Example

introscope.enterprisemanager.harvest.duration=3500

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.smartstor.duration

This property defines the limit, in milliseconds (ms), at which Enterprise Manager SmartStor duration is considered an Important Event. When this threshold is met or exceeded, the Enterprise Manager notifies the APM Status console Important Events table view.

This property is based on Enterprise Manager monitoring of the SmartStor Duration (ms) metric. A SmartStor duration longer than 3500 ms indicates Enterprise Manager resource problems. For information about SmartStor duration and Enterprise Manager performance, see the *CA APM Sizing and Performance Guide*.

This property applies to all the Collectors in a cluster or to a Standalone Enterprise Manager.

Default

3500

Example

introscope.enterprisemanager.smarstor.duration=3500

ResourceMetricMap.properties

Resource metric categories from various agents are reported in the Workstation. The resource metric categories are predefined. While you cannot add or delete the predefined resource metric groups, you can add, change, or remove the resource metric mappings that specify the data reported for each resource metric group if you have administrator rights.

To map resource metrics to resource metric categories:

1. Verify that JMX reporting is enabled on the agent. For more information, see *Enabling JMX Reporting* in the *CA APM Java Agent Implementation Guide*.
2. Edit the *ResourceMetricMap.properties* file located on the MOM in the *<EM_Home>/Config* directory by mapping your resource metrics to the resource metric categories. For more information about mapping resource metrics, see your application server documentation.
   a. Add metric entries using a *<categoryName>.path<recordID>:<agentMetricPath>* format. For example:
   `threads.used.path.1=Tomcat|ThreadPool|getCurrentThreadsBusy.path.1=CPU:Processor Count`

   Where *memory.utilization* is the name of the metric, and *CPU:Processor Count* is its default metric path.

CA Technologies Pre-release Documentation
Authorized Use Only.

b. Add incremental record ID metrics using numbers. For example:

```
cpu.utilization.path.1=Non Existent metric
cpu.utilization.path.2=CPU|Processor 0:Utilization % (aggregate)
```

You use multiple paths in the same metric group to configure metrics with different paths under the same name. This allows you to, for example, get the same metric from different application servers where a frontend is deployed.

3. Restart the MOM and collectors.

The newly mapped resource metric values are reported in the Workstation.

Example of a ResourceMetricMap.properties file configured for multiple instances of the same metric

Here is a section of a ResourceMetricMap.properties file where multiple instances of the Threads in Use metric have been configured to point to three collectors in a cluster.

```
#CPU Utilization
cpu.utilization.path.1=CPU:Utilization % (process)
#Time Spent in GC
memory.utilization.path.1=GC Monitor|Percentage of Time Spent in GC
during last 15 minutes
#Threads in Use
threads.used.path.1=WebSpherePMI|threadPoolModule|WebContainer:Act
iveCount
threads.used.path.2=WebLogic|JMX Aggregate|Thread Pool:Waiting
Request Count
threads.used.path.3=Tomcat|ThreadPool|default thread
pool:getCurrentThreadsBusy
#Threads Available
#JDBC Connections in Use
connections.used.path.1=WebSpherePMI|connectionPoolModule:WaitingT
hreadCount
connections.used.path.2=WebLogic|JMX Aggregate|JDBC Connection
Pool:Waiting Thread Count
#JDBC Connections Available
```

You can map resource metrics to these resource metric categories:

CPU Utilization (see page 463)
Time Spent in GC (see page 463)
Threads in Use (see page 463)
Threads Available (see page 463)
JDBC Connections in Use (see page 463)
JDBC Connections Available (see page 464)

CA Technologies Prerelease Documentation
Unauthorized Use Only

CPU Utilization

CPU utilization measures the amount of CPUs being used.

Mapping example:

```
cpu.utilization.path.1=CPU:Utilization % (process)
cpu.utilization.path.2=Some other metrics path
```

Time Spent in GC

Time Spent in GC measures the time spent in the garbage collector to help you understand whether the amount of work is significant for application performance.

Mapping example:

```
memory.utilization.path.1=memory utilization path
```

Threads in Use

Threads in Use measures the number of threads in use.

Mapping example:

```
threads.used.path.1=threads used path
```

Threads Available

Threads Available measures the number of threads in use.

Mapping example:

```
threads.available.path.1=threads available path
```

JDBC Connections in Use

JDBC Connections in Use measures the number of Java Database Connectivity instances are active.

Mapping example:

```
connections.used.path.1=connections used path
```

CA Technologies Prerelease Documentation. Authorized Use Only.

JDBC Connections Available

JDBC Connections in Use measures the number of Java Database Connectivity are available to be used.

Mapping example:

```
connections.available.path.1=connections available path
```

Catalyst.properties

Catalyst properties allow you to configure CA APM to load CA Catalyst data and filter the alerts reported from Catalyst Data.

More Information:

catalyst.host (see page 465)
catalyst.protocol (see page 465)
catalyst.port (see page 465)
catalyst.login (see page 465)
catalyst.password (see page 465)
catalyst.rest.entrypoint (see page 466)
catalyst.rest.alerts.pollingperiodsec (see page 466)
catalyst.rest.virtualhosts.pollingperiodsec (see page 466)
catalyst.product.baseurl (see page 467)
catalyst.alert.filter.total (see page 467)
catalyst.alert.filter.time (see page 467)
catalyst.alert.filter.products (see page 467)
catalyst.alert.filter.types (see page 468)
catalyst.alert.filter.severities (see page 469)
catalyst.rest.maxFailures (see page 469)
catalyst.rest.maxAlertsPerPage (see page 470)
catalyst.rest.alerts.threads (see page 470)
catalyst.rest.virtualhosts.threads (see page 470)
catalyst.entity.cache.size (see page 471)
catalyst.entity.cache.expirationsec (see page 471)
catalyst.database.alerts.enabled (see page 471)
catalyst.database.map.name (see page 472)
Configuring APM to Load Catalyst Data (see page 227)
Configure CA APM to Load CA Catalyst Data (see page 228)
Configure the CatalystPolicy.xml File (see page 229)

CA Technologies Pre-Release Documentation. Internal Use Only.

catalyst.host

Set this required property with the CA Catalyst RESTful API host name. After you uncomment this line, add the value to CA Catalyst.

Default

This property is disabled by default.

catalyst.protocol

Set this required property with the protocol for the CA Catalyst web access.

Example

catalyst.protocol = http

catalyst.port

Set this required property with the port number for the CA Catalyst web access.

Example

catalyst.port = 7070

catalyst.login

Set this required property with the user name used to log in to CA Catalyst.

Example

catalyst.login =<enter your CA Catalyst user name>

catalyst.password

Set this required property with the password used to log in to CA Catalyst.

Example

catalyst.password =<enter your CA Catalyst password>

CA Technologies Prerelease Documentation. Authorized Use Only.

catalyst.rest.entrypoint

Set this required property with URL context for the CA Catalyst RESTful API entry point.

Example

catalyst.rest.entrypoint = /ssaweb/search

catalyst.rest.alerts.pollingperiodsec

Set this optional polling property to determine how often CA Catalyst is checked for new alerts.

Default

15 (seconds)

Example

catalyst.rest.alerts.pollingperiodsec = 15

catalyst.rest.virtualhosts.pollingperiodsec

Set this optional property to determine how often CA Catalyst is checked for virtual host updates.

Default

60 (seconds)

Example

catalyst.rest.virtualhosts.pollingperiodsec = 60

CA Technologies Prerelease Documentation. Authorized Use Only.

catalyst.product.baseurl

Set this property with the base URL addresses for the contributing products.

Example

catalyst.product.baseurl = CA:09998-><http://www.website1.com/>

catalyst.product.baseurl = CA:00048-><http://www.website2.com/>

catalyst.product.baseurl = CA:00002-><http://www.website3.com/>

catalyst.alert.filter.total

Set this optional property to set the maximum number of CA Catalyst alerts to load into the Workstation.

Default

250

Example

catalyst.alert.filter.total = 250

catalyst.alert.filter.time

Set this optional property to set the length of time for CA Catalyst alerts appear in the Workstation.

Default

72 (hours)

Example

catalyst.alert.filter.time = 72

catalyst.alert.filter.products

Set this optional property to exclude all alerts for the specified products.

Example

catalyst.alert.filter.products = CA:00030

CA Technologies Prerelease Documentation. Authorized Use Only.

catalyst.alert.filter.types

Set this optional property to identify the CA Catalyst alert types to exclude from appearing in the Workstation. Separate alert types with a comma.

Options

- Risk
- Risk-Capacity
- Risk-Change
- Risk-Fault
- Risk-Security
- Risk-Performance
- Risk-Root Cause
- Quality
- Compliance
- Cost

Example

catalyst.alert.filter.types = Risk-Security,Compliance,Cost

CA Technologies Prerelease Documentation. Authorized Use Only.

catalyst.alert.filter.severities

Set this optional property to determine which CA Catalyst alert severities should be excluded from appearing in the Workstation. Define a range using a hyphen (-) between two severity values, open range using a hyphen to the right or to the left of severity value, or comma to set multiple values.

Default

Unknown

Options

The following options can be set for this property and allows you to use open ranges:

- Normal
- Informational
- Unknown
- Minor
- Major
- Critical
- Fatal

Example

catalyst.alert.filter.severities = -Unknown

catalyst.rest.maxFailures

Set this optional property to specify the maximum number of CA Catalyst failures to allow before canceling the polling.

Default

10

Example

catalyst.rest.maxFailures = 10

CA Technologies Prerelease Documentation. Authorized Use Only.

catalyst.rest.maxAlertsPerPage

Set this optional property to specify the maximum number of CA Catalyst alerts to return in a single REST query.

Default

250

Example

catalyst.rest.maxAlertsPerPage = 250

catalyst.rest.alerts.threads

Set this optional property to specify the maximum number of threads to use for alerts.

Default

2

Example

catalyst.rest.alerts.threads = 2

catalyst.rest.virtualhosts.threads

Set this optional polling property to specify the maximum number of threads to use for virtual hosts.

Default

1

Example

catalyst.rest.virtualhosts.threads = 1

CA Technologies Prerelease Documentation. Authorized Use Only.

catalyst.entity.cache.size

Set this optional property to specify the size of entity in-memory cache to use for CA Catalyst.

Default

3000 (megabytes)

Example

catalyst.entity.cache.size = 3000

catalyst.entity.cache.expirationsec

Set this optional property to specify the expiration time in seconds of the entity cache.

Default

36000 (seconds)

Example

catalyst.entity.cache.expirationsec = 36000

catalyst.database.alerts.enabled

Set this property to send CA Catalyst database alerts to the Workstation. By default CA Catalyst database alerts are sent to the Workstation. Change the setting to *false* to stop sending CA Catalyst database alerts.

Default

true

Default

catalyst.database.alerts.enabled=true

CA Technologies Prerelease Documentation. Authorized Use Only.

catalyst.database.map.name

Set this optional property to specify the location of the file that maps database connection strings to CA Catalyst database parameters. Set this property only if out-of-the-box mapping does not work properly.

Example

catalyst.database.map.name = CatalystDatabaseMap.txt

EMservice.conf

The *EMService.conf* file, in <*EM_Home*>\bin contains settings used when the Enterprise Manager is run as a Windows Service.

Note: Before configuring these properties, you must unregister the Enterprise Manager Service. After configuring properties, re-register the Enterprise Manager Service.

wrapper.java.command

The Java application wrapper command.

Default

jre\\bin\\java.exe

wrapper.java.mainclass

The Java main class wrapper command.

Default

org.tanukisoftware.wrapper.WrapperSimpleApp

wrapper.java.classpath.x

The Java Classpath that includes the *wrapper.jar*. Class path elements should be added in a sequential manner, such as follows:

- *wrapper.java.classpath.1=./lib/ServiceWrapper.jar*
- *wrapper.java.classpath.2=./launcher.jar*
- *wrapper.java.classpath.3=*

CA Technologies Prerelease Documentation. Authorized Use Only

wrapper.java.library.path.1

The Java library path is the location of the *Wrapper.DLL* or *libwrapper.so*.

Default

./lib

wrapper.java.additional.x=

Additional Java parameters that are applicable only on 64-bit Windows.

The first additional parameter must be *-Xrs*. From there you can add parameters as needed. For example:

- *wrapper.java.additional.1=-Xrs*
- *wrapper.java.additional.1=-Djava.awt.headless=false*
- *wrapper.java.additional.2=-showversion*
- *wrapper.java.additional.3=-XX:+UseConcMarkSweepGC*
- *wrapper.java.additional.4=-XX:+UseParNewGC*

Default

-Xrs

wrapper.java.initmemory

Wrapper's initial memory size in megabytes.

Default

128

wrapper.java.maxmemory

Wrapper's maximum memory size in megabytes.

Default

512

CA Technologies Prerelease Documentation. Authorized Use Only.

wrapper.app.parameter.x

Application parameters. You can add parameters as needed. For example:

- *wrapper.app.parameter.1=org.eclipse.core.launcher.Main*
- *wrapper.app.parameter.2=-consolelog*
- *wrapper.app.parameter.3=-noExit*
- *wrapper.app.parameter.4=-product*
- *wrapper.app.parameter.5=com.wily.introscope.em.product*
- *wrapper.app.parameter.6=-install*
- *wrapper.app.parameter.7=./product/enterprisemanager*
- *wrapper.app.parameter.8=-configuration*
- *wrapper.app.parameter.9=./product/enterprisemanager/configuration*

wrapper.startup.timeout

The number of seconds to allow between the time that the wrapper launches the JVM process and the time that the JVM side of the wrapper responds that the application has started. Entering a 0 value means the system will never time out.

Default

30 (seconds)

wrapper.console.format

The format of output for the console.

Default

PM

wrapper.console.loglevel

The log level for the console output.

Default

INFO

CA Technologies Prerelease Documentation. Authorized Use Only.

wrapper.logfile

The name and location of the log file used for wrapper output logging.

Default

./logs/EMService.log

wrapper.logfile.format

The format of the output for the log file.

Default

LPTM

wrapper.logfile.loglevel

The log level for log file output.

Default

INFO

wrapper.logfile.maxsize

The maximum size that the log file will be allowed to grow to before the log is rolled. Size is specified in bytes. With the *0* default value disabling the log roll. You can abbreviate with a *k* (*kb*) or *m* (*mb*) suffix. For example, *10m = 10 megabytes*.

Default

0

wrapper.logfile.maxfiles

The maximum number of rolled log files that will be allowed before old files are deleted. The default value of 0 implies no limit.

Default

0

CA Technologies Prerelease Documentation. Authorized Use Only.

wrapper.syslog.loglevel

The log level for the system and event log output.

Default

NONE

Notes

Do not modify any of the following *ntservice* properties when an application using this configuration file has been installed as a service. Please uninstall the service before modifying this section. After you have made modifications, the service can then be reinstalled.

wrapper.ntservice.name

Name of the service.

Default

IScopeEM

wrapper.ntservice.displayname

Display name of the service.

Default

Introscope Enterprise Manager

wrapper.ntservice.description

Description of the service.

Default

Introscope Enterprise Manager

CA Technologies Prerelease Documentation. Authorized Use Only.

wrapper.ntservice.dependency.x

Service dependencies. Add dependencies as needed starting from 1. For example:

- *wrapper.ntservice.dependency.1=*
- *wrapper.ntservice.dependency.2=*
- *wrapper.ntservice.dependency.3=*

wrapper.ntservice.starttype

Mode in which the service is installed. **Settings** include *AUTO_START* or *DEMAND_START*.

Default

AUTO_START

wrapper.ntservice.interactive

Setting this value to *true*, starts the service to interact with the desktop.

Default

false

Introscope Enterprise Manager.lax

The *Introscope Enterprise Manager.lax file* is located in the *<EM_Home>* directory. The Enterprise Manager runtime behavior is controlled by properties defined in this file.

lax.application.name

The default name of the Enterprise Manager executable.

WARNING! Do not edit.

Default

Introscope Enterprise Manager.exe

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.class.path

The Java classpath necessary to run this application.

Can be appended with a JDBC driver and license.

Can be separated by colons (Mac OS/UNIX) or semicolons (Windows).

Example
lax.class.path=launcher.jar;lax.jar

Default

<blank>

lax.command.line.args

This is passed to the main method. When entering this value, be sure to quote arguments with spaces. For example:
lax.command.line.args=$CMD_LINE_ARGUMENTS$ -consolelog -noExit -product com.wily.introscope.em.product -install ".\\product\\enterprisemanager" -configuration ".\\product\\enterprisemanager\\configuration"

Default

<blank>

lax.dir

The path to the directory holding LaunchAnywhere's native launcher.

Default

<blank>

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.main.class

The class that contains the main method for the application.

Example

lax.main.class=org.eclipse.core.launcher.Main

Default

main

lax.main.method

The main method of LaunchAnywhere's Java launcher. This property should not be adjusted from its default value: *lax.nl.java.launcher.main.method=main*

Default

main

lax.nl.current.vm

This property represents the virtual machine to be used the next time the Enterprise Manager is started. It can be set to use an installed JRE. See the *CA APM Installation and Upgrade Guide* for JVM version requirements.

Example

If you are configuring this property to use the bundled JVM provided, it would display as follows:

```
lax.nl.current.vm=jre\bin\java.exe
```

Default

The default value varies by operating system.

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.nl.java.launcher.main.class

The main class of LaunchAnywhere's java launcher. This property should not be adjusted from its default value: *lax.nl.java.launcher.main.class=com.zerog.lax.LAX*

Default

com.zerog.lax.LAX

lax.nl.java.launcher.main.method

The main method of LaunchAnywhere's Java launcher. This property should not be adjusted from its default value: *lax.nl.java.launcher.main.method=main*

Default

main

lax.nl.java.option.additional

Optional Java arguments (such as heap size).

Example

lax.nl.java.option.additional=-Xms512m -Xmx1024m -Djava.awt.headless=false -XX:MaxPermSize=256m -XX:+UseConcMarkSweepGC -XX:+UseParNewGC -XX:CMSInitiatingOccupancyFraction=50 -Dosgi.clean=false

The following switches are recommended for 32-bit Sun JVMs:

```
#   -showversion -XX:+UseConcMarkSweepGC -XX:+UseParNewGC
-XX:CMSInitiatingOccupancyFraction=50 -XX:MaxPermSize=256m
lax.nl.java.option.additional=-Xms512m -Xmx1024m
-Djava.awt.headless=false -XX:MaxPermSize=256m
-showversion -XX:+UseConcMarkSweepGC
-XX:+UseParNewGC -XX:CMSInitiatingOccupancyFraction=50
```

Default

<blank>

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.nl.message.vm.not.loaded

The message displayed in a user dialog if no virtual machine from the *lax.nl.valid.vm.*list can be found. This property is internal to the InstallAnywhere launcher. Introscope has separate Java requirements. See the *CA APM Installation and Upgrade Guide* for the JVM requirements.

The lax.nl.message.vm.not.loaded=LaunchAnywhere either could not find a Java VM, or the Java VM on this system is too old. LaunchAnywhere requires a Java 2 VM in order to launch Introscope.

lax.nl.valid.vm.list

A string containing one or more of *[ALL JDK JRE J1 J2 JRE_J1 JDK_J1 JRE_J2 JDK_J2 MSJ]* delimited by spaces or commas. If the native launcher cannot find the current virtual machine, it will search for values in this list: *lax.nl.valid.vm.list=J2.*

Default

J2

lax.nl.win32.microsoftvm.min.version

The minimum version of Microsoft's virtual machine with which this launcher will run. For example:

lax.nl.win32.microsoftvm.min.version=2750.

Default

2750

lax.root.install.dir

The path to the *installdir* magic folder. For example:

lax.root.install.dir=.

Default

<blank>

CA Technologies Prerelease Documentation
Authorized Use Only.

lax.stderr.redirect

Standard Error Output.
Set to:

- blank for no output.
- *console* to send to a console window.
- any path to a file to save to the file.

Default

console

lax.stdin.redirect

Set to:

- blank for no output.
- *console* to send to a console window.
- any path to a file to save to the file.

Notes

If you are running the Enterprise Manager in nohup mode for UNIX platforms, do NOT change this setting, or the Enterprise Manager may fail to start.

If local requirements dictate that you set this property to blank, you MUST set introscope.enterprisemanager.disableInteractiveMode (see page 408) to true.

Default

console

CA Technologies Pereleases Documentation. Authorized Use Only.

lax.stdout.redirect

Standard Output.

Set to:

- blank for no output.
- *console* to send to a console window.
- any path to a file to save to the file.

Default

console

lax.user.dir

If left blank, this property will cause the native launcher to not alter the platform default behavior for setting the user directory.
To override this you may set this property to a relative or absolute path, with therelative paths being relative to the launcher.

Example
`lax.user.dir=.`

Default

<blank>

lax.version

The version of LaunchAnywhere that created the properties file. For example, *lax.version=<VersionNumber>*

Default

10.0

IntroscopeWorkstation.properties

A variety of Workstation behaviors are controlled by properties in this file.

CA Technologies Prerelease Documentation. Authorized Use Only

log4j.logger.Workstation

Controls the amount of logging detail and the output location of the Workstation log.

Level of detail value can be:

- *FATAL*
- *WARN*
- *INFO*
- *VERBOSE#com.wily.util.feedback.Log4JSeverityLevel*
  (This value increases the level of detail logged.)
- *DEBUG*
- *TRACE*

Destination value can be:

- *console*
- *logfile*

Default

INFO, console

log4j.appender.logfile.File

Name and location of Workstation logfile.

Default

IntroscopeWorkstation.log

introscope.workstation.graph.antialiased

This controls the anti-aliasing of lines and symbols in graphs. If this value is set to false, it disables anti-aliasing.

Default

true

CA Technologies Prerelease Documentation. Authorized Use Only.

transport.http.proxy.host

An HTTP Tunneling Proxy Server property that specifies the host the Workstation is tunneling over HTTP. If this method is chosen, it 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.

Default

<blank>

transport.http.proxy.port

An HTTP Tunneling Proxy Server property that specifies the port the Workstation is tunneling over HTTP. If this method is chosen, it 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.

Default

<blank>

transport.http.proxy.username

The username applies only if the proxy server requires authentication, at which point you would enter a value to represent a user name.

Default

<blank>

transport.http.proxy.password

The username applies only if the proxy server requires authentication, at which point you would enter a value to represent a password, that a user must enter.

Default

<blank>

CA Technologies Prerelease Documentation. Authorized Use Only

transport.tcp.truststore

If you configure SSL this optional property specifies the location of the truststore that contains the trusted Enterprise Manager certificates. This value is either an absolute path or a path relative to the Workstation's working directory.

If you are on Windows, backslashes must be escaped. For example:

transport.tcp.truststore=C:\\Introscope\\config\\internal\\server\\keystore

If no truststore is specified, the Workstation trusts all certificates.

Default

<blank>

transport.tcp.trustpassword

The password for the truststore.

Default

password

transport.tcp.truststore

The location of a keystore containing the workstation's certificate. Either an absolute path or a path relative to the Workstation's working directory. If you are on Windows, backslashes must be escaped.

Example

..\\Introscope\\config\\internal\\server\\keystore

transport.tcp.keystore

The property for the location of the keystore. This is needed only if the Enterprise Manager requires client authentication.

Default

internal/server/keystore

CA Technologies Prerelease Documentation. Authorized Use Only.

transport.tcp.keypassword

The password for the keystore.

Default

password

transport.tcp.ciphersuites

Set this to enable cipher suites using a comma-separated list of cipher suites. If this property is not specified, the default enabled cipher suites will be used.

Default

<blank>

Notes

While this file contains the following properties, they should not be modified for normal use:

- log4j.additivity.Workstation=false
- log4j.appender.console=com.wily.org.apache.log4j.ConsoleAppender
- log4j.appender.console.layout=com.wily.org.apache.log4j.PatternLayout
- log4j.appender.console.layout.ConversionPattern=%d{M/dd/yy hh:mm:ss a z} [%-3p] [%c] %m%n
- log4j.appender.logfile=com.wily.org.apache.log4j.RollingFileAppender
- log4j.appender.logfile.layout=com.wily.org.apache.log4j.PatternLayout
- log4j.appender.logfile.layout.ConversionPattern=%d{M/dd/yy hh:mm:ss a z} [%-3p] [%c] %m%n
- log4j.appender.logfile.MaxBackupIndex=4
- log4j.appender.logfile.MaxFileSize=200MB

CA Technologies Prerelease Documentation. Authorized Use Only.

Introscope Workstation.lax

A variety of Workstation behaviors are controlled by properties defined in this file.

lax.application.name

Important! It is important not to edit the default name of this executable.

Default

Introscope Workstation.exe

lax.class.path

The Java classpath necessary to run this application. It can be separated by colons (Mac OS/UNIX) or semicolons (Windows).

Default

launcher.jar;lax.jar

lax.command.line.args

What command line arguments will be passed to the main method. Can use this property to supply Workstation with login values, bypassing the Login screen when starting Workstation.

Default

$CMD_LINE_ARGUMENTS$

lax.dir

The path to the directory holding LaunchAnywhere's native launcher.

Default

<blank>

lax.main.class

The class that contains the main method for the application.

Default

org.eclipse.core.launcher.Main

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.main.method

The method in the main class that will be invoked.

Default

main

lax.nl.current.vm

VM to use the next time the Workstation is started. Can be set to any installed JRE version. See the *CA APM Installation and Upgrade Guide* for JVM version requirements.

Default

Varies by operating system.

lax.nl.java.launcher.main.class

Notes

It is important that you do not modify this property because it is the main class of LaunchAnywhere's java launcher.

Default

com.zerog.lax.LAX

lax.nl.java.launcher.main.method

Notes

It is important that you do not modify this property because it is the main method of LaunchAnywhere's java launcher.

Default

main

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.nl.java.option.additional

Other option settings for JVM.

Initial minimum and maximum heap sizes.

Default

-client -Xms64m -Xmx256m -Dsun.java2d.noddraw=true

lax.nl.message.vm.not.loaded

The message displayed in a user dialog if no virtual machine from the *lax.nl.valid.vm.list* can be found.

Notes

This property is internal to the InstallAnywhere launcher. Introscope has separate Java requirements. For more information, see the *System Requirements* in the *CA APM Installation and Upgrade Guide*.

Default

LaunchAnywhere either could not find a Java VM, or the Java VM on this system is too old. LaunchAnywhere requires a Java 2 VM in order to launch Introscope.

lax.nl.valid.vm.list

A string containing one or more of [ALL JDK JRE J1 J2 JRE_J1 JDK_J1 JRE_J2 JDK_J2 MSJ] delimited by spaces or commas. If the native launcher cannot find the current virtual machine, it will search for ones defined in this list.

Default

J2

lax.nl.win32.microsoftvm.min.version

The minimum version with which this version of Microsoft's virtual machine launcher will run.

Default

2750

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.root.install.dir

The path to the *installdir* magic folder.

Default

<blank>

lax.stderr.redirect

Standard Error Output.

Possible values:

- <blank>—no output
- *console*—sends output to a console window
- *<path_to_a_file>*—saves output to a sepcified file.

Default

<blank>

lax.stdin.redirect

Standard Input.

Possible values:

- <blank>—no input
- *console*—reads from a console window
- *<path_to_a_file>*—reads input from a sepcified file

Default

<blank>

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.stdout.redirect

Standard Output.

Possible values:

- <blank>—no output
- *console*—sends output to a console window
- *<path_to_a_file>*—saves output to a sepcified file.

Default

<blank>

lax.user.dir

This property will cause the native launcher to not alter the platform default behavior for setting the user directory.

To override this you may set this property to a relative or absolute path. Relative paths are relative to the launcher.

Default

<blank>

lax.version

The version of LaunchAnywhere that created this properties file.

Default

9.0

IntroscopeWebView.properties

A variety of WebView behaviors are controlled by properties defined in this file.

CA Technologies Prerelease Documentation. Authorized Use Only

introscope.webview.jetty.configurationFile

The path to a Jetty 6.0.2 XML configuration file. Either an absolute path, or a path relative to the config directory. This file creates an SSL listener on port 8443 for HTTPS traffic.

For advanced Jetty configuration, uncomment this property and modify the settings in the Jetty configuration file.

If uncommented, the http listener specified by the introscope.webview.tcp.port property is not created, but if desired it can be configured in the Jetty configuration file.

Default

webview-jetty-config.xml

introscope.webview.tcp.ipaddress

A WebView HTTP server setting property for binding WebView to a specific local IP address. When it is not configured, WebView will accept incoming connections on all local addresses (the wildcard address).

Default

(set to a valid IP address on WebView machine)

introscope.webview.tcp.port

The port on which WebView listens for HTTP traffic.

Default

8080

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.webview.context.path

The application context path for WebView. The context path is appended to the fully qualified hostname and port number of the WebView server, to form the WebView root—the URL for accessing WebView. This feature allows the user to define an application name for WebView.

If the context path is:

/webview

Then the URL to access WebView is:

http://host:port/webview/

Default

/

Note: If you are using WebView in an SAP environment, then the default is /webview

introscope.webview.enterprisemanager.tcp.host

Enterprise Manager hostname to connect to.

Default

localhost

introscope.webview.enterprisemanager.tcp.port

Enterprise Manager port to connect to.

Default

5001

CA Technologies Prerelease Documentation. Authorized Use Only

log4j.logger.WebView

Amount of logging detail for the WebView log.

Level of detail value can be:

- *INFO*
- *VERBOSE#com.wily.util.feedback.Log4JSeverityLevel*

Destination value can be:

- *console*
- *logfile*

Default

INFO, console

log4j.logger.WebServer

Amount of logging detail for the WebServer log.

Level of detail value can be:

- *INFO*
- *VERBOSE com.wily.util.feedback.Log4JSeverityLevel*

Destination value can be:

- *console*
- *logfile*

Default

INFO, console

CA Technologies Prerelease Documentation. Authorized Use Only.

log4j.logger

The logging property for the third-party software WebView uses.

Level of detail value can be:

- *INFO*
- *VERBOSE com.wily.util.feedback.Log4JSeverityLevel*

Destination value can be:

- *console*
- *logfile*

Default

INFO, console

log4j.logger.com.wily

Amount of logging detail for the CA APM log.

Level of detail value can be:

- *INFO*
- *VERBOSE com.wily.util.feedback.Log4JSeverityLevel*

Destination value can be:

- *console*
- *logfile*

Default

INFO, console

log4j.logger.com.isomorphic

Amount of logging detail for the isomorphic log.

If logfile is specified above, the location of the log file is configured with the following property: *log4j.appender.logfile.File*. Full paths can be used if desired.

Default

Error, console

CA Technologies Prerelease Documentation. Authorized Use Only.

log4j.logger.org.apache.myfaces

Amount of logging detail for the Apache Myfaces log.

If logfile is specified above, the location of the log file is configured with the following property: *log4j.appender.logfile.File*. Full paths can be used if desired.

Default

ERROR,console

log4j.logger.org.apache.jasper

Amount of logging detail for the Apache Jasper log.

If logfile is specified above, the location of the log file is configured with the following property: *log4j.appender.logfile.File*. Full paths can be used if desired.

Default

ERROR,console

log4j.logger.org.apache.struts

Amount of logging detail for the Apache Struts log.

If logfile is specified above, the location of the log file is configured with the following property: *log4j.appender.logfile.File*. Full paths can be used if desired.

Default

ERROR,console

log4j.logger.org.mortbay

Amount of logging detail for the Mortbay log.

If logfile is specified above, the location of the log file is configured with the following property: *log4j.appender.logfile.File*. Full paths can be used if desired.

Default

ERROR,console

CA Technologies Prerelease Documentation. Authorized Use Only.

Logging Configuration Properties

The following properties should not be modified for normal use:

- log4j.additivity.WebView=false
- log4j.additivity.WebServer=false
- log4j.additivity.com.wily=false
- log4j.appender.console=com.wily.org.apache.log4j.ConsoleAppender
- log4j.appender.console.layout=com.wily.org.apache.log4j.PatternLayout
- log4j.appender.console.layout.ConversionPattern=%d{M/dd/yy hh:mm:ss a z} [%-3p] [%c] %m%n
- log4j.appender.logfile=com.wily.org.apache.log4j.RollingFileAppender
- log4j.appender.logfile.layout=com.wily.org.apache.log4j.PatternLayout
- log4j.appender.logfile.layout.ConversionPattern=%d{M/dd/yy hh:mm:ss a z} [%-3p] [%c] %m%n
- log4j.appender.logfile.MaxBackupIndex=4
- log4j.appender.logfile.MaxFileSize=200MB

Note

If the value are set to true, graphs will be antialiased (better looking). For example, introscope.workstation.graph.antialiased=true

Default

The default is dependent on the property.

introscope.workstation.graph.antialiased

Graphs will be antialiased—better looking—if you keep the default true value.

Default

true

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.explorer.refresh.interval

Frequency (in seconds) at which the Investigator tree refreshes.

Default

15

introscope.webview.svg.text.as.images

Set to true to enable WebView to display non-Latin character sets, such as Kanji.

Default

false

log4j.appender.console

This property is added to the WebView properties file, to write messages to the console.

Default

logs/IntroscopeWebView.log

log4j.appender.console.encoding

This property is added to the WebView properties file, and set to the value "IBM-1047" to enable WebView to write messages in EBCDIC format when running under z/OS.

Default

<blank>

log4j.appender.logfile

This property is added to the WebView properties file, to append messages to the log file.

Default

com.wily.org.apache.log4j.RollingFileAppender

CA Technologies Prerelease Documentation. Authorized Use Only.

log4j.appender.logfile.encoding

This property is added to the WebView properties file, and set to the value "IBM-1047" to enable WebView to write messages in EBCDIC format when running under z/OS.

Default

<blank>

log4j.appender.logfile.File

Name and location, relative to the Enterprise Manager home directory, of the WebView log file.

Default

logs/IntroscopeWebView.log

Introscope WebView.lax

The runtime behavior of WebView is controlled by properties defined in this file.

lax.application.name

LaunchAnywhere Executable Properties File - Macrovision Corp.

Default

<blank>

lax.class.path

The Java classpath necessary to run this application.

Default

lib/WebServices.jar:lib/IntroscopeClient.jar:lib/Workstation.jar:lib/WebViewAPI.jar:lib/WebViewServer.jar:lib/IntroscopeServices.jar

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.command.line.args

Command line arguments to pass to the main method. Can use this property to supply WebView with login values, bypassing the Login screen when starting WebView.

Default

$CMD_LINE_ARGUMENTS$

lax.dir

The path to the directory where the LaunchAnywhere's native launcher resides.

Default

<blank>

lax.main.class

The class that contains the main method for the application.

Default

org.eclipse.core.launcher.Main

lax.main.method

The method in the main class that will be invoked.

Default

main

lax.nl.current.vm

VM to use the next time the Workstation is started. Can be set to any installed JDK or JRE version. See the *CA APM Installation and Upgrade Guide* for JVM version requirements.

Default

Varies by operating system

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.nl.java.launcher.main.class

The main class of LaunchAnywhere's java launcher.

Notes

This property should not be adjusted.

Default

com.zerog.lax.LAX

lax.nl.java.launcher.main.method

The main method of LaunchAnywhere's java launcher.

Notes

This property should not be adjusted.

Default

main

lax.nl.java.option.additional

Optional settings for JVM, such as heap size.

Default

-Xms128m -Xmx512m -Djava.awt.headless=false -Dsun.java2d.noddraw=true

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.nl.message.vm.not.loaded

The message displayed in a user dialog if no VM from the *lax.nl.valid.vm.list* can be found.

Notes

This property is internal to the InstallAnywhere launcher. Introscope has separate Java requirements. For more information, consult the *System Requirements* section of the *CA APM Installation and Upgrade Guide*.

Default

n/a

lax.nl.valid.vm.list

A string containing one or more of [ALL JDK JRE J1 J2 JRE_J1 JDK_J1 JRE_J2 JDK_J2 MSJ] delimited by spaces or commas. If the native launcher cannot find the current virtual machines, it will search for ones defined in this list.

Default

J2

lax.nl.win32.microsoftvm.min.version

The minimum version of Microsoft's VM this application will run against.

Default

2750

lax.root.installation.dir

The path to the installdir magic folder.

Default

<blank>

CA Technologies Prerelease Documentation Authorized Use Only

lax.stderr.redirect

Standard Output.

Possible values:

- <blank>—no output
- *console*—sends output to a console window
- *<path_to_a_file>*—saves output to a sepcified file.

Default

console

lax.stdin.redirect

Standard Input.

Possible values:

- <blank>—no input
- *console*—reads from a console window
- *<path_to_a_file>*—reads input from a sepcified file

Default

console

lax.stdout.redirect

Standard Output.

Possible values:

- <blank>—no output
- *console*—sends output to a console window
- *<path_to_a_file>*—saves output to a sepcified file.

Default

console

CA Technologies Prerelease Documentation. Authorized Use Only.

lax.user.dir

This is left blank because the property will cause the native launcher to not alter the platform default behavior for setting the user dir. To override this behavior, set this property to a relative or absolute path. Relative paths are relative to the launcher.

Default

<blank>

lax.version

The version of LaunchAnywhere that created this properties file.

Default

9.0

CA Technologies Prerelease Documentation. Authorized Use Only.

Appendix B: JDBC API

Introscope includes an API that enables you to use a data analysis tool such as Excel or SAS to view Introscope data. Using the JDBC interface, you can extract metric and event data for historical ranges. The data is returned in a table with a single column containing an XML string that describes the table metadata, and contains the table of results.

This section contains the following topics:

Configuring the JDBC API (see page 507)
Using the JDBC API (see page 507)

Configuring the JDBC API

The file *IntroscopeJDBC.jar* contains the JDBC driver class and all classes needed by the driver. It is the only file needed on the classpath of the client application.

The driver class is *com.wily.introscope.jdbc.IntroscopeDriver*.

The URL format is (note that case is important):

```
jdbc:introscope:net//{username}{:password}@{host}:{em port}
```

For example:

```
jdbc:introscope:net//Admin:@localhost:5001
```

Using the JDBC API

The syntax for the Introscope JDBC API supports these virtual views:

- metric_data
- trace_headers
- traces

CA Technologies Pre-release Documentation. Authorized Use Only.

Metric queries

When you perform a metric query use one of the following date formats supported by the Introscope JDBC driver:

- *yy/mm/dd hh:mm:ss*
- *EEE MMM dd HH:mm:ss zzz yyyy*

For example, a metric query may look like:

```
select * from metric_data where agent={regex} and metric={regex} and timestamp
between {starttime} and {endtime}
```

or:

```
select * from metric_data where agent='.*Custom Metric Host.*' and metric='.*'
and timestamp between '01/01/01 00:00:00' and '01/01/07 00:00:00'
```

These qualifiers support the Enterprise Manager's metric aggregation features:

This qualifier	Does this
sort by value {a|d}	Flags the Enterprise Manager to do an ascending or descending *N* sort on the metric values *for each timeslice*. Use this qualifier with the *maxmatches* flag to specify the metric limit for the Top N match.
maxmatches	Flags the Enterprise Manager to only return data for up to this number of metrics. If *maxmatches=2*, only data for two metrics are returned. This is different from the sql *limit* keyword, which limits the number of rows returned. Our *maxmatches=2* example might return 50 timeslices for 2 metrics, for a total of 100 rows of data.
period	Flags the Enterprise Manager to aggregate data points to the given period. For example, the Enterprise Manager usually reports 15-second data values, but can aggregate to 60 seconds if *period=60*.
aggregateall	Flags the Enterprise Manager to aggregate all metric data into one summarized timeslice. For example, you could use this qualifier to give a single Average Response Time for a metric across a one-week time period. Avoid using *aggregateall* with the *period* qualifier—they both deal with aggregation and would conflict.

These qualifiers must appear in this order in the SQL statement:

```
[sort by value {ascending|descending}] [maxmatches={max metrics}] [period={period in
seconds}] [aggregateall]
```

CA Technologies Prerelease Documentation
Authorized Use Only

Metric query examples

```
select * from metric_data where agent='.*' and metric='.*Average Response.*' and
timestamp between '01/01/01 00:00:00' and '01/01/07 00:00:00' maxmatches=10

select * from metric_data where agent='.*' and metric='.*Average Response.*' and
timestamp between '01/01/01 00:00:00' and '01/01/07 00:00:00' maxmatches=1

select * from metric_data where agent='.*' and metric='.*Average Response.*' and
timestamp between '01/01/01 00:00:00' and '01/01/07 00:00:00' maxmatches=1 period=120

// match one metric_data and aggregate all values. This returns one row of data select
* from metrics where agent='.*' and metric='.*Average Response.*' and timestamp
between '01/01/01 00:00:00' and '01/01/07 00:00:00' maxmatches=1 aggregateall

// this is a top 10 query select * from metric_data where agent='.*' and
metric='.*Average Response.*' and timestamp between '01/01/01 00:00:00' and '01/01/07
00:00:00' sort by value d maxmatches=10
```

Transaction trace queries

The basic transaction header query syntax is:

```
select * from trace_headers where timestamp between {starttime} and {endtime} and
query={lucene query string} [limit {rows}]
```

where *lucene query string* is a valid Lucene query string, as in the Workstation's historical query (see the *CA APM Workstation User Guide* for more information.)

The optional *limit* flag limits the number of rows returned by the query. If no limit is specified, the default transaction limit defined in the server is used.

The query for trace bodies is identical, but references a different view:

```
select * from traces where timestamp between {starttime} and {endtime} and
query={lucene query string} [limit {rows}]
```

Transaction trace query examples:

```
select * from trace_headers where timestamp between '01/01/01 00:00:00' and '01/01/07
00:00:00' and query='s*'

select * from trace_headers where timestamp between '01/01/01 00:00:00' and '01/01/07
00:00:00' and query='host:blinky'

select * from traces where timestamp between '01/01/01 00:00:00' and '01/01/07
00:00:00' and query='host:blinky' limit 100
```

The rows of data returned by trace queries include agent information like the domain, host, process and agent name, and the XML representation of the transaction trace.

CA Technologies Proprietary Use Document Authorized Use Only.

Appendix C: Configuring Introscope Extensions

If you have enabled the Introscope extensions when installing, you will have to configure them. For more information about installing, see the *CA APM Installation and Upgrade Guide*.

This section contains the following topics:

Configuring the SNMP Adapter (see page 511)
Configuring the Introscope alerts web services extension (see page 520)

Configuring the 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. For more general information about the SNMP Adapter, see the *CA APM Workstation User Guide*.

SNMP Adapter limitations

The ability to read Introscope-specific metrics in an SNMP Manager is a convenient way to monitor data from all your network components, from one location. Note, however, these limitations to Introscope SNMP functionality:

- Only Introscope metrics are available for viewing in the SNMP Manager.
- Other Introscope features, such as alerts and actions, are not published as part of the Introscope-specific information.
- Introscope information is read-only.
- "Set" is not supported.

SNMP Adapter components

The SNMP Adapter components are:

SNMP Collections

An SNMP collection is an Introscope element that defines which metrics will be published to a MIB file. SNMP collections are stored in a Management Module.

CA Technologies Prerelease Documentation Authorized Use Only.

MIB File (Management Information Base)

A MIB is a text file that describes the information reported by an SNMP agent. A MIB file can be loaded into an SNMP Manager, so that Manager knows what information the SNMP agent can report.

An Introscope-generated MIB file is a snapshot in time of the currently active SNMP collections, and describes the matching Introscope metrics in SNMP-compliant format.

SNMP Agent

An Introscope SNMP agent is installed in an Enterprise Manager deployment. The agent collects Introscope-specific metric data from the Enterprise Manager, converts it to SNMP-specific format, and sends it to the SNMP Manager.

Alert notifications

If an Alert is triggered, and fires an SNMP notification action, the SNMP Adapter sends an SNMP notification (trap) describing the alert to the configured SNMP manager specified in the Enterprise Manager properties file.

Alerts

The Enterprise Manager now sends SNMP notifications (traps) to configured SNMP managers when alerts are triggered, if an SNMP notification action has been configured for the alert. The notification that holds the Alert message has a value object of type *octet string*. The content is the same as for a Console Notification Action.

These SNMP Objects are automatically published in the MIB.

Notifications use fixed OIDs (unique Object IDs), described as follows, where *xxxx* is the randomly-chosen SnmpAdapter identifier (e.g. 5527):

Notification (numeric format)

Value: 1.3.6.1.4.1.10480.1.1.xxxx.5000

Notification (name format)

Value:
iso.org.dod.internet.private.enterprises.wily.products.introscope.snmpAdapterxxxx.wilyAlertNotif

Alert (numeric format)

Value: *1.3.6.1.4.1.10480.1.1.xxxx.6000*

CA Technologies Pre-release Documentation. Authorized Use Only

Alert Message (name format)

Value:
iso.org.dod.internet.private.enterprises.wily.products.introscope.snmpAdapterxxxx.wilyAlertNotifMessage

Configuring notifications and alerts

To configure notifications (traps):

- Edit *IntroscopeEnterpriseManager.properties* and set the following properties.

introscope.agent.enterprisemanager.snmp.target.port

Port of the SNMP manager

Default

162

introscope.agent.enterprisemanager.snmp.target.destination

SNMP manager's hostname or IP address

Default

localhost

introscope.agent.enterprisemanager.snmp.target.community

SNMPv2 community of the SNMP manager

Default

public

CA Technologies Prerelease Documentation. Authorized Use Only.

introscope.enterprisemanager.snmp.notification.enable

Turns notification on *(true)* or off *(false)*

Default

true

Notes

The *...snmp.notification.enable* property is supported by Hot Config, so it is not necessary to restart the Enterprise Manager after setting the property. The Enterprise Manager must be restarted before it will pick up the other properties.

You can also configure alerts and SNMP notification actions on the Workstation/Enterprise Manager. For more information, see the *CA APM Workstation User Guide* and various topics in this guide.

Configuring SNMP functionality

To prepare Introscope-specific metrics to be viewable in an SNMP framework:

- enable SNMP in the Enterprise Manager properties file—see Enabling SNMP by configuring property settings (see page 515)
- create an SNMP Collection—see the *CA APM Workstation User Guide* for more information
- publish an MIB—see the *CA APM Workstation User Guide* for more information
- load the MIB into the SNMP Manager—see Loading the MIB file into an SNMP Framework (see page 516)

CA Technologies Pre-release Documentation. Authorized Use Only.

Enabling SNMP by configuring property settings

To enable SNMP, you configure two settings in the file *IntroscopeEnterpriseManager.properties*.

To configure the settings:

1. Shut down the Enterprise Manager if it is running.
2. Open the file *<EM_Home>/config/IntroscopeEnterpriseManager.properties*.
3. In the setting *introscope.enterprisemanager.snmp.enable*, enter a value of *true*.
4. In the setting *introscope.enterprisemanager.snmp.agent.port,* accept the default value of 161.

Note: If you are using a UNIX platform you must be logged in as root, or have root privileges, to use any port lower than 1000 (the default SNMP agent port is 161).

1. Save the changes.
2. Restart the Enterprise Manager.

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.

CA Technologies Prerelease Documentation
Authorized Use Only.

Loading the MIB file into an SNMP Framework

The SNMP Manager must have a copy of the MIB, to identify what metrics are available from the SNMP agent. After you publish a MIB, you must load the MIB file into your SNMP Manager.

You must upload the MIB file on the SNMP Manager when:

- you publish a new MIB.
- the Enterprise Manager or the managed application is moved to another host.

To use the Introscope-generated MIB in your SNMP framework, consult the documentation for your specific SNMP Framework to see how to prepare the MIB. Contact the CA Technologies Professional Services for assistance.

Uploading to HP LoadRunner

The following information was current as of publication. Contact HP for updates.

To upload an MIB file generated by Introscope into HP LoadRunner:

1. Install SNMP Adapter:
   a. Copy *SNMPAdapter.jar* to *<EM_Home>/product/enterprisemanager/plugins*.
   b. Edit the *IntroscopeEnterpriseManager.properties* file (located in *<EM_Home>/config/*) to set the following property:
   `introscope.enterprisemanager.snmp.enable=true`
2. Restart the Enterprise Manager.
3. Create an SNMP collection. See the *CA APM Workstation User Guide* for more information.
4. Publish the MIB file:
   a. If it is not open, launch the Management Module Editor in the Introscope Investigator by selecting File > New Management Module Editor.
   b. Publish the MIB file by clicking Manager > Publish MIB...
   c. Choose LoadRunner.
   d. Click Publish.
5. Copy the MIB file:
   - from: *<EM_Home>/snmp*
   - to: SiteScope: *<SiteScope>/templates.mib*
6. Restart the SiteScope Windows service.

CA Technologies Prerelease Documentation. Authorized Use Only.

7. Use the SiteScope SNMP Browser to verify that it has loaded the MIB file:
   a. Go to Start > Programs > HP SiteScope > Open HP SiteScope LR.
   b. Click Open SiteScope.
   c. Open Monitors tab.
   d. Select: SiteScope > Tools > SNMP Browser.
   e. Enter the Enterprise Manager hostname, version V2.
   f. Select the MIB file.
   g. Click Browse.
8. Add measurements:
   a. Open the LoadRunner Controller. Do Start > Programs > HP LoadRunner > Applications > Controller.
   b. Create a Manual Scenario.
   c. On the Scenario window, click the Run tab.
   d. In the Available Graphs > System Resource Graphs list, double-click SNMP.
   e. Select SNMP Graph.
   f. Click Monitors > Add Measurements.
9. In the SNMP dialog:
   a. Click Advanced.
   b. Select SiteScope and click OK.
   c. Click Add to add new monitored server machine.
   d. Set the Enterprise Manager host as Monitored Server and click OK.
   e. Click Add to add Resource Measurements.
   f. Verify that Server is set to the Enterprise Manager machine name.
   g. Change SNMP version to V2.
   h. Update Every... to higher value. If this setting is too low, HP LoadRunner may show errors.
10. Expand the tree and select desired OIDs (metrics).
11. Verify data in HP LoadRunner.

Notes:

If you use HP LoadRunner:

- HP LoadRunner SNMP Monitor is required to view the published data.
- If the MIB does not show up in SiteScope SNMP Browser, review the error logfile.

CA Technologies Prerelease Documentation. Authorized Use Only.

- Because HP LoadRunner does not allow the same name for any node in the MIB, metrics in the MIB with the same name appear with numbers appended. For example:
  `fileInputRateBytesPerSecond23`
- HP LoadRunner does not graph negative numbers.

Uploading to HP OpenView

To upload an MIB file generated by Introscope into HP OpenView:

1. Install SNMP Adapter:
   a. Copy *SNMPAdapter.jar* to *<EM_Home>/product/enterprisemanager/plugins*.
   b. Edit the *IntroscopeEnterpriseManager.properties* file (located in *<EM_Home>/config/*) to set the following property:
   `introscope.enterprisemanager.snmp.enable=true`
2. Restart the Enterprise Manager.
3. Create an SNMP collection. See the *CA APM Workstation User Guide* for more information.
4. Publish the MIB file:
   a. If it is not open, launch the Management Module Editor in the Introscope Investigator by doing File > New Management Module Editor.
   b. Publish the MIB file by clicking Manager > Publish MIB...
   c. Choose HP OpenView.
   d. Click Publish.
5. Copy the MIB file:
   - from: *<EM_Home>/snmp*
   - to: the machine where HP OpenView installed
6. On the HP OpenView machine, do Start > Programs > HP OpenView > Network Node Manager Admin > Network Node Manager
7. In the root console, go to Settings > Load.
8. Under Unload the .mibs, choose SNMP.
9. Click Load.
10. Select the Introscope MIB file.
    The file name will include the name of the Enterprise Manager that generated the MIB file.

CA Technologies Pre-release Documentation. Authorized Use Only.

11. After the upload is complete, click Close.
12. Open the MIB browser. Do Tools > SNMP.
    a. In the Name or Address field, enter the Enterprise Manager hostname.

    This will be the same hostname referred to in the name of the MIB file from Introscope.

    b. In the Community Name field, enter *public*.

    c. Expand the node:
    `iso > org > dod > private > enterprises> wily > products > introscope > snmpAdapter > hostname`

To display a metric in a graph:

1. Select a metric.
2. Click Graph.

To find a metric in the MIB:

1. Select any hostname, agentname, resource name from the MIB.
2. Click Start Query.

Converting Introscope metrics to SNMP metrics

Introscope converts both the Introscope metric type and Introscope metric name into SNMP-compliant information.

Metric type conversion

The SNMP agent converts Introscope metrics into SMIv2 metric types as shown here:

Introscope Metric Type	Generic MIB SMIv 2 Metric Type	HP OpenView or LoadRunner MIB SMIv 2 Metric Type	BMC MIB SMIv 2 Metric Type
Int	INTEGER32(-2147483648 …2147483647	INTEGER32(-2147483648… 2147483647	COUNTER32(-2147483648… 2147483647
Long	INTEGER32 (same as above)	INTEGER32 (same as above)	COUNTER32 (same as above)
String	OCTETSTRING	OCTETSTRING	OCTETSTRING

CA Technologies Pre-release Documentation. Authorized Use Only.

SNMP does not support integers greater than 32 bits. Longs are converted to 32 bit integers using this algorithm:

- if the long is greater than Integer.MAX_VALUE, set it to Integer_MAX_VALUE
- if the long is less than Integer.MIN_VALUE, set it to Integer.MIN_VALUE
- otherwise, simply cast it to an Int

Metric name conversion

The SNMP agent uses these rules to convert Introscope metric name segments into MIB identifiers:

1. Segments must contain only legal characters (A-Z|a-z|-|0-9). Any illegal characters in a segment are removed.
2. Each segment must begin with a lowercase letter. If a segment begins with an uppercase letter, it will be converted to lowercase. If a segment begins with a number, it will be prepended with the letter "a."
3. A segment cannot end with a hyphen. An ending hyphen will be removed.
4. A segment cannot contain two hyphens in a row. One hyphen will be removed.

A segment must be 64 characters or fewer. If a segment is more than 64 characters, it is truncated to 64 characters. If truncation results in identical segments, the segment is further truncated, and appended with a dash and a number—for example, (-2).

Configuring the Introscope alerts web services extension

Introscope provides the CA APM Web Services APIs so you can use third-party applications to use CA Introscope as a source for alerts and metrics. You configure the CA Introscope Alerts Web Services extension to configure the CA Introscope alerts available to your third-party applications.

The properties below do not appear in IntroscopeEnterpriseManager.properties. You must manually add them in the file to configure them.

For details about the web services files that are installed along with the Enterprise Manager, see the *CA APM Installation and Upgrade Guide*.

To configure the Introscope alerts web services extension:

1. Open the *IntroscopeEnterpriseManager.properties* file, which is located in the *<EM_Home>/config* directory.
2. Add the *introscope.alerts.extension.managementmodules.enable* property, and list the Management Modules for which you want to provide alerts. If you leave the property value empty or specify "ALL," CA Introscope provides alerts from all Management Modules.

CA Technologies Proprietary and Confidential Information. Unauthorized Use Prohibited. Copyright © CA.

3. Add this property: *introscope.alerts.extension.inventory.updatecheckintervalsecs.* This property specifies how often the alerts extension checks for alerts, in seconds. If you do not specify a value, Introscope checks every 60 seconds by default. You cannot specify a value lower than 15 seconds; if you do so, Introscope defaults to 15 seconds.

   If you add agents, remove agents, or save changes to the Management Module during the above time interval, then CA Introscope does an internal inventory check for agents and Management Modules. Introscope updates the inventory before processing the alert status changes.

4. Restart the Enterprise Manager for the property file changes to take effect.

CA Technologies Prerelease Documentation. Authorized Use Only.

Appendix D: The Default tess-default.properties File

The *tess-default.properties* file contains assorted properties that are rarely modified and that should only be modified by experienced CA CEM administrators or as directed by CA Support.

This section contains the following topics:

Modifying the tess-default.properties Property Default File (see page 523)
Setting the Perl or Python path for evidence collection (see page 524)
Setting the directory path for the Additional Defect Information link (see page 525)
Improving defect processing performance (see page 527)
Changing the time of the daily APM database purge (see page 529)
Changing the directory locations for startup, download and bad files (see page 530)
Other CA CEM Default Properties (see page 532)

Modifying the tess-default.properties Property Default File

The CA CEM tess-default.properties file is in the *<EM_Home>/config* directory.

Important! If you need to edit any of the default properties, you need to create your own version of the file; name it:
tess-customer.properties

By creating your own properties file called *tess-customer.properties*, you ensure that your properties' modifications are not lost during a subsequent CA CEM software upgrade.

WARNING! In a clustered environment, after creating or editing your *tess-customer.properties* file, you must copy this file to each Enterprise Manager.

CA Technologies Internal Use Only Authorized Use Only Documentation

Each Enterprise Manager has a *tess-default.properties* file. The properties that an Enterprise Manager reads from its local file depends on which services are running on that Enterprise Manager. For example, a property related to customer experience metrics is read from the Enterprise Manager running the TIM Collection Service, and a property related to the maximum number of user groups is read from the Manager of Managers.

Follow these steps:

1. Locate the file:

   ```
   <EM_Home>/config/tess-default.properties
   ```

2. Create your own version of the *tess-default.properties* file in the same directory, and name it *tess-customer.properties*.

3. Modify your *tess-customer.properties* file, adding only the properties you need to modify for your CA CEM environment.

   a. Edit your new properties file.

      For example, using the vi editor:

      ```
      vi tess-customer.properties
      ```

   Important! The *tess-customer.properties* file should contain *only* those properties with values that differ from the CA CEM system defaults. If you leave all the properties in your *tess-customer.properties* file, then system upgrades might not take advantage of future system performance or tuning updates to the properties.

   b. Save the *tess-customer.properties* file.

4. Copy the *tess-customer.properties* file to all the Enterprise Managers in your cluster.

5. Restart each Enterprise Manager in your cluster.

CA Technologies Proprietary Documentation. Authorized Use Only.

Setting the Perl or Python path for evidence collection

To use evidence collection (Setting business impact and collecting evidence for incidents (see page 276)), you need to have installed Python or Perl on the MOM or standalone Enterprise Manager. Edit the python.path or perl.path properties to specify the location of Python or Perl.

perl.path

The directory where you have installed Perl.

Options

A directory path to the Perl executable

Default

/usr/bin/perl

Example

C:\\perl\\bin\\perl.exe

Notes

If a backslash is used for a Windows platform, a double slash (\\) is required because a single slash (\) is interpreted as an escape sequence.

python.path

The directory where you have installed Python.

Options

A directory path to the Python executable

Default

/usr/bin/python

Example

C:\\python\\bin\\python.exe

Notes

If a backslash is used for a Windows platform, a double slash (\\) is required because a single slash (\) is interpreted as an escape sequence.

CA Technologies Prerelease Documentation. Authorized Use Only.

Setting the directory path for the Additional Defect Information link

To make the Additional Defect Information link appear on a defect details page, there must be an HTML file named with the defect number.

The defect details page can link to a custom web page that can contain additional defect information. The current defect details page displays various details about a defect, such as time, throughput, and web server address. However, there might be additional details about a defect—details more specific to your company's needs—that you would also like to add to the defect details page.

CA Technologies Professional Services can provide this defect information and deposit it in a directory for CA CEM. This additional defect information, accessed from the Additional Defect Information link on the defect details page, can help in triaging end-user problems and show contextual information related to each individual defect.

For the Additional Defect Information link to appear at the bottom of the defect details page for a given defect, the following must be true:

- There must be a file named using this convention: defect-<defectId>.htm
- The file must be placed in the directory specified in the additionalDefectDetailFileLocation property.

For example, if this file exists on the MOM *<EM_Home>/cem/tess/additionaldefectdetails/defect-600000000000000026.htm*, then a link to this page appears on the defect details page for defect number 600000000000000026.

If you have read permission for the HTML file, when you click the Additional Defect Information link, the HTLM page appears.

Important! Verify that you write a script to delete old defect-<defectId>.htm files. These files are not deleted automatically when the defects are aged out.

additionalDefectDetailFileLocation

If a file named defect-<DefectID>.htm is in this directory on the MOM, then a link to the file appears at the bottom of the defect detail page for the defect.

In a clustered environment:

- The directory must be on the MOM.
- The additionalDefectDetailFileLocation property is read from the MOM, the *tess-customer.properties* file or the *tess-default.properties* file.

Options

A directory path. This can be an absolute path or a path relative to the MOM <EM_Home>.

Default

cem/tess/additionaldefectdetails

Example

C:\\test\\evidence_for_defects (Windows example)

/test/evidence_for_defects (UNIX example)

Notes

If a backslash is used for a Windows platform, a double slash (\\) is required because a single slash (\) is interpreted as an escape sequence.

CA Technologies Preferred Documentation
Authorized Use Only.

Improving defect processing performance

Querying the database to determine the number of user groups affected by an incident can be CPU intensive. This is especially true if the incidents contain many defects. Querying for the affected user groups can delay the processing of defects. These delays can be reduced by querying less often.

By default, the affected-user-group query is performed whenever a defect is added to an incident. However, you can choose to have this query performed less often—at a specific *update interval*—instead of each time a defect is found.

If the affected-user-group query is performed at intervals, then a cache is created to hold the current affected-user-group count during the interval. In addition to setting the update interval, you set an *ageout interval* property to determine how often idle incidents are removed from the cache.

Note: The downside to performing this query less often is that the affected-user-group count can be out-of-date by the length of the update interval, or, on rare occasions, by the length of the ageout interval. However, if it results in faster performance for processing defects, the slightly late count might be a good trade off for you.

To reduce the frequency of the query for affected-user-groups:

1. (Optional) Confirm that there are delays in processing defects.

   For example, select CEM > Incident Management > Defects and show all defects. Are defects arriving when expected or are they delayed?

2. Follow the procedure in Modifying the tess-default.properties File (see page 523). Specify values for these properties:

   - defects.affectedUserGroupUpdateInterval_Secs (the update interval)
   - defects.affectedUserGroupAgeoutInterval_Mins (the ageout interval)

   For example:

   ```
   defects.affectedUserGroupUpdateInterval_Secs=180
   defects.affectedUserGroupAgeoutInterval_Mins=10
   ```

   CA Technologies recommends initial values of 180 seconds for the update interval and 10 minutes for the ageout interval. If the affected-user-group query is causing delays in defect processing, then you should see an improvement with these values.

   For details about these properties, see defects.affectedUserGroupUpdateInterval_Secs (see page 528) and defects.affectedUserGroupAgeoutInterval_Mins (see page 528).

CA Technologies Proprietary and Confidential Information
Authorized Use Only

defects.affectedUserGroupUpdateInterval_Secs

The interval, in seconds, between queries for the number of affected user groups associated with an incident.

Options

Numeric value, representing seconds:
— Minimum value, 0, specifies that the affected-user-group query is performed whenever a defect is added to an incident.
— Maximum suggested value is 3600 (1 hour).

Default

0

Example

defects.affectedUserGroupUpdateInterval_Secs=180

Notes

The affected-user-group count reported can be out-of-date by as much as the interval value, or, on rare occasions, by the length of the ageout interval.

defects.affectedUserGroupAgeoutInterval_Mins

If an incident has been idle (not updated) for longer than this value, then the incident is removed (aged out) from the affected-user-group cache.

Options

Numeric value, representing minutes:
— Minimum value, set the ageout interval to more than twice the value of the update interval. If you don't, then a warning appears in the IntroscopeEnterpriseManager log and the value is reset to twice the update interval.
— Maximum suggested value is 121 minutes (just over 2 hours).

Default

1

CA Technologies Pre-release Documentation. Authorized Use Only.

Example

defects.affectedUserGroupAgeoutInterval_Mins=10

Notes

This value should be more than twice and less than four times the value of update interval. For example, if the update interval is 180 seconds, the ageout interval should be between 7 and 11 minutes.

When an incident is removed from the cache, the affected-user-group query is run and the incident's affected-user-group count is updated. This prevents the count from being out-of-date longer than the ageout interval.

Changing the time of the daily APM database purge

By default, the daily purge of stale statistics from the database begins soon after 03:00. However, you can change the hour during which the purge begins.

stats.purgeStatsAt

Sets the hour for the daily purge of statistics data from the database. The age of the statistics purged is determined by the Data Retention Settings.

Options

integer value (0 to 23)

Default

3

Example

stats.purgeStatsAt=4

Notes

Schedule your daily backup to start after statistics have been purged from the APM database.

If you upgraded from CA CEM 4.5 and your tess-customer.properties file contains *stats.purgeStatsAndBackupDbAt*, then this property determines the time for the daily statistics purge and overwrites *stats.purgeStatsAt*.

CA Technologies Prerelease Documentation
Authorized Use Only.

Changing the directory locations for startup, download and bad files

By default, the data files sent to the Enterprise Manager from the TIMs are stored in these directories:

- cem/tess/downloadfiles (This directory contains data files from the TIM that are waiting to be processed.)
- cem/tess/startupfiles (When the Enterprise Manager starts up, old files from the TIM and old files in the downloadfiles directory are transferred to this directory for processing. These files are processed in a different thread from those in downloadfiles directory.)
- cem/tess/badfiles (If there is a problem with a file so that it cannot be processed, it is placed in this directory. The log file reports when files are moved to this directory. Files are not automatically deleted from this directory. If there are many files in this directory, delete them.)

You can change these directories by Modifying the tess-default.properties File (see page 523).

These are reasons you might want to configure these directories:

- Disk space concerns: to move the directory to a larger drive.
- If you move the TIM Collection service from one Enterprise Manager to another and if the tess-customer.properties files on both Enterprise Managers point to the same directory on the same (third) machine, then no data is lost when you move the TIM Collection service.

  By default: When the TIM Collection Service is moved to a different Enterprise Manager, any unprocessed data files remaining on the original Enterprise Manager do not get processed (unless the service gets moved back to the original Enterprise Manager).

CA Technologies Prerelease Documentation. Internal use Only.

collectors.downloadDataDirectory

The directory to which files downloaded from a TIM.

Options

A directory path. If the path is relative, it is relative <EM_Home>. If an absolute path is specified, it is used as is.

Default

cem/tess/downloadfiles

Examples

collectors.downloadDataDirectory=cem/TIMfiles/downloads

collectors.downloadDataDirectory=/opt/cemfiles/downloads

collectors.downloadDataDirectory=N:\\cem\\downloadfiles

collectors.downloadDataDirectory=\\\\server2\\c$\\cem\\tess\\downloadfiles

Notes

If a backslash is used for a Windows platform, a double slash (\\) is required because a single slash (\) is interpreted as an escape sequence.

If the directory does not exist, it is created.

collectors.startupDataDirectory

Files are moved to this directory for processing when the Enterprise Manager is started.

Options

A directory path. If the path is relative, it is relative <EM_Home>. If an absolute path is specified, it is used as is.

Default

`cem/tess/startupfiles`

Examples

```
collectors.startupDataDirectory=cem/tess/afterrestart
collectors.startupDataDirectory=/opt/cemfiles/afterrestart
collectors.startupDataDirectory=N:\\cem\\startupfiles
collectors.startupDataDirectory=\\\\server2\\c$\\cem\\tess\\startupfiles
```

Notes

If a backslash is used for a Windows platform, a double slash (\\) is required because a single slash (\) is interpreted as an escape sequence.

If the directory does not exist, it is created.

collectors.badDataDirectory

If files that contain errors that prevent them from being processed, they are moved into this directory.

Options

A directory path. If the path is relative, it is relative <EM_Home>. If an absolute path is specified, it is used as is.

CA Technologies Prerelease Documentation. Authorized Use Only.

Default

cem/tess/badfiles

Examples

collectors.badDataDirectory=cem/TIMdatafiles/badfiles

collectors.badDataDirectory=/tmp/badfiles

collectors.badDataDirectory=N:\\cem\\badfilesfromtess

collectors.badDataDirectory=\\\\server2\\c$\\cem\\tess\\badfiles

Notes

If a backslash is used for a Windows platform, a double slash (\\) is required because a single slash (\) is interpreted as an escape sequence.

If the directory does not exist, it is created.

Other CA CEM Default Properties

This topic provides information about some other <wcem> default properties:

baseline.durationInDays (see page 532)
BTStatsMetricName.backwardCompatible (see page 533)
downloadfiles.minutes.collect (see page 533)
email.defaultFromAddress (see page 534)
invalid.importexport.filename.chars (see page 535)
timTessCommunication.useSsl (see page 535)

CA Technologies Product Documentation.
Authorized Use Only.

baseline.durationInDays

The number of days of hourly data you want to use for the baseline calculation.

The default value is 28. This means that the baseline calculation is based on the hourly data collected in the last 28 days.

Usage example: You know that CA CEM has been collecting data for 10 days, but the first 3 days are not valid because of experimenting and changing transaction definitions. Therefore, you set the value to 7. You later increase this value as more days of valid data are collected.

Options

numerical value

Default

28

Example

baseline.durationInDays=7

Notes

Use a number that is multiple of 7 because, in most cases, a regular business cycle of monitored transactions is one week.

For general information about baselines, see the *CA APM Transaction Definition Guide*.

BTStatsMetricName.backwardCompatible

In CA CEM 5.0, the term "business process" changed to "business service." This property allows you to continue to use the CA CEM 4.5 dashboard metric names for customer experience metrics.

Options

true or false

If true, then the CA CEM 4.x metric name "Business Process" is used.

If false, then the name "Business Service" is used.

Default

false

Example

BTStatsMetricName.backwardCompatible=false

Notes

If you change this setting to true, you must use the Javascript calculators from CA CEM 4.5.

downloadfiles.minutes.collect

This sets a time-lag threshold for deleting old data files.

When an Enterprise Manager falls behind in processing files, then defect, automatic transaction discovery, and customer experience metrics files that are older than this are deleted (unprocessed) from TIMs and the Enterprise Manager.

For example, by default, when an Enterprise Manager restarts:

1. Files in cem/tess/startupfiles and cem/tess/downloadfiles that are more than a day old are deleted. (Those remaining in cem/tess/downloadfiles are moved to cem/tess/startupfiles.)
2. Files on TIMs that are older than a day are deleted.

CA Technologies Prerelease Documentation. Authorized Use Only.

Options

positive integer (minutes)

If 0, then all existing defect, automatic transaction discovery, and customer experience metrics files are deleted from the Enterprise Manager and TIMs when the Enterprise Manager restarts.

If -1, then all files are kept. Use this with caution and make sure that the Enterprise Manager does not fall behind in processing the data.

Default

1440

(24 hours)

Example

downloadfiles.minutes.collect=720

Notes

CA CEM events are generated when old data files are deleted from TIMs or the Enterprise Manager.

If defect files are deleted by the Enterprise Manager, the defect count (for example, on the Incident Management > Defects page) won't be the same as the number of defects reported on reports that plot statistics.

email.defaultFromAddress

If Embedded Entitlements Manager is not used for authentication or if the user's email address is not specified in Embedded Entitlements Manager, then this address is used for testing the SMTP server settings and possibly as the From address for emailed reports.

Options

A valid email address

Default

support@catechnologies.com

Example

email.defaultFromAddress=my.email@company.com

Notes

If used, this property appears in the From Address and To fields of the System > Email Settings page.

CA Technologies Prerelease Documentation. Authorized Use Only

invalid.importexport.filename.chars

Specifies which special characters can be used in exported business service names.

Options

Special characters that are not allowed in Linux file names.

Default

/ \\ : * ? < > & ; |

Example

invalid.importexport.filename.chars=/ \\ : * ? < > & ; |

Notes

When business services are exported, the name of the export file is based on the name of the business service. If the name contains characters that cannot be used in file names, then the export fails.

To prevent this failure, ensure that the characters are in this property list. When the export file is created, the characters are replaced by underscores. For example, when a business service called my&BP is exported, the export file name is BTExport_my_BP_<datestamp>.zip.

timTessCommunication.useSsl

Determines if Enterprise Manager/TIM communication is by HTTP or HTTPS.

Options

0 or 1

If 0, then Enterprise Manager and TIM communicate via HTTP.

If 1, then Enterprise Manager and TIM communicate via HTTPS.

Default

0

Example

timTessCommunication.useSsl=1

Notes

This property does not exist in the tess-default.properties file. You need to add it to your *tess-customer.properties* file.

CA Technologies Prerelease Documentation. Authorized Use Only

Index

A

access policies, for business services • 283
ACTION, argument for evidence collection script • 278
actions
  commands • 155
activate actions commands • 156
activate alerts command • 158, 160
activate calculators command • 161
activate dashboards command • 163
activate management modules command • 154
activate report templates command • 166
activate snmp collections command • 170
active scripting, required for CA CEM • 233
affected user group, update interval • 527
agent properties
  reloading CA Introscope • 367, 371
  verifying Introscope • 366, 370
agents
  CA CEM, in SiteMinder • 389
  defining for SiteMinder on CA CEM • 390
  TIM, in SiteMinder • 384
age-out rules for incidents
  configuring • 279
aggregation • 319
aggregation statistics, setting defaults • 262
aggregation, defects to incidents • 270
aggregation, effects on reports • 322
alerts
  commands • 157, 159
antivirus software, running on EMs • 260
Apache web servers, • 334
APM database • 63
  backup • 64
  exporting configuration • 69
  importing configuration • 69
  restore • 64
appliance health • 314
arguments for evidence collection scripts • 278
automatic transaction discovery
  selecting TIMs for • 257
AutoProbe, configuring tracing for • 365

B

backlogs in processing • 260
backup, APM database • 64
badfiles directory • 530
baseline, setting for defect type conditions • 532
bootstrap
  Management Module • 422
bulk editing
  user settings • 301
Business application name is not valid event • 308
business impact
  calculate • 268
  configuring • 268
  configuring default levels • 276
  incident cost • 268
  settings • 270
  versus incident cost • 268
  weighting levels • 270
business service
  access policies for • 283
  modifying with bulk editing • 301

C

CA EEM • 283
CA Service Desk
  about • 393
  configuring a request area • 397
  configuring the plug-in • 394
  integrating with CA CEM • 393
  integration overview • 393
  terminology • 400
  troubleshooting • 398
  using with CA CEM • 397
calculators
  commands • 161
capacity protection for user groups • 293
capture comprehensive defect details • 262
client IP addresses
  client IP addresses, basing on an HTTP header • 296
CLIENTIP, argument for evidence collection script • 278
Clock Skew Tolerance
  setting for Enterprise Manager • 213

CA Technologies Prerelease Documentation. Authorized Use Only

clustered environment
distributing Enterprise Manager services • 255
tess-customer.properties warning • 523
CLW
see Command Line Workstation • 141
collection
about SNMP • 511
Command Line Workstation
about • 141
command reference • 145
command syntax • 143
Enterprise Manager commands • 145
prerequisites • 141
running • 142
Command Line Workstation commands
actions • 155
activate actions • 156
activate alerts • 158
activate calculators • 161
activate dashboards • 163
activate management modules • 154
activate report templates • 166
activate snmp collections • 170
alerts • 157, 159
calculator • 161
copy management module • 155
dashboard • 163
deactivate actions • 157
deactivate alerts • 159, 160
deactivate calculators • 162
deactivate dashboards • 164
deactivate management modules • 154
deactivate report templates • 167
deactivate snmp collections • 171
delete actions • 157
delete alert • 159
delete alert downtime schedules • 161
delete calculators • 162
delete dashboards • 164
delete management module • 155
delete report templates • 167
delete snmp collections • 171
generate report • 168
list actions • 156
list agents • 147
list alert downtime schedules • 159
list alerts • 157
list calculators • 161
list dashboards • 163
list management modules • 153
list metric groups • 165
list report templates • 166
list snmp collections • 170
management modules • 153
metric groupings • 165
persistent data N previous minutes • 146
persistent data time range • 146
query alert downtime schedules • 159
query state of actions • 157
query state of alerts • 159
query state of calculators • 162
query state of dashboards • 164
query state of management modules • 154
query state of report templates • 167
query state of snmp collections • 171
readonly management modules • 155
rename action • 157
rename alert • 159
rename alert downtime schedules • 161
rename calculator • 162
rename dashboard • 164
rename report template • 167
rename snmp collection • 170, 171
report templates • 166
send config file • 148
shutdown • 145
SNMP collections • 169
specifying pathname • 147
transaction trace • 172
turn off agents • 148
turn on agents • 147
configuration file, domain • See domain configuration file
configuration settings
Management Module settings • 422
configuring the SNMP Adapter • 511
contacting technical support • 5
copy management module command • 155
correlational SLAs • 302
deleting • 305
modifying • 304
custom agents
creating for SiteMinder • 389
for each TIM, on CA SiteMinder • 384
type for CA CEM on CA SiteMinder • 389
customer support, contacting • 5

CA Technologies Prerelease Doc Information
Authorized Use Only

D

daily statistics not aggregated • 260
dashboards
  commands • 163
data files, locations • 530
data retention
  defects • 274
  effects on reports • 322
  incidents, pending • 279
  settings • 262
  time of day to purge statistics • 529
database • See APM database
Database Cleanup Service has started event • 308
Database Cleanup Service has stopped event • 308
DB Cleanup service • 252
  stopped • 260
DB2, z/OS • 41
deactivate actions command • 157, 171
deactivate alerts command • 159
deactivate calculators command • 162
deactivate dashboards command • 164
deactivate management modules command • 154
deactivate report templates command • 167
Default User Settings (SLAs) • 262
defect count discrepancy, reasons for • 533
DEFECTID, argument for evidence collection script • 278
defects
  aggregation • 270
  bulk editing conditions for • 301
  processing delays • 527
  report • 270
Defects dropped event • 308
delete actions command • 157
delete alert command • 159
delete alert downtime schedules command • 161
delete calculators command • 162
delete dashboards command • 164
delete management module command • 155
delete report templates command • 167
delete snmp collections command • 171
deregistering Enterprise Manager • 35
documentation • 27
  Introscope • 30
domain
  defining • 262
  name • 262
domain configuration file • 288
Domain Configuration History table • 289
domains.xml • 110
Down, EM Running State • 260
downloadfiles directory • 530
Duplicate session ids detected event • 308

E

EEM • See CA EEM
email
  configuring SMTP • 281
  reports • 323
  sending for incidents • 280
  settings, defining • 281
Embedded Entitlements Manager • See CA EEM
Enterprise Manager
  automatic shutdown • 50
  automatic Windows Service • 33
  changing IP address • 258
  configuring JVM • 47
  configuring Windows Service • 34
  configuring z/OS with • 39
  deregistering Windows Service • 35
  missing data from TIMs • 260
  modifying Windows Services • 35
  removing from cluster • 259
  restarting • 52
  running control scirpts on UNIX • 35
  setting Clock Skew Tolerance • 213
  shutdown command • 145
  starting OS/400 • 37
  starting UNIX • 35
  starting Windows • 33
  stopping • 50
  stopping and starting Windows Services • 35
  troubleshooting crashes • 52
  UNIX nohup mode on • 37
  Windows Service prerequisites • 33
Enterprise Manager cannot reach TIM event • 308
Enterprise Manager deleted some data files event • 308
Enterprise Manager deleted some TIM data • 308
Enterprise Manager restarted event • 308
Enterprise Manager services
  Collector appears down • 260
  distributing • 255
  modifying configuration • 258
  removing unused Enterprise Managers • 259
  tasks required for • 252

CA Technologies Prerelease Documentation. Authorized Use Only

typical distributions • 254
evidence collection
Additional Defect Information • 525
configuring severity levels • 276
script arguments • 278
scripts • 276
exporting
APM database configuration • 69
extension files
where to deploy files in a cluster • 104
extract persistent data command
previous N minutes • 146
time range • 146

F

feedback, providing on documentation • 27
filtering, TIM traffic with Napatech adapter • 344
firewall
connecting components behind • 58

G

generate report command • 168
grouping users
behind proxy servers • 296
by network address • 262, 295
by request attributes • 294
manually • 298
groups, user • See user groups

H

health, TIM machine • 314
HP LoadRunner • 515, 516
HP OpenView • 514, 518
HTTP analyzer plug-in error event • 308
HTTP communications from browser to Enterprise Manager • 245
HTTP request headers, grouping users by • 294

I

Ignore Applications in User Recognition • 262
impact leaders report • 270
impact level
configuring severity • 276
defining • 262
inheriting from the domain • 299
transactions • 262
user group • 299
importing
APM database configuration • 69
incident cost
amount per user per minute • 268
calculate • 268
configuring • 268
versus business impact • 268
incidents
closing • 270
configuring age-out rules • 279
report • 268
rules for generating • 275
sending email for • 280
states of • 270
integrating
CA Introscope with CA CEM • 349
CA Service Desk with CA CEM • 393
CA SiteMinder with CA CEM • 387
Internet Protocol version 6 • See IPv6
IP addresses
editing Enterprise Managers • 258
TIMs • 234
IP subnet
add users to a group via • 299
create user groups by • 262
iPlanet web servers, SSL private keys • 335
IPv6 • 245
and IPv4 (dual stack) • 247
IT value, defining • 262
iv-remote-address header, for identifying client IP • 296

J

Java agents, configuring in CA Introscope • 365
JDBC API
configuring • 507
metric queries • 508
transaction trace queries • 509
JDBC, z/OS • 42
JVM, configuring Enterprise Manager • 47

L

lax files
Introscope WebView.lax • 500
Introscope Workstation.lax • 487
Introsocpe Enterprise Manager.lax • 477
Limit exceeded adding new components event • 308
list actions command • 156
list agents command • 147

CA Technologies Prerelease Documentation. Authorized Use Only.

list alert downtime schedules command • 159
list alerts command • 157
list calculators command • 161
list dashboards command • 163
list management modules command • 153
list metric groups command • 165
list report templates command • 166
list snmp collections command • 170
LoadRunner • 515, 516
log files
  Napatech adapter driver • 342
  Napatech adapter driver, viewing • 345
  Napatech adapter messages in TIM log • 347
  TIM Napatech hardware filter configuration • 345
Log4J
  Workstation settings • 211
Logins dropped event • 308
logo, using on PDF reports • 323

M

machine health, TIMs • 314
Management Module Bootstrapping
  configuring settings • 422
management modules
  commands • 153
Maximum Data Point Retrieval Property
  setting • 215
  setting Workstation property • 215
metric groupings
  commands • 165
metrics
  converting • 519
  name conversion • 520
MIB
  MIB file • 515
  MIB, loading into an SNMP framework • 516
  multibyte • 212
mirrored ports • 241
  versus network taps • 240
monitors
  enabling • 287
  synchronizing • 288
multiple port network cards • 241

N

Napatech adapter
  driver log • 342
  enabling • 343
  ethernet port behavior • 348
  filtering with • 344
  filters, removing • 345
  installing • 342
  not getting traffic • 347, 348
  part number • 342
  removing • 345
  serial number, finding • 348
Netscape web servers, SSL private keys • 335
network
  alternative configurations • 239
  monitoring traffic on • 240
  TIMs placement on • 239
Network Allies, for Napatech adapter • 342
network cards
  multiple port • 241
  napatech adapter • 341
network taps
  using • 240
  versus mirrored ports • 240
new users group, managing • 298
NICs, multiple port • 241
No network data event • 308
nohup
  UNIX mode on Enterprise Manager • 37
non-identifying components
  maximum for automatic discovery property • 457

O

OpenSSL-based web servers, private keys • 334
OS/400
  command line with Enterprise Manager • 37
  control script with Enterprise Manager • 37
  SSL • 59
  start script with Enterprise Manager • 37
  starting Enterprise Manager • 37
OUTPUT, argument for evidence collection script • 278

P

path parameters
  setting delimiter • 262
pathname command
  specifying for authentication • 147
PDF reports • 323
Perl scripts for evidence collection • 276
  arguments for • 278

BMC Technologies Prerelease Documentation. Authorized Use Only.

path to executable • 524
plug-in
CA Service Desk • 394
CA SiteMinder • 388
policy servers
defining for CA SiteMinder • 389
defining on CA CEM • 389
ports, mirrored • 241
PowerPacks
in clustered environments • 104
printing reports • 323
private keys, SSL • See SSL
processing backlogs • 260
Professional Services, additional defect information link • 525
properties files
about (reference) • 403
Introscope Enterprise Manager.lax • 477
Introscope WebView.lax • 500
Introscope Workstation.lax • 487
IntroscopeEnterpriseManager.properties • 403
IntroscopeWebView.properties • 492
IntroscopeWorkstation.properties • 483
proxy server
identifying users behind • 296
ProxyForwardHeader, TIM setting • 296
Python scripts for evidence collection • 276
arguments for • 278
path to executable • 524

Q

query alert downtime schedules command • 160
query state of alerts command • 159
query state of calculators command • 162
query state of dashboards command • 164
query state of management modules command • 154
query state of report templates command • 167
query state of snmp collections command • 171

R

readonly management modules command • 155
redirect output
Workstation • 211
reloading agent properties • 367, 371
rename calculator command • 162
rename action command • 157
rename alert command • 159
rename alert downtime schedules command • 161
rename dashboard command • 164
rename report template command • 167
rename snmp collection command • 171
report templates
commands • 166
reports
aggregation, about • 319
defects • 270
delivery options • 320
formats • 321
impact leaders • 270
incident • 268
methods for capturing • 319
scheduling • 325
securing data • 320
settings, email • 281
time frames • 322
request area, configuring with CA Service Desk • 397
restarting Enterprise Manager • 52
restore, APM database • 64
routers, mirrored port on • 241
rules
age-out for incidents • 279
generating incidents • 275

S

scheduling reports • 325
scheduling tasks
using Command Line Workstation • 141
scripts for evidence collection • See Python scripts for evidence collection, See Perl scripts for evidence collection
security groups, for business service access • 283
and report data • 320
send config file command • 148
serial number, for Napatech adapter • 348
SERVER, argument for evidence collection script • 278
sessions, transaction trace • 374
SEVERITY, argument for evidence collection script • 278
shell commands
Enterprise Manager • 53
shutdown, automatic Enterprise Manager • 51
Sigma (Six)
defining for transaction SLAs • 262
user group SLAs • 299

CA Technologies Prerelease Documentation. Authorized Use Only

SiteMinder
checklist for integration with CA CEM • 387
configuring with CA CEM • 388
creating agents • 389
defining agents on CA CEM • 390
defining policy servers • 389
enabling • 390
integrating with CA CEM • 387
troubleshooting integration • 391
SLAs • See correlational SLAs
by user group, by business transaction • 302
correlational • 302
defining defaults • 262
inheriting from the domain • 299
transaction • 262
transaction success rate • 262
user groups • 299
SmartStor • 77
requirements • 79
SMTP, configuring for email • 281
SNMP Adapter
configuring • 511
converting Introscope metric names • 519
converting Introscope metric types • 519
loading the MIB • 516
MIB • 512
SNMP Agent • 512
SNMP Collections • 421, 511
commands • 169
SPAN ports • See mirrored ports
special characters, in SmartStor tools commands • 84
SSL • 333
Enterprise Manager on OS/400 • 59
private keys • 333
startupfiles directory • 530
states of incidents • 270
statistics
aggregation, setting defaults • 262
data retention • 262
time of day to purge • 529
Stats Aggregation service • 252
stopped • 260
Stats Aggregation Service has started event • 308
Stats Aggregation Service has stopped event • 308
stats.purgeStatsAt, TESS properties • 529
subnet mask
for defining user groups • 262, 295
for troubleshooting defects by IP subnet • 262
success rates
correlational SLAs • 303
for transaction SLAs, defining • 262
for user group SLAs • 299
SunOne web servers, SSL private keys • 335
support, contacting • 5
switches, mirrored port on • 241
synchronization, view changes since last • 289
synchronizing TIMs and agents • 288

T

taps, network • 240
technical support, contacting • 5
term mapping, CA Service Desk CA CEM • 401
terminal window
commands, Enterprise Manager • 53
stopping Enterprise Manager from • 51
TessCollectorIpAddr • 258
tess-customer.properties file • 523
tess-default.properties file • 523
editing • 523
TessIpAddr • 258
testing email connections • 281
thread dumps, properties • 439
TIM Appliance Health report • 314
TIM Collection Service • 252
restarting after shutdown • 52
stopped • 260
TIM Collection Service has started event • 308
TIM Collection Service has stopped event • 308
TIM deleted some of its data files event • 308
TIM restarted event • 308
TIM Settings page • 296
TIMs
appliance health • 314
between the aggregation switch and the network load balancer • 244
between the firewall and the aggregation switch • 244
between the Internet and the router • 243
between the network load balancer and the Web server • 245
between the router and the firewall • 243
data files on Collectors • 530
inserting in the network • 241
IP addresses • 234
network placement • 239

CA Technologies Prerelease Documentation. Authorized Use Only

not receiving traffic through Napatech adapter • 347, 348
preferred insertion in the network • 244
selecting for automatic transaction discovery • 257
synchronizing • 288
TIMs, enabling monitors • 287
using with Napatech adapter • 341
tracing,
.NET • 369
Java • 365
TRANDEFID, argument for evidence collection script • 278
TRANID, argument for evidence collection script • 278
transaction change control • 289
Transaction definitions dropped event • 308
Transaction Discovery Collection service • 257
Transaction Event Database
settings • 80, 419
transaction settings, configuring defaults • 262
transaction times
SLAs, defining • 262
user groups • 299
transaction trace • 172
commands examples • 176
commands output file • 173
commands trace mode • 174
queries from transaction event database • 177
session • 374
sessions sample scripts for all agents • 186
Transaction Trace Clamp • 175
transactions
impact level, defining • 262
modifying with bulk editing • 301
SLAs, defining • 262
Troubleshoot Defects by IP Subnet • 262
troubleshooting
CA Introscope CA CEM integration • 391
CA Service Desk CA CEM integration • 398
defect count discrepancy • 533
defect processing • 527
Enterprise Manager crashes • 52
Enterprise Manager Services • 260
logging in to CEM console • 233
Napatech adapter in TIMs • 347
SiteMinder CA CEM integration • 391
TIM health • 314
WatchDog.sh on z/OS • 202
turn off agents command • 148
turn on agents command • 147

U

UNIX
running Enterprise Manager control scripts • 36
starting Enterprise Manager • 35
Unknown, EM Running State • 260
User group limit reached event • 308
user groups • 291
add user to group via IP address • 299
assigning users to groups by request attribute • 294
assigning users to groups by subnet • 295
assigning users to new users group • 298
bulk editing • 301
configuring with correlational SLAs • 302
create by IP subnet • 262
creating for users behind proxy servers • 296
creating subnet user groups • 295
deleting • 305
impact level, defining • 299
managing • 291
manually assigning users to groups • 298
maximum • 293
modifying • 299
SLAs, defining • 299
users
deactivating or deleting • 306
managing • 291
reactivating • 306
settings, modifying with bulk editing • 301
users, when to deactivate • 262

V

value settings, defining • 262
verbose mode
running Workstation in • 211
virtual memory
z/OS • 40

W

web servers, SSL private keys • 334
WebView
settings • 422
weight, business impact levels • 270
Windows Service
deregistering Enterprise Manager • 35

CA Technologies Prerelease Documentation. Authorized Use Only.

modifying Enterprise Manager • 35
prerequisites for Enterprise Manager • 34
running Enterprise Manager • 34
stopping and starting Enterprise Manager • 35
Windows Service, configuring Enterprise Manager • 34
Workstation
command line functions • 141
connecting to Enterprise Manager behind firewall • 58
stopping Enterprise Manager from • 51

X

X-Forwarded-For header, for identifying client IP address • 296

Z

z/OS
base structure for Java • 39
configuring as batch job • 42
customizing Enterprise Manager for • 43
customizing PROC • 43
customizing UNIX shell script • 43
DB2 • 41
Enterprise Manager installation overview • 42
executing PROC • 43
JDBC • 42
partitioned dataset • 43
running Enterprise Manager as a started task • 45
running Enterprise Manager on • 45
security and authorization • 40
software prerequisites for Java • 39
starting Enterprise Manager automatically • 46
verifying Enterprise Manager • 46
virtual memory • 40
WatchDog.sh, running on • 202
WILYPROC with • 43

CA Technologies Prerelease Documentation. Authorized Use Only.