

# CA NSM

## Administration Guide

r11.2 SP2



This documentation and any related computer software help programs (hereinafter referred to as the "Documentation") are for your informational purposes only and are 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 used or disclosed by you except as may be permitted in 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 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 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 THE END USER 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 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 © 2010 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

## CA Product References

This document references the following CA components and products:

- CA 7<sup>®</sup> Workload Automation
- CA Access Control
- CA ADS<sup>™</sup> (CA ADS)
- CA Advanced Systems Management (CA ASM)
- CA Cohesion Application Configuration Manager (ACM)
- CA ASM2<sup>®</sup> Backup and Restore
- CA eHealth Performance Manager
- CA Jobtrac<sup>™</sup> Job Management (CA Jobtrac JM Workstation)
- CA NSM
- CA NSM Job Management Option (CA NSM JMO)
- CA San Manager
- CA Scheduler<sup>®</sup> Job Management (CA Scheduler JM)
- CA Security Command Center (CA SCC)
- CA Service Desk
- CA Service Desk Knowledge Tools
- CA Software Delivery
- CA Spectrum<sup>®</sup> Infrastructure Manager
- CA Virtual Performance Management (CA VPM)

# Contact CA

## Contact Technical Support

For your convenience, CA provides one site where you can access the information you need for your Home Office, Small Business, and Enterprise CA 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

## Provide Feedback

If you have comments or questions about CA product documentation, you can send a message to [techpubs@ca.com](mailto:techpubs@ca.com).

If you would like to provide feedback about CA product documentation, complete our short [customer survey](#), which is also available on the CA Support website, found at <http://ca.com/docs>.

# Contents

---

<b>Chapter 1: Introduction</b>	<b>17</b>
About CA NSM	17
About This Guide	18
UNIX and Linux Support	18
CA NSM Databases	19
Management Data Base	19
Distributed Intelligence Architecture	20
Discovery	20
Visualizing Your Enterprise	21
Management Command Center	21
Other CA NSM User Interfaces	23
Discovery Classic	25
WorldView	25
Business Process View Management	25
Smart BPV	26
Customizing Your Business Views Using Unicenter Management Portal	26
Monitoring Your Enterprise	26
Unicenter Configuration Manager	26
Unicenter Remote Monitoring	27
Administering Critical Events	28
Event Management	28
Alert Management System	29
Analyzing Systems Performance	29
Correlating Important Events	29
Advanced Event Correlation	30
Unicenter Notification Services	30
Creating Customized Reports	30
Trap Manager	31
System Monitoring for z/OS	32
Integration with Other Products	32
Unicenter Service Desk	32
Unicenter Management for MOM	33
Unicenter Cisco Integration	33
Intel Active Management Technology	34
eHealth Integration with the Management Command Center	34
SPECTRUM Integration	35
Related Publications	36

---

## Chapter 2: Securing CA NSM 39

Role-based Security .....	39
Securing the MDB .....	39
MDB Users (Microsoft SQL Server Databases) .....	40
MDB User Groups (Ingres Databases) .....	40
MDB Users (Ingres Databases) .....	41
Operating System Users .....	42
Ingres Virtual Node Names (VNODES) .....	43
Component-Level Security .....	44
What is Security Management .....	46
Administrators or Power Users Group .....	47
How You Change the CA NSM Administrator Password On Windows .....	47
Change the Password for the Severity Propagation Engine User Accounts (Windows) .....	47
How You Change File Privileges on Microsoft Windows 2003 Server .....	48
Run Utilities Requiring Administrator Privileges on Windows Vista .....	49
Create Additional Users with Administrator Privileges to Run Discovery (Microsoft SQL Server Databases) .....	49
How You Create Additional Users Without Administrator Privileges (SQL Server Databases) .....	50
Create Additional Users with Administrator Privileges to Run Discovery (Ingres Databases) .....	50
How You Create Additional Users Without Administrator Privileges (Ingres Databases) .....	51
WorldView Security .....	51
Management Command Center Security .....	56
Integrating with eTrust Access Control .....	59
How Integration and Migration Works .....	60
Rules and Statistics Not Migrated .....	61
Attributes Not Migrated .....	62
Protecting and Filtering MDB Data Using Data Scoping .....	62
Data Scoping Rules .....	63
How Data Scoping Rules are Inherited .....	65
Rule Performance Issues .....	68
Data Scoping Security on Windows .....	69
Data Scoping Security on UNIX/Linux .....	69
User IDs Required for Data Scoping Rule Evaluations .....	69
Data Scoping Limitations When the MDB Resides on UNIX/Linux .....	74
Data Scoping Limitations on UNIX/Linux When the MDB Resides on Windows .....	74
Data Scoping in the 2D Map (Windows) .....	75
Activate Data Scoping on Windows .....	75
Deactivate Data Scoping on Windows .....	76
Activate Data Scoping on UNIX/Linux .....	76
Deactivate Data Scoping on UNIX or Linux .....	77
DataScope Rule Editor .....	77
Implement End-to-End Data Scoping .....	78

---

Communication Protocol Security .....	79
Encryption Levels .....	79
Agent to Manager Communication Security .....	80
Common Communications Interface (CAICCI) .....	81

## **Chapter 3: Discovering Your Enterprise** **99**

Discovery .....	99
How You Can Combine Running Classic and Continuous Discovery .....	101
Classic Discovery Multi-Homed Device Support .....	102
Discovery Classification Engine .....	102
Discovery Timestamp .....	102
How Subnet Filters Work .....	103
How Timeout Values Affect Discovery .....	103
Discovery Object Creation Rules .....	104
Types of Discovery Methods .....	104
How You Modify or Write Classification Rules .....	106
How to Enable Classification of New Classes .....	106
methods.xml file--Configure Classification Methods .....	106
classifyrule.xml--Configure Classification Rules .....	112
Device Not Discovered .....	114
Discovering Your Network Devices Continuously in Real-Time Mode .....	115
Continuous Discovery Architecture .....	115
How Continuous Discovery Monitors Your Network .....	116
How Continuous Discovery Discovers and Monitors Subnets .....	117
Continuous Discovery Default Configuration .....	117
DHCP Engine Configuration .....	118
Set the Admin Status Property for an Object Using Continuous Discovery .....	118
Exclude Classes from Discovery .....	119
How You Set Up SNMP Community Strings for Continuous Discovery .....	120
Discovery Managers .....	121
Discovery Events Reported to the Event Console .....	123
Discovery Agents .....	123
Discovery and Firewalls .....	127
Continuous Discovery Rapidly Consumes Memory .....	128
Discovering Your Network Devices on Demand Using Classic Discovery .....	129
Discovery Methods .....	130
How Agent Discovery Works .....	131
How IPX Discovery Works .....	132
How SAN Discovery Works .....	133
How Discovery Uses Subnets .....	134
How You Prepare to Run Discovery .....	135
How You Discover a Single Network .....	136

---

How You Determine the Time Required to Ping a Class B Network .....	136
How Names of Discovered Devices are Determined .....	137
Discovery Creates Incorrect Subnets .....	138
Discovering IPv6 Network Devices using Common Discovery .....	138
Common Discovery .....	139
Using Common Discovery GUI .....	141
Understanding IPv6 Discovery .....	147

## **Chapter 4: Visualizing Your Enterprise** **149**

WorldView Components .....	149
Managed Objects .....	150
Viewing Your Network Topology Using the 2D Map .....	151
Business Process Views .....	157
Determining the Relative Importance of an Object in Your Network .....	161
Set Policies for a Managed Object's Severity Using Alarmsets .....	164
Severity Propagation Service .....	164
How You Correctly Stop and Restart the Microsoft SQL Server Database .....	165
Viewing Object Details and Properties .....	166
Modifying Class Properties with the Class Editor .....	166
Viewing MIBs and WBEM Data with ObjectView .....	167
Viewing Relationships Among Objects Using the Association Browser .....	169
Viewing Links Between Objects .....	170
Viewing Historical Information about Your Network .....	171
Importing and Exporting Objects to and from WorldView .....	171
Understanding IPv6 Discovery .....	174
Registering and Updating Unicenter Components Using Unicenter Registration Services .....	177
Configuring Business Process Objects Using Business Process Views .....	180
Business Process Objects .....	180
Rules .....	180
Integration with Event Management .....	183
Creating Business Process Views Using SmartBPV .....	184
Business Process Views .....	185
Benefits of SmartBPV .....	185
How SmartBPV Works .....	185
SmartBPV Examples .....	186
How Optimizing SmartBPV Enhances Implementation .....	186

## **Chapter 5: Customizing Your Business Views** **189**

Why You Need Unicenter Management Portal .....	189
CleverPath Portal Technology .....	190
Users, Workgroups, and Security Profiles .....	191



---

Scoreboards and Dashboards .....	191
Scoreboards and Dashboards Distributed with Unicenter MP .....	192
Unicenter MP Administration .....	193
Administration Wizard .....	194
Task 1: Manage Components .....	195
Workplace Templates .....	196
Create Workplaces from Templates .....	197
Working with Components .....	198
Working with Unicenter WorldView .....	198
Working with Agent Management .....	203
Working with Unicenter Event Management .....	206
Working with Unicenter Alert Management .....	209
Working with Unicenter MP Notification .....	213
Working with Unicenter MP Reports .....	214
Working with Unicenter Service Metric Analysis .....	215
Working with Unicenter Service Desk .....	216
eHealth Integration with Unicenter MP .....	217
Working with SPECTRUM .....	219
Additional Component Integrations .....	220

## **Chapter 6: Monitoring Your Enterprise** **221**

Using Agent Technology to Monitor Resources .....	221
Understanding Unicenter Remote Monitoring .....	222
Remote Monitoring Architecture .....	223
Resource Types You Can Monitor .....	225
Securing Access to Remote Monitoring .....	227
Understanding Resource Monitoring .....	227
Basic Concepts .....	227
General Functions .....	228
Monitoring System Resources .....	236

## **Chapter 7: Host Resources MIB** **239**

Understanding Systems Management .....	243
Understanding the Architecture .....	244
Tools to Configure Managed Resources .....	250
Configuring Managed Nodes .....	255
Configuring a DSM Environment .....	258
Monitoring the Health of your DSM .....	261
Understanding Configuration Manager .....	264
Resource Model Groups .....	264
Base Profiles .....	265

---

Differential Profiles .....	267
File Packages .....	268
Delivery Schedules .....	269
Configuration Bundles .....	270
Reporting Feature .....	273

## **Chapter 8: Administering Critical Events** **275**

Event Management .....	275
Events .....	276
Event Management Policies .....	276
Event Agent .....	277
Dates and Times for Automated Event Processing .....	280
Automatic Responses to Event Messages .....	280
Event Console .....	289
SNMP Traps .....	292
Event Policy Packs .....	302
Wireless Message Delivery .....	306
Alert Management System .....	312
What Are Alerts? .....	312
How Alert Management Works .....	313
Viewing and Responding to Alerts in the Management Command Center .....	320
Integrating with Unicenter Service Desk .....	321

## **Chapter 9: Correlating Important Events** **323**

Unicenter Notification Services .....	323
How Unicenter Notification Services Works .....	324
Features of Unicenter Notification Services .....	325
Configuration and Diagnostics .....	329
Advanced Event Correlation .....	339
Why Use AEC? .....	340
How AEC Works .....	340
Alert Management Integration .....	341
Event Definitions .....	341
Configure AEC .....	342
Impact Analysis .....	348
Implement AEC .....	349
Understanding the AEC Components .....	351

## **Chapter 10: Improving Systems Performance** **363**

Analyzing Systems Performance .....	363
Performance Scope Usage .....	364

---

Working with Performance Trend .....	365
Effective Reporting with Performance Reporting .....	365
Charging for Resource Usage with Performance Chargeback .....	366
Data Fundamentals .....	366
Real-time Data Gathering .....	366
Historical Data Gathering .....	367
Performance Architecture .....	368
Data Accessibility and Management by the Performance Data Grid .....	370
Configuration Services .....	371
Main Performance Architecture Components .....	371
Administrative Tools .....	375
Secure, Centralized Configuration with Performance Configuration .....	375
Command-Line Utilities .....	375

## **Chapter 11: Creating Customized Reports** **377**

Types of Reports .....	377
Report Templates .....	378

## **Chapter 12: Securing CA NSM Objects** **379**

What is Security Management .....	379
How Security Management Works .....	380
Security Policies .....	381
How the Commit Process Works .....	381
How Security Management Is Implemented .....	382
Phase 1: Customize Security Management Options .....	383
How You Modify Windows Security Management Option Settings .....	383
How You Modify UNIX/Linux Security Management Option Settings .....	383
Options to Consider for Your Operations .....	383
Additional Options for UNIX/Linux Platforms .....	387
Set Certain Options to Absolute Values .....	388
Phase 2: Start Security in QUIET Mode .....	388
Phase 3: Create Rules for Production in WARN Mode .....	389
Defining User Groups .....	389
Defining Asset Groups .....	391
Asset Permissions .....	392
Defining Access Permissions .....	395
How CAISSF Scoping Options Work .....	396
Phase 4: Set Options for Production, FAIL Mode .....	399
How You Commit Rules in Fail Mode .....	399
How You Deactivate Security Management .....	399
Security Management Reports .....	399

---

Access Violations Written to the Event Console Log .....	400
UNIX/Linux Reports .....	400
<b>Appendix A: Unicenter NSM r11.2 UNIX and Linux Support</b>	<b>403</b>
UNIX and Linux Support .....	403
Supported Components .....	404
UNIX and Linux Support Quick Reference .....	406
<b>Appendix B: FIPS-140-2 Encryption</b>	<b>409</b>
CA NSM FIPS 140-2 Compliance .....	409
Compliant Components .....	409
Systems Performance .....	410
Active Directory Management .....	422
Agent Technology .....	423
Common Communications Interface .....	424
Management Command Center .....	426
Unicenter Management Portal .....	429
Web Reporting Server .....	430
<b>Appendix C: Managing Traps Using the Trap Manager</b>	<b>433</b>
Trap Daemon .....	433
Trap Filters .....	434
Local Versus Remote Installation .....	434
<b>Appendix D: Managing Cisco Devices Using Cisco Integration</b>	<b>435</b>
Analyzing CISCO Integration .....	435
Cisco Device Recognition .....	435
<b>Appendix E: Replicating Objects in the WorldView Repository</b>	<b>437</b>
Analyzing Repository Bridge .....	437
How Repository Bridge Works .....	438
Repository Bridge Architectures .....	439
Fanout Architecture .....	439
Aggregation Architecture .....	440
How to Determine Which Architecture to Use .....	441
Repository Bridge Components .....	442
Bridge Configuration .....	442
Bridge Control .....	443
Bridge Instances .....	444

---

Repository Bridge Supported Platforms .....	444
Repository Bridge in a Distributed Organization .....	444
Repository Bridge for a Restricted View of Resources .....	445
Repository Bridge for Problem Notification .....	445
Troubleshooting .....	445
View Repository Bridge Log Files .....	446
How to Create a Bridge Configuration File (Windows Only) .....	446
Bridging Rules (Windows) .....	448
Bridging Objects to A Repository Where a DSM is Running .....	448
Start the Bridge Configuration GUI (Windows Only) .....	448
Manage Repository Bridge Instances Using a Windows Service (Windows Only) .....	449
Create a Configuration File (UNIX/Linux) .....	450
Rule File Parameters for UNIX/Linux .....	451

## **Appendix F: Support for DMI, MOM, and SCOM** **453**

Desktop Management Interface (DMI) .....	453
DMI Service Provider .....	454
Unicenter Support for Desktop Management Interface (DMI) .....	455
Install the DMI Manager and DMI Agent .....	455
Set SNMP Destinations in the CA DMI Agent .....	456
Unicenter Management for Microsoft Operations Manager .....	457
MOM Terminology .....	457
How MOM Management Works .....	458
MOM Alerts as Event Messages .....	459
Status of MOM Entities in WorldView .....	460
Using MOM Management .....	461
Integration with Microsoft System Center Operations Manager (SCOM) .....	461
Minimum Software Requirements .....	462
SCOM Terminology .....	463
How the SCOM Integration Works .....	464
SCOM Alerts as Event Messages .....	465
Status of SCOM Entities in WorldView .....	466
SCOMMsgconfig Utility .....	466

## **Appendix G: Scanning the Systems for Viruses** **469**

Virus Scan .....	469
Downloading Virus Signature Updates .....	469
Deleting Old Scan Logs .....	470

## **Appendix H: Using Ports to Transfer Data** **471**

Utilizing and Configuring Ports .....	471
---------------------------------------	-----

---

Required Open Ports .....	472
Optional Ports .....	473
Configure the DIA Communications Port.....	474
CA Message Queuing Service (CAM) .....	476
Supported Transport Layer Protocols .....	476
Components That Use CAM/CAFT .....	477
CAM/CAFT Configuration Files .....	478
CAM/CAFT Binaries .....	478
How to Encrypt the MCC Data Transport (CAM) for AIS Providers .....	479

## **Appendix I: Integrating with CA Spectrum Service Assurance** **483**

CA NSM Connector Import .....	483
-------------------------------	-----

## **Appendix J: Integrating with CA Spectrum** **485**

CA Spectrum-NSM Integration Kit .....	485
CA Spectrum Infrastructure Manager and CA NSM Integration Guide .....	485

## **Appendix K: Integrating with CA Virtual Performance Management 11.7 VC AIM** **487**

Introduction to CA Virtual Performance Management .....	487
CA SystemEDGE Agent.....	488
Logical Partition (LPAR) AIM .....	488
Service Response Monitor (SRM) AIM .....	488
VMware vCenter (VC) AIM .....	489
Xen AIM .....	489
Zones AIM .....	490
Integration with CA Virtual Performance Management .....	490
Discover VPM Resources .....	491
IBM LPAR Object Discovered .....	491
Start the LPAR AIM Agent View .....	492
Sun Zones Objects Discovered .....	492
Start the Zones AIM Agent View .....	493
Citrix XenServer Objects Discovered .....	494
Start the Citrix XenServer AIM View .....	495
VMware Objects Discovered .....	495
Start the VC AIM Agent View .....	496
Enable AIMs in VPM integration .....	496

## **Appendix L: Job Management Option** **497**

How CA NSM Job Management Option Works .....	497
--	-----

---

CA NSM Job Management Option Job Server .....	498
Unicenter Universal Job Management Agent .....	498
CA NSM JM Option Profiles .....	499
CA NSM JM Option Variables .....	500
Types of Job Scheduling .....	500
How to Specify Where to Perform Work .....	500
How to Identify Resource Requirements for Workload Balancing .....	501
How to Schedule Work by Dates .....	502
Expanded Calendar Processing .....	503
How to Form Groups of Related Tasks (Jobsets) .....	504
Jobset Resources .....	504
Jobset Predecessors .....	505
How to Identify Work to Perform .....	507
Jobset Membership .....	507
How to Schedule Work by Special Events .....	513
Use caevent .....	514
Run a Job on Demand .....	516
How to Test Your CA NSM JM Option Policy Definitions .....	517
How to Run Additional CA NSM JM Option Reports .....	518
Autoscan .....	518
How a Job or Jobset Qualifies for Selection During Autoscan .....	519
Cleanup and Backlogging .....	519
Workload Processing .....	520
Maintenance Considerations .....	521
Job Management Logs (UNIX/Linux) .....	521
Tracking File .....	522
Undefined Calendars During Autoscan .....	523
Purge Old History Records (UNIX/Linux) .....	523
Unload the CA NSM JM Option Database Definitions to a Text File .....	523
How to Submit Jobs on Behalf of Another User .....	524
Agent/Server Configurations .....	524
Single Server .....	525
Cross-Platform Scheduling .....	526
Job Management Managers and Agents .....	527
Implementation .....	528
Windows Configuration Environment Variables .....	530
UNIX/Linux Configuration Environment Variables .....	532
Environment Variables for Jobs and Actions .....	533
Monitor Workload Status .....	534
Jobflow Tracking on Windows .....	535





# Chapter 1: Introduction

---

This section contains the following topics:

[About CA NSM](#) (see page 17)

[About This Guide](#) (see page 18)

[UNIX and Linux Support](#) (see page 18)

[CA NSM Databases](#) (see page 19)

[Management Data Base](#) (see page 19)

[Distributed Intelligence Architecture](#) (see page 20)

[Discovery](#) (see page 20)

[Visualizing Your Enterprise](#) (see page 21)

[Customizing Your Business Views Using Unicenter Management Portal](#) (see page 26)

[Monitoring Your Enterprise](#) (see page 26)

[Administering Critical Events](#) (see page 28)

[Analyzing Systems Performance](#) (see page 29)

[Correlating Important Events](#) (see page 29)

[Creating Customized Reports](#) (see page 30)

[Trap Manager](#) (see page 31)

[System Monitoring for z/OS](#) (see page 32)

[Integration with Other Products](#) (see page 32)

[Related Publications](#) (see page 36)

## About CA NSM

CA NSM delivers innovative, secure, and platform-independent management to let you deploy single platform or heterogeneous business applications. CA NSM solutions help you sustain an optimized, on-demand infrastructure, maximizing your IT investment by continuously assessing and self-managing network and systems elements.

CA NSM lets organizations deploy and maintain a complex, secure, and reliable infrastructure that supports business objectives. It helps ensure the continuous health and performance of your critical infrastructure through innovative and intelligent techniques to help you control costs while maintaining or increasing responsiveness to changing business priorities. Its ability to integrate with other solutions in the CA portfolio and share information using a common database provides unparalleled intelligence for CA's EITM strategy.

## About This Guide

This guide is intended for use by system administrators and contains general information about how to customize, configure, and maintain CA NSM after installation and implementation. For more detailed information, including specific procedures, see the CA NSM Management Command Center online help.

The topics that follow describe the components that are included with or that can be integrated with your CA NSM installation.

**Note:** For detailed information about installing and implementing CA NSM if it has not yet been installed, see the *Implementation Guide*.

## UNIX and Linux Support

Unicenter Network and Systems Management r11.2 provides support on UNIX and Linux platforms for key CA NSM manager components. This release provides an upgrade path for UNIX and Linux users with Unicenter NSM 3.1 and r11 managers installed.

The components supported on UNIX and Linux platforms mirror those components in the base CA NSM r11.2 product, with the inclusion of UNIX and Linux support. Therefore, the areas in this guide for the base CA NSM r11.2 components supported on UNIX and Linux also apply to UNIX and Linux users. For more information about Unicenter NSM UNIX and Linux support and a listing of the components supported and the applicable areas of this guide, see the appendix "Unicenter NSM r11.2 UNIX and Linux Support."

**Note:** CA NSM r11.2 for UNIX and Linux does not support Ingres. Any UNIX and Linux information in the CA NSM documentation set pertaining to Ingres databases does not apply to CA NSM r11.2 users.

## CA NSM Databases

In CA NSM r11, the database tool used for the MDB is Ingres for both Windows and UNIX/Linux. In CA NSM r11.1 and r11.2, however, the database tool used for the MDB on Windows platforms is Microsoft SQL Server. The documentation for CA NSM r11.2 has information for both Ingres databases and Microsoft SQL Server databases, so be aware that some of it may not apply, depending on the CA NSM version you are running.

On UNIX and Linux platforms, Unicenter NSM r11.2 does not use Ingres for the MDB. Unicenter NSM r11.2 for UNIX and Linux platforms supports a free embedded database PostgreSQL. Therefore any Ingres or Microsoft SQL Server information does not apply to Unicenter NSM r11.2 for UNIX and Linux. For more information about the PostgreSQL database, see the *MDB Overview*.

CA NSM r11 users can migrate the Ingres database to the r11.2 PostgreSQL database. For more information, see the *Migration Guide*.

## Management Data Base

An integrated Management Database (MDB) is the critical foundation for achieving effective management solutions through information centralization and product integrations. CA delivers an integrated MDB as the foundation for integration across all management solutions. The MDB combines all data from currently distinct disciplines--operations, storage, security, life cycle, and service management--and provides the foundation necessary to manage and optimize an organization's IT infrastructures. Customers and third-party partners can extend the MDB to store related IT management data from non-CA software products and tools.

The MDB provides a single integrated database schema for the management data stored by all CA products, both distributed and mainframe. The MDB is delivered with all CA products at no additional cost and runs on a high performance Ingres, Microsoft SQL Server, or PostgreSQL database. The single schema of the MDB enables integration of CA products without API-level programming efforts.

You can use standard database utilities to backup and restore the MDB.

**Note:** For more information about implementing the MDB for CA NSM, see the *Implementation Guide*.

## Distributed Intelligence Architecture

Distributed Intelligence Architecture (DIA) allows a central location to manage all components and aspects of a network. DIA makes data requests and retrievals standard across different Unicenter components by providing a generic mechanism that permits the dynamic deployment of necessary files to facilitate the correct monitoring of any given system. Deployment is generic and open to growth. DIA allows for high speed, secure communications to transport data while providing remote node management and inherent failover capabilities. All out-bound communications from all DIA components can use the secure sockets provided if you enable public or private key encryption.

**Note:** For more information about DIA architecture, configuration, and encryption, see the *Implementation Guide*.

## Discovery

Discovery discovers and classifies devices on IP and IPX networks. It provides both an ad hoc (on demand) and continuous (real-time) mode. It provides discovery services to other CA Common Services components and updates the MDB with newly discovered and classified network objects.

When you install your product, you can use any of the following types of Discovery:

### **Classic Discovery**

Provides on demand discovery that lets you decide which subnets you want to discover and when. You can also configure Classic Discovery to run at regular intervals, which can be used as an alternative to Continuous Discovery and ensures that your discovered environment in the MDB is always current. You can start a Classic Discovery from the Discovery Classic GUI, the Management Command Center, the Unicenter Browser Interface, or the command line.

### **Continuous Discovery**

Provides event-driven and ongoing discovery. Continuous Discovery employs a manager and agents that continuously scan your network in real-time mode for new devices or changes in IP addressing of existing IP devices. You can configure Continuous Discovery for optimal load balancing between the Discovery Agents and the Discovery Manager. If you choose this method of discovery, you must install the Discovery Agents and the Discovery Manager.

### **Common Discovery**

Discovers IPv6 networks. The Common Discovery Import utility discovers IPv6 networks using Common Discovery technology and imports IPv6 addresses into WorldView, where they are integrated with existing networks.

**Note:** For more information about Discovery, see the "Discovering Your Enterprise" chapter in this guide. For more information about Common Discovery and the Common Discovery Import utility, see the chapters "Discovering Your Enterprise" and "Visualizing Your Enterprise."

## Visualizing Your Enterprise

CA has performed extensive analysis on user interfaces and how various roles within a data center work effectively. The result is a variety of interfaces tailored for specific users. Using this role-based management methodology, you can easily navigate IT complexity using scoped and meaningful visualizations.

**Note:** For more information about WorldView, Business Process View Manager, and Smart BPV, see the chapter "Visualizing Your Enterprise." For more information about the Desktop Management Interface, see the appendix "Support for DMI, MOM, and SCOM."

## Management Command Center

The Management Command Center (Unicenter MCC) user interface integrates all Unicenter enterprise and network monitoring functionality into a single console. The Management Command Center provides dynamic multi-viewer content relevant to any asset in the MDB by providing a workplace that integrates relevant plugins, such as Alert Management System alerts, Event Management System events, Agent events, Dashboards, and Web Reporting Services, for the tree node you select.

**Note:** For specific information about using Unicenter MCC, see the Unicenter MCC online help.

## Providing Access to CA NSM Components in the Management Command Center

Every CA NSM installation has one or more directories of installed components known as Global Catalogs. Created at installation time, a Global Catalog provides information about the availability and location of CA NSM components to the Management Command Center, and provides access to WorldView, the Distributed State Machines, and Enterprise Management Managers. At installation, WorldView, Enterprise Management, DSM, and other information sources register with a Global Catalog as *providers*. Providers enable applications, data, and business components to appear as native within the object infrastructure so that all data can be accessed and integrated, regardless of origin, format, or location. Depending upon the size and configuration of your network, you may choose to set up a single Global Catalog that contains information about all the providers (CA NSM components) at your site, or you may choose to set up multiple Global Catalogs, where each catalog provides access to a subset of the available providers.

Any Management Command Center client pointing to a Global Catalog can see the entire CA NSM installation that has registered to that particular Global Catalog. The Management Command Center client running on your computer, for example, may be using a Global Catalog on a server somewhere else. Your local browser can access any provider known to that Global Catalog. Although a Management Command Center client has the ability to access more than one Global Catalog, the Management Command Center is configured to use one Global Catalog as its *master catalog* for the purpose of accessing current information about available providers. When you open the Management Command Center, it updates the local catalog with information about all currently published providers in the master catalog and removes information about providers that may have been "unpublished" (removed) from the Global Catalog since the last time the application was started. This process is referred to as the catalog "synchronization" process.

**Note:** For more information about Global Catalogs, see the Management Command Center online help.

## Start the Management Command Center

You can open the MCC user interface in several different ways.

To start Unicenter MCC on Windows, do one of the following:

- Select Start, Programs, CA, Unicenter, NSM, Management Command Center.
- Run the `tndbrowser.bat` command.
- Run the `showinmcc` command to open the Management Command Center for a specified node. If Unicenter MCC is already open when you run this command, then navigate to the specified object.

To start the Management Command Center on UNIX or Linux, run the `camcc` command from the `$JI_SYSTEM/bin` directory.

You can route the `camcc` display by setting the `DISPLAY` environment variable to the proper hostname or IP address. By default, only one Management Command Center instance is permitted per UNIX/Linux server, but you can edit the `$JI_SYSTEM/.max_ue` file to change this limit to reflect the number of instances of Unicenter MCC that you want to run simultaneously on the server. The new limit takes effect when all instances of the Unicenter MCC are restarted.

**Note:** For more information about the `tndbrowser.bat`, `showinmcc`, and `camcc` commands, see the online *CA Reference*.

## Other CA NSM User Interfaces

In addition to the Management Command Center, CA NSM provides the following other user interfaces:

- Unicenter Browser
- Unicenter Classic
- WorldView Classic
- Discovery Classic
- Agent dashboards

## Unicenter Browser Interface

The Unicenter Browser Interface is an Internet and intranet-based user interface that provides virtually all of the functionality found in the WorldView Classic user interface. Since it is a Java applet, it does not require client installation, which makes it accessible from virtually anywhere. Windows users can access it locally on the system where the MDB resides through Start, Programs, CA, Unicenter, NSM, Unicenter Browser Interface. Remote systems can access it by entering a URL address on any Java-enabled web browser in the following form:

```
http://wvserver/ubi/ubi.html
```

### **wvserver**

Specifies the name or IP address of the web server on which the server components reside.

You can use many of the procedures for the WorldView Classic GUI, contained in the CA Procedures located in the Online Books program group, with this interface.

**Note:** You must install the Unicenter Browser from the Product Explorer before it appears in the NSM program group.

## Unicenter Classic

Unicenter Classic refers to the traditional Windows-based user interface delivered with previous versions of CA NSM. Unicenter Classic includes the WorldView, Enterprise Management, and Discovery program groups accessed through Start, Programs, CA, Unicenter, NSM. Procedures based on the Unicenter Classic GUI are contained in the online CA Procedures.

Unicenter Classic also includes the cautil command line interface.

## WorldView Classic

WorldView Classic refers to the traditional Windows-based user interface. WorldView Classic includes the WorldView program group accessed through Start, Programs, CA, Unicenter, NSM, WorldView. Procedures based on the WorldView Classic GUI are contained in CA Procedures located in the Online Books program group.

WorldView Classic also includes the cautil command line interface.

**Note:** For more information about the cautil command line interface, see CA Reference in the Online Books program group.

## Agent Dashboards

Dashboards display real-time information from CA NSM agents. A dashboard lets you combine on one screen multiple metrics from one or many agents and one or many hosts. Each metric is presented in an individual tile. Dashboards poll the data from the agents and show the metrics "as is."

The Management Command Center supports two types of dashboards:

- **Agent dashboards** display information about a single agent, which consists of a number of chart titles each of which reflects the state of a particular variable/group monitored by the agent on a host.
- **Server dashboards** display information about each agent installed on the host.

To display dashboards, a CA NSM Web Reports and Dashboards server must be installed and running on a host that the Management Command Center can access. If a dashboard server is found, a Dashboard viewer option becomes available for agent objects in the Topology and DSM view trees. After selecting an agent object, you can open the Dashboard viewer using the Add or Open Viewer context menu (available by right-clicking the object). You can also click the right pane drop-down list and choose Dashboards.



The first time you request a dashboard a connection dialog appears, which allows you to select the dashboard server you want to use. The connection dialog also contains user name and password fields for specifying the credentials to use when the server is accessed. The information you enter is saved and used for subsequent access to the same server for the remainder of your session.

## Discovery Classic

Discovery Classic refers to the traditional Windows-based user interface. Discovery Classic includes the Discovery program group accessed through Start, Programs, CA, Unicenter, NSM, Discovery. Procedures based on the Discovery Classic GUI are contained in the CA Procedures located in the Online Books program group.

## WorldView

Unicenter WorldView offers a highly visual and intuitive approach to enterprise management with the 2D Map available through the Management Command Center, the Unicenter Browser Interface, and the WorldView Classic GUI. The 2D Map works as an infrastructure navigator, allowing you to view any part of your enterprise with the click of a button. For example, you can view all assets in your network—from the global network to the local subnets, hubs, bridges links, the servers and workstations connected to them, their processors and drives, all the way down to the databases and applications.

WorldView provides support for the Desktop Management Interface (DMI) specification. This feature lets you manage the installed hardware and software on your PCs. This can be accomplished locally, as well as remotely across your network. DMI is available on Windows only.

## Business Process View Management

Business Process View Management (BPVM) lets you use Business Process objects to view and manage your network. You can use Business Process objects to apply new rules to your system to determine how states are propagated from existing WorldView objects using methods that include simple counters and complex correlation rules. BPVM is available on Windows only.

**Note:** The BPVM service manages all BusinessProcessObjects so If it is not created by the service, the object is deleted upon restart of the service.

## Smart BPV

Smart Business Process View Management (SmartBPV) lets you automatically create and dynamically update Business Process Views. Through analysis of network activity, SmartBPV identifies infrastructure elements that support a specific application and automatically builds and continuously updates a focused view for management. SmartBPV is available on Windows only.

## Customizing Your Business Views Using Unicenter Management Portal

Unicenter Management Portal provides role-based, dynamic, and personalized views of management information securely over the web. Integration with Unicenter Management Portal lets CA NSM components and other CA solutions consolidate data from numerous and disparate data sources, querying, reporting, and presenting it in a unified view that suits each viewer's needs.

**Note:** For more information, see the chapter "Customizing Your Business Views."

## Monitoring Your Enterprise

To facilitate comprehensive and integrated network polling and administration, CA NSM uses Agent Technology to automate manager tasks and responses to events. Agent Technology monitors and reports the state and status of machines and applications. Agent Technology lets you monitor resources (also called managed objects) as well as the applications. The status of these devices is displayed on the 2D Map.

**Note:** For more information about Agent Technology, agent dashboards, Unicenter Configuration Manager, and Unicenter Remote Monitoring, see the chapter "Monitoring Your Enterprise."

## Unicenter Configuration Manager

You can manage agent configurations centrally and automatically using the Unicenter Configuration Manager.

Unicenter Configuration Manager is a stand-alone utility for the scheduled auditing of existing, or delivery of legitimate configuration data, for Unicenter product components.

To access the Unicenter Configuration Manager using a web browser, enter the following URL:

`http://UCMServerName:port/wiser`

**UCMServerName**

Specifies the name of the computer on which Unicenter Configuration Manager is installed.

**port**

Specifies the port for the Unicenter Configuration Manager server.

To access the Unicenter Configuration Manager agent configuration tool from the Management Command Center, a Unicenter Configuration Manager server must be installed and running on a host that the Management Command Center can access.

## Unicenter Remote Monitoring

Unicenter Remote Monitoring is a component of CA NSM that provides you the ability to remotely monitor the health and availability of your network resources, including production servers and workstations.

The key distinguishing feature of Remote Monitoring is that it works by using non-intrusive technology. This technology lets you monitor network resources without installing an agent on each monitored device. Instead, the Remote Monitoring Agent is installed on a single remote machine that probes the monitored resources for data, status, and other information used to assess the health and availability of that resource. This type of monitoring is particularly useful in these situations:

- You need to deploy a resource monitoring solution quickly or temporarily
- You need to monitor systems where the installation of conventional agents would be intrusive, is prohibited by corporate policy, or is simply impractical
- You need to monitor resources at a remote location through the WAN

Unicenter Remote Monitoring can monitor the following resource types:

- Windows
- UNIX
- Linux
- Mac OS X
- IP

## Administering Critical Events

CA NSM lets you administer critical events throughout your network using the Event Management component and the Alert Management System (AMS).

**Note:** See the chapter "Administering Critical Events" for more information about Event Management and the Alert Management System.

### Event Management

Event Management, the focal point for integrated message management throughout your network, can monitor and consolidate message activity from a variety of sources. It lets you identify event messages that require special handling and initiate a list of actions for handling an event. Through support of industry-standard facilities, you can channel event messages from any node in your network to one or more monitoring nodes. You can centralize management of many servers and ensure the detection and appropriate routing of important events.

For example, you may want to route message traffic to different event managers:

- Event and workload messages to the production control event manager
- Security messages to the security administrator's event manager
- Problem messages to the help desk administrator's event manager

By filtering messages that appear on each console, you can retrieve specific information about a particular node, user, or workstation.

Wireless Messaging provides alternate channels for operator input in situations where the operator cannot access a CA Event Console. The supported messaging protocols are email and pager. Using the SMTP/POP3 mail messaging protocol, you can send and receive pager messages from two-way pager devices. An incoming message can trigger any series of actions you define for Event Console to perform in response to it.

You can install the Event Manager on Windows and UNIX or Linux platforms. For more information about installation options, see the *Implementation Guide*.

## Alert Management System

The Alert Management System (AMS) is a tool for organizing and tracking the most important events in an enterprise or a logical segment of an enterprise. It lets you focus on and manage the highest severity IT events.

AMS provides tools for defining alert policies and multiple panes in the Management Command Center for viewing alerts. AMS also lets you link to Unicenter Service Desk and Unicenter Service Desk Knowledge Tools, which is a customer support application that manages calls and IT assets, resolves problems, and shares corporate knowledge.

## Analyzing Systems Performance

Systems Performance lets you gain control over your computer systems through value-added information related to performance. Systems Performance lets you monitor key parameters that influence performance of your systems and applications, and compare what is actually happening against a set of predefined operational parameters. Systems Performance also lets you determine the most efficient use of your computer resources by letting you examine how busy each server is, which servers are overused or underused, when computers are being used, whether usage is increasing or decreasing, and where bottlenecks occur.

**Note:** For more information about Systems Performance, see the chapter "Analyzing Systems Performance."

## Correlating Important Events

CA NSM lets you analyze network, system, and application events to identify the true cause within a series of related events. You can then set up policies to identify, notify, and respond to problems that may impact end-user service. Escalation and notification policies can be tailored for each group or individual to reduce resolution time. These tasks are accomplished using Advanced Event Correlation (AEC), and Unicenter Notification Services.

**Note:** For more information about Advanced Event Correlation and Unicenter Notification Services, see the chapter "Correlating Important Events."

## Advanced Event Correlation

Advanced Event Correlation (AEC) integrates seamlessly with Unicenter Event Management to provide powerful Event Correlation, Root Cause, and Impact Analysis capabilities. When used in combination with existing Unicenter features, it lets you rapidly identify the root cause of problems being reported to the event console, by increasing the quality and reducing the quantity of the data that you have to process. Event reformatting and suppression capabilities also help to ensure that existing management procedures implemented through Message Records and Actions are only invoked when applicable, significantly reducing the number of false alarms encountered.

Using AEC, you can do the following:

- Distinguish between failure messages
- Determine the root cause of failure
- Provide an impact analysis of a failure
- Diagnose and filter unwanted messages
- Respond to dynamically changing environments

## Unicenter Notification Services

Unicenter Notification Services lets you send wired and wireless messages using various protocols and devices to get the attention of operators or administrators, wherever they are, who must resolve problems or attend to emergencies.

## Creating Customized Reports

CA NSM reports are based on a common technology called the Web Reporting Server (WRS).

CA NSM provides a set of predefined scoreboards to report various groups of managed resources according to their status, and the number of events generated from these managed resources. CA NSM also provides report templates that you can use to create a customized report by providing a required parameter. You can also modify the predefined scoreboards to create a customized report.

CA NSM also provides a set of canned WorldView and agent reports that let you view managed resource status, resource inventory, and so forth.

To create reports in CA NSM, you use the Report Viewer. The Report Viewer is a reporting feature of the Unicenter MCC that displays canned reports for the following types of information:

- Administration
- Documentation
- Agent Technology
- WorldView
- Unicenter Scoreboards

Canned reports are visible when you select Reports from the left pane drop-down list. When you select a report, WRS opens in the right pane viewer. Reports are viewed as HTML in the right pane using a web browser window.

**Note:** For more information about using WRS, see the WRS online help.

For those customers who choose not to install Unicenter MCC, CA NSM provides the Report Explorer. You can use the Report Explorer to create customized reports just as you can using WRS. The Report Explorer uses the Windows Explorer interface to view, print, edit, and create reports. To open the Report Explorer, choose Start, Programs, CA, Unicenter, NSM, Utilities, Report Explorer.

**Note:** For more information about using the Report Explorer, see the Report Explorer online help.

## Trap Manager

The Trap Manager is a component of CA NSM that lets you perform sophisticated trap database and trap filter file management. You can use the Trap Manager to manage trap information and translation messages stored in the Management Database (MDB) and trap filters stored in the trap filter file.

To sign on to the Trap Database, go to Start, Programs, CA, Unicenter, Trap Manager, Unicenter Trap Manager. The Enter window appears. For more information, see the online help.

## System Monitoring for z/OS

The z/OS system agent enables you to monitor key resources of your z/OS system and provides status, event, and configuration information. The agent can monitor individual resources as well as the health of an entire system, allowing you to quickly determine the cause of a problem. The z/OS system agent also monitors UNIX System Services (USS) resources.

The z/OS system agent puts you in control by allowing you to determine the warning and critical thresholds for each monitored resource. The agent monitors these resources and, whenever a user-defined threshold is exceeded, sends an SNMP trap.

## Integration with Other Products

CA NSM is tightly integrated with other CA products, such as Unicenter Service Desk and eTrust Security Command Center. CA NSM also provides integration with third-party products such as Microsoft Operations Manager and Cisco devices. This section gives you a brief overview of the products you can use to gather information from a wide variety of platforms and architectures.

### Unicenter Service Desk

CA NSM provides a connection to Unicenter Service Desk (Service Desk), which is a customer support application that manages calls, tracks problem resolution, shares corporate knowledge, and manages IT assets.

The integration with Service Desk is through the Alert Management System (AMS) and the Management Command Center (Unicenter MCC) interface. The connection is installed automatically with CA NSM. Interaction with the Service Desk reduces the workload of the customer support staff because some manual tasks are eliminated.

You can also access Unicenter Service Desk in Unicenter Management Portal by publishing a Service Desk portlet, or by performing a Service Desk action on an alert in the Alert Console.



## Unicenter Management for MOM

Unicenter Management for MOM (MOM Management) integrates CA NSM with Microsoft Operations Manager (MOM).

Microsoft Operations Manager delivers operations management by providing event management, proactive monitoring and alerting, reporting, and trend analysis. It helps administrators monitor and manage the events and performance of Windows 2000 or 2003 server systems. MOM is similar to CA NSM Event Management.

The integration between MOM and CA NSM provides a central location for performing management functions. You can see the status of MOM servers and managed PCs using WorldView, and you can change the status from there. You can also view open MOM alerts and CA NSM events in one location, the Event Console.

**Note:** CA NSM provides integration kits to both of Microsoft's management applications, Microsoft Operations Manager (MOM) and System Center Operations Manager 2007 (SCOM). Although the integrations to MOM and SCOM can coexist on the same management server, each one integrates only with its Microsoft counterpart.

## Unicenter Cisco Integration

Cisco Integration is a component of CA NSM that lets you manage your Cisco devices. CA NSM provides class definitions for Cisco routers and switches, but does not identify the Cisco device model. Cisco Integration provides object identifiers for each Cisco device, which lets CA NSM automatically discover and classify Cisco devices. You can then use CA NSM to visualize and monitor these Cisco devices.

## Intel Active Management Technology

The advantages of using Intel AMT with CA NSM are that now you can do the following:

- Perform enhanced discovery of computers regardless of the health and state of their operating system.
- Control the power state of any operating system.
- Perform remote problem determination, which lets you troubleshoot hardware errors for Intel AMT-enabled devices.
- Resolve hardware-related problems quickly and efficiently using out-of-band communication with network devices.
- Provide enhanced security threat protection for all the systems on the LAN.
- Discover hardware asset information using out-of-band communications.

## eHealth Integration with the Management Command Center

The Unicenter MCC provides access to the eHealth Suite of products so that you can monitor the performance of eHealth objects, generate reports, and create and close alerts associated with eHealth alarms. The eHealth Suite delivers comprehensive fault, availability, and performance management across complex, heterogeneous systems and application environments. eHealth collects a wide variety of data from your network infrastructure to generate alarms and reports.

From Unicenter MCC you can access the following features of eHealth:

### **Business Service Console**

The eHealth *Business Service Console* (BSC) is a Web-based tool that provides a high-level view of the availability and performance of business services across an organization. The BSC offers customized business views, immediate notification of performance problems, and drill-down capability for fault resolution.

### **eHealth Report Server**

The Report Server is part of the eHealth Web interface. The eHealth *Web interface* provides access to reports and applications through Web browsers on local systems.

### **At-a-Glance Reports**

The Unicenter MCC provides access to eHealth At-a-Glance reports for Alert Management System (AMS) alerts that were created from eHealth alarms. The reports are also available for eHealth objects in the Topology and DSM views. eHealth *At-a-Glance reports* show the overall performance of an eHealth object for a specified time period. The reports consist of several charts that display different performance statistics on one page.

### Alarm Detail Reports

The Unicenter MCC provides access to eHealth Alarm Detail reports for AMS alerts that were created from eHealth alarms. eHealth *Alarm Detail reports* show the availability and performance history over time of an eHealth object that caused an alarm to be generated.

### Trend Reports

The Unicenter MCC provides access to eHealth Trend reports for eHealth objects in the WorldView Topology and DSM Views. eHealth *Trend reports* are charts that plot a variable for an object over a period of time. Trend reports can also show variables for groups of objects. The reports can reveal patterns over time and relationships between objects and between variables. The available Trend reports are Availability, Bandwidth, and Error, depending on the type of managed object.

### eHealth Alarms and netHealth Exceptions Create AMS Alerts

Based on policy that you deploy, eHealth alarms and netHealth exceptions create alerts automatically. When alarms are closed, the associated alerts are closed. Likewise, if an alert associated with an eHealth alarm is closed through AMS, the alarm is also closed.

**Note:** If you receive a security error when closing an alert associated with an eHealth alarm or netHealth exception, see *Authorize Users to Run Commands*.

## SPECTRUM Integration

CA NSM provides integration with CA Spectrum Infrastructure Manager, which is a network fault management tool that provides proactive management of your network infrastructure through root cause analysis, impact analysis, event correlation, and service level management.

CA NSM integrates with CA Spectrum through an Integration Kit that you can install from the Unicenter Product Explorer. After you install the kit, you can view CA Spectrum device model alarms from the MCC, 2D Map, and the Event Console. You can also launch the CA Spectrum OneClick interface from the MCC, 2D Map, or Management Portal.

For more information about integrating with CA Spectrum, see the *CA Spectrum Infrastructure Manager and CA NSM Integration Guide (5147)*, which is included with CA NSM and CA Spectrum.

## Related Publications

The following guides provide information that you will find useful. Most are available on the CA NSM installation media.

### **Administration Guide**

Is intended for use by system administrators and contains general information and procedures about how to secure, customize, configure, and maintain CA NSM after installation and implementation. Individual chapters describe the components that are included with or that can be integrated with your CA NSM installation.

### **Agent Technology Support for SNMPv3**

Provides information about how Agent Technology can take advantage of the SNMPv3 protocol. Documents how the security information is handled on the manager and agent side as well as how it is applied to the managed systems. SNMPv3 configuration and usage details are provided in this guide.

### **CA Procedures**

Contains procedures and processes for all components of CA NSM, including WorldView, Agent Technology, Enterprise Management, Event Management, CAICCI, Data Scoping, Discovery, Notification Services, Wireless Messaging, Security Management, and CA NSM Job Management Option.

### **CA Reference**

Contains commands, parameters, and environment variables for all components of CA NSM, including Advanced Event Correlation, Agent Technology, Enterprise Management, Event Management, CAICCI, Data Scoping, Discovery, Notification Services, Wireless Messaging, Security Management, CA NSM Job Management Option, and WorldView.

### **Implementation Guide**

Contains architecture considerations, pre-installation tasks, installation instructions, post-installation configuration information, and implementation scenarios. Appendixes include in-depth information about Distributed Intelligence Architecture (DIA), the MDB, and the CA High Availability Service (HAS) for cluster aware environments. This guide is intended for users who are implementing CA NSM on a new system.

### **Inside Active Directory Management**

Provides general information, installation scenarios, and configuration procedures for Active Directory Management.

### **Inside Event Management and Alert Management**

Provides detailed information about Event Management (message records and actions), Advanced Event Correlation, and Alert Management.

**Inside the Performance Agent**

Contains detailed information about the configuration and use of the Performance Agent.

**Inside Systems Management**

Describes systems management from the CA NSM architecture perspective. The guide describes the different layers (WorldView, Management Layer, Monitoring Layer) and associated components, for example: Distributed State Machine (DSM), Unicenter Configuration Manager, dashboards, and so on.

**Inside Systems Monitoring**

Explores how to use and configure the system agents of CA NSM to monitor the system resources in your environment. The chapters guide you through the process of configuring and optimizing the agent for your special requirements.

**Inside Systems Performance**

Contains detailed information about the three architectural layers of Systems Performance, and provides guidance in the deployment, configuration, use, and best practices of the Systems Performance components.

**MDB Overview**

Provides a generic overview of the Management Database (MDB), a common enterprise data repository that integrates CA product suites. The MDB provides a unified database schema for the management data stored by all CA products (mainframe and distributed). The MDB integrates management data from all IT disciplines and CA products. The guide includes implementation considerations for the database systems that support the MDB and information specific to the CA NSM implementation of the MDB.

**MIB Reference Guide**

Provides detailed information about each MIB attribute of the CA NSM system agents.

**Migration Guide**

Provides detailed upgrade and migration instructions. This guide is only available on the CA Support website: <http://ca.com/support>

**Programming Guide**

Provides details for constructing applications by CA development teams and by third parties and their clients. The guide is intended for developers who use one or more of the application programming interfaces (APIs) in the SDK to develop applications for use with CA NSM. Key among these APIs are the WorldView API, the Agent Technology API, and the Enterprise Management API.

### **Readme Files**

Provides information about known issues and information discovered after CA NSM publication. The following readme files are available:

- The CA NSM r11.2 SP2 for UNIX and Linux readme
- The CA NSM r11.2 SP2 Windows readme
- The Unicenter Management Portal readme

### **Release Notes**

Provides information about operating system support, system requirements, new and changed features, published fixes, international support, and the documentation roadmap. The following release notes are available:

- The CA NSM r11.2 SP2 for UNIX and Linux release notes
- The CA NSM r11.2 SP2 release notes
- The Unicenter Management Portal release notes

### **Unicenter Management Portal Implementation Guide**

Provides installation, deployment, and basic administrative instructions for Unicenter Management Portal.

### **CA Green Book, Systems Management**

Identifies the CA solution for managing challenges involved in maintaining the performance and availability of complex server infrastructures. The CA solution provides proactive management of servers as part of an overall effort to improve service levels, and minimize the costs of managing the computing infrastructure through increased automation. It provides a view of the requirements for systems management and best practices for deployment. This guide is available online at:

<https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/common/greenbooks.html>.

### **CA Green Book, Service Availability Management**

Describes how to deliver integrated end-to-end performance and event management that is centered on services. The CA Service Availability Management solution leverages the Manager of Managers integration capabilities of CA NSM and eHealth and explains how to take advantage of those capabilities. It includes details on how to install and configure a variety of management solutions to provide simpler and more comprehensive management and monitoring of IT services. This guide is available online at:

<https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/common/greenbooks.html>.

# Chapter 2: Securing CA NSM

---

This section contains the following topics:

[Role-based Security](#) (see page 39)

[Securing the MDB](#) (see page 39)

[Component-Level Security](#) (see page 44)

[Integrating with eTrust Access Control](#) (see page 59)

[Protecting and Filtering MDB Data Using Data Scoping](#) (see page 62)

[Communication Protocol Security](#) (see page 79)

## Role-based Security

CA NSM was developed with detailed security and now uses a role-based approach so that the management station is not a point of concern for today's security-conscious IT environments.

CA NSM and its options are unique in providing a security methodology that protects corporate assets and also makes the system easier to manage because CA NSM security lets you segregate security according to a user's role within the organization.

CA NSM can be secured at the following levels:

- Database security (MDB)
- CA NSM component-level security, such as securing WorldView tables, Agent Technology security, Enterprise Management security
- General product security for the primary communications protocols

**Note:** For information about using the Security Management component to secure CA NSM objects such as calendar and event, see the chapter "Securing CA NSM Objects."

## Securing the MDB

The Management Database (MDB) maintains information about all objects in your enterprise, including their properties and relationships to other objects. You must ensure the integrity and availability of this information.

The Management Database (MDB) creates a hierarchy that you must understand so that you can access it correctly from CA NSM and its components. The CA NSM database security model uses one of the following ways to connect to the MDB, depending on which database you are using:

- For Ingres databases, private VNODES instead of an installation password.  
All connections to the MDB require a valid operating system user ID and password. That user ID must also be defined to Ingres.
- For Microsoft SQL Server databases, Microsoft SQL Server authentication or Windows authentication.  
All connections to the MDB require either a valid operating system ID or Microsoft SQL Server user ID. Different applications require each method of authentication.

This section defines the preferred way of connecting to and accessing database objects in the MDB. Topics covered include the following:

- MDB User Groups (Ingres)
- MDB User Roles (Microsoft SQL Server)
- MDB Users
- Operating System Users
- Virtual Node Names (VNODEs) (Ingres)

## MDB Users (Microsoft SQL Server Databases)

For a CA NSM user to connect to the MDB, that user must be one of the following types of users:

- A valid Microsoft SQL Server user that has a default group assigned to a CA NSM Microsoft SQL Server role, such as uniadmin or uniuser at the product level or wvadmin, wvuser, emadmin, or emuser at the component level.
- A valid operating system user on the Microsoft SQL Server computer that is Windows-authenticated to Microsoft SQL Server.

A Microsoft SQL Server user can be set up with a password, or a Windows-authenticated user can be set up by a Microsoft SQL Server user who has the system administrator database role.

## MDB User Groups (Ingres Databases)

User group identifiers enable the database administrator (or user that has the security privilege) to grant identical privileges to a group of users, which simplifies the task of user ID administration. A group can contain any number of users.



As part of the MDB definition, the following user groups are defined for CA NSM:

- An administrator group called uniadmin at the product level and wvadmin and emadmin at the component level.
  - Table privileges: Insert, Update, Delete, Select
  - Users assigned to these groups have the Security privilege that allows uniadmin to "impersonate" user mdbadmin, which allows WorldView and Enterprise Management tables to be owned by user mdbadmin and to grant access to other users accordingly. This security privilege is required for creating and updating WorldView classes and Enterprise Management data.
- A read-only user group called uniuser at the product level and wvuser and emuser at the component level.
  - Table privileges: Select
- For each CA NSM component, a user group could be created with grants for the tables within that subcomponent.
  - For an administrator group
  - For a read-only user group

By default, no users are defined for these groups. These groups are available if you want to protect tables at the component level.

- Other component groups are defined for other components of CA NSM.

**Note:** For detailed information about administering Enterprise Management Database privileges in Ingres on UNIX and Linux operating systems, see the *Inside Event Management and Alert Management* guide.

## MDB Users (Ingres Databases)

For a CA NSM user to connect to the MDB, that user must be one of the following types of users:

- A valid Ingres user that has a default group assigned to a CA NSM Ingres group, such as uniadmin or uniuser at the product level or wvadmin, wvuser, emadmin, or emuser at the component level.
- A valid operating system user on the Ingres server computer.

Ingres users are defined without Ingres passwords. Ingres verifies the operating system user before checking whether the user is defined to Ingres.

An Ingres user can be set up with a password, at the time the user is created or by an Ingres user who has the `maintain_users` privilege. This password has no connection with the operating system user's password. Ingres users with privileges can change their own passwords using the `ALTER USER SQL` command by specifying the old and new passwords. Only an Ingres user with the correct privileges can change another user's password. Ingres user passwords are not currently used to connect to the MDB.

**Important!** CA NSM components may not be able to connect if the Ingres user has been assigned a password.

An Ingres user can be set up with an expiration date. Once that date is past, the Ingres user cannot connect to Ingres until the expiration date is reset. Only an Ingres user with the correct privileges can reset the expiration date.

**Note:** Ingres user expiration dates are not currently used to connect to the MDB.

For security reasons, the Ingres user `mdbadmin` owns all database objects, does not have a corresponding operating system user ID, and should not be used by any application.

## Operating System Users

CA NSM prompts you to create a CA NSM administrator account with a password (Ingres databases) or a Microsoft SQL Server account with a password (Microsoft SQL Server databases). The default is `nsmadmin`.

**Note:** The WorldView Manager component must be installed on the MDB server computer.

### How You Create Additional CA NSM Administrators (Microsoft SQL Server Databases)

When you install the CA NSM Server component, you are prompted to create a CA NSM Microsoft SQL Server account with a password. You can create another user who will have CA NSM administrator privileges for the MDB.

1. Create a Microsoft SQL Server user with a password.
2. Assign the user to a default user role for the tablespace for which that user needs access.

**For example:**

```
sp_adduser 'nsmadmin', 'uniadmin'
```

## How You Create Additional CA NSM Administrators (Ingres Databases)

When you install the CA NSM Server component, you are prompted to create a CA NSM administrator account with a password. You can create another user who will have CA NSM administrator privileges for the MDB.

1. Create an operating system user with a password.
2. Create an Ingres user using the same name.
3. Assign the user to a default user group for the tablespace for which that user needs access.

### For example:

```
CREATE USER nsm_admin_user WITH group = uniadmin
```

**Note:** You can create an operating system user with a password expiration date, which may be a requirement for your organization. The Ingres VNODE entry on the client will not be able to connect to the server until the password entry for the VNODE is reset.

**Important!** For security reasons, do not create an operating system user called `mdbadmin`.

## Ingres Virtual Node Names (VNODES)

Remote connections to the MDB using Ingres use an Ingres VNODE.

During the CA NSM installation on UNIX and Linux operating systems, a VNODE called `nsm_servername` is created to provide connectivity to the MDB by daemons and other utilities running as root. Only processes running under the root ID have access to this VNODE.

### WorldView

For WorldView on Windows, the following naming standard applies:

`WV_servername`

Each logical repository name in WorldView, which is created by running the Define Logical Repository utility, has a VNODE created with the same name.

When WorldView attempts to connect to the MDB, a connection dialog appears, which prompts for a server name to connect to.

The WorldView registry is scanned to look for a valid VNODE for the server and user combination. If one is found, WorldView connects with that VNODE.

If the connection fails, the connection dialog prompts for a user ID and password, and from this information the VNODE is updated and the connection is attempted again. If the connection fails again, this cycle is repeated until the connection succeeds, or the user clicks Cancel. If the connection succeeds, the VNODE is saved for subsequent connections to WorldView.

When you are using the WorldView Classic GUI (Windows), the user ID and password you provide on the Repository Sign On dialog is saved and stored in the VNODE. When you start any additional WorldView component, such as Object Browser or Severity Browser, you are not prompted again for MDB credentials because the credentials saved in the VNODE are used.

For WorldView on UNIX and Linux, most WorldView components have input parameters that let you specify the server to connect to and the user name and password to use. Components that are run without specifying the server use the DefaultRepository registry entry, which is set at installation, to determine the server.

## Component-Level Security

Security at the product level helps keep unauthorized users from causing problems with key infrastructure components. Component level security focuses on improving the following two aspects of security:

- Unintentional problems caused by users having more access or authority than required to do their jobs
- Efficiencies that can be gained by having users presented only with the information required to do their jobs properly

CA NSM security provides about 100 rules, 9 roles (also known as user groups), and about 100 assettypes. CA NSM provides embedded security, which is a DENY mode security engine that, by default, is turned on. The following components use CA NSM embedded security:

- Calendar Management
- Embedded Security (protects itself)
- Job Management Option
- Alert Management
- Notifications Services
- Agent dashboards
- Web Reporting Service
- Unicenter Configuration Manager
- Event Management
- Management interfaces, which include Management Command Center, Unicenter Management Portal, Unicenter Browser Interface, and Unicenter for Pocket PC (logon only)

**Note:** Security is not installed by default, nor is it a selectable option. If you select any of the components that use security, the installation asks if you would like to enable security, and installs security if you answer "yes." For Windows, the question appears only when you are installing CA NSM in non-Express mode.

Without specific "permit" security rules for a given role or user, access to a component is denied. Default permit rules are created and activated for each of the components that uses embedded security for the following roles and types of access:

- Systems administrators (SYSADMIN) have *full* access to most components. By default, these users include only "administrator," "root," and the installing user.
- Network administrators (NETADMIN) have *full* access to most components. By default, these users include only "administrator" and "root."
- Operators (OPERATOR) have *read* access to most components. By default, these users include a "dummy" user for place-marker purposes.
- Application administrators (APPADMIN), database administrators (DBADMIN), mail administrators (MAILADMIN), web administrators (WEBADMIN), and business users (USER) have no users assigned and *no access* to most components. By default, these users include a "dummy" user for place-marker purposes.
- General users (PUBLIC) have no users assigned.

Exceptions to the above rules are as follows:

- The embedded security component allows full access for only systems administrators.
- Unicenter Management Portal access varies by role. For example, application administrators may have access that systems administrators do not.
- Logon access for most user interfaces, such as Unicenter Management Portal, Unicenter Configuration Manager, Unicenter Browser Interface, Management Command Center, and Unicenter for Pocket PC, is available for all roles. By default these roles include "administrator," "root," and the installing user.
- Windows embedded security does not provide granularity for some components, such as Calendar Management, Event Management, Security Management, and the Job Management Option. These components have access that is all or nothing, and therefore, systems administrator, network administrator, and operator roles have identical access.

## What is Security Management

Security Management provides a policy-based security system that protects against unauthorized access to vital CA NSM components and products.

To protect your management systems, you must augment physical security with software that can do the following:

- Prevent unauthorized individuals from gaining access (logging on) to your systems being managed by CA NSM.
- Ensure that only authorized personnel can access management data and resources.
- Protect against access abuse by users with administrative privileges.
- Provide a way to review and audit the use of data and resources.

Enhanced and simplified Security Management means reduced errors, greater responsiveness, and increased flexibility in meeting the needs of your organization. Most importantly, it means you can implement thorough and effective security policies without disrupting your work environment.

**Note:** The CA NSM Security Management components no longer provide file access authorization. If you need this type of additional security, you may want to evaluate eTrust Access Control. For more information, see Integration with eTrust Access Control.

## Administrators or Power Users Group

Certain CA NSM components and applications update configuration information. To run such components or applications, a user must be a member of the Windows Administrators or Power Users group, or root on UNIX/Linux. These components include the following:

- Classic Discovery (dscvrbe)  
**Note:** On UNIX and Linux, members of the Power Users group that have update, delete, insert, and select privileges can also run dscvrbe.
- 2D Map (catng2d)
- Repository Import/Export Utility (trix)
- Define Logical Repository utility (iirepdef)
- Most WorldView Classic GUI user interfaces

## How You Change the CA NSM Administrator Password On Windows

When CA NSM is installed on Windows, a prompt appears that lets you create a CA NSM operating system user ID and password.

To change the operating system password for this account, follow these steps:

1. Stop all the services on the WorldView server (the MDB server) and the remote server where Unicenter Services are running.
2. Change the Unicenter Administrator Password on the WorldView server by running the modp command.
3. Run the modp command on all servers that contain CA NSM services that connect to the WorldView tables in the MDB on the server where you changed the password.

**Note:** For more information about running modp, see the online *CA Reference*.

## Change the Password for the Severity Propagation Engine User Accounts (Windows)

When CA NSM is installed, the Severity Propagation Service is registered and a SeverityPropagation user account with a strong password is automatically created. A RunAs user account with the same password is also added to the dcomcfg utility. These user IDs are created so that the Severity Propagation Engine can stay connected when the user logs off of the computer.

You may want to change the password for these user accounts for security reasons. To do this, you must deregister the Severity Propagation Service and re-register it with a new password.

**Important!** Failure to deregister and re-register the Severity Propagation Service correctly will result in a catastrophic failure of many CA NSM subsystems. Any errors that occur during registration and deregistration are written to the application event log in the operating system's event viewer.

#### **To change the password for the SeverityPropagation and RunAs user accounts**

1. Stop the Severity Propagation Service (sevprop.exe) using the Windows Service Manager.
2. Run the following command from a command line:  

```
sevpropcom /unregister
```

The Severity Propagation Service is removed from the dcomcfg utility and the SeverityPropagation user account is removed.
3. Run the following command from a command line:  

```
sevpropcom /regserver
```
4. The Severity Propagation Service is re-registered and the SeverityPropagation user account is created with a proprietary password. The password conforms to Microsoft's most rigorous password complexity methodology, using Microsoft's LSA policy to ensure the security of the password.  
**Note:** You can use the `sevpropcom /regserver /password` command to register the DCOM server with a user-defined password. You must ensure that all password requirements are met if you enter your own password.
5. Start the Severity Propagation Service (sevprop.exe) using the Windows Service Manager.

## **How You Change File Privileges on Microsoft Windows 2003 Server**

Microsoft Windows 2003 Server has strict security controls in place for user accounts other than the user under which a product was installed. If you installed CA NSM while logged on as a particular user, for example, administrator, but want to run CA NSM using another user account, you may need to allow write permissions to the other user account from the administrator account.

1. Log in as the user that installed CA NSM, for example administrator.
2. Use Windows Explorer and navigate to the folder where you installed CA NSM.
3. Right-click the folder and select Properties.
4. Select the Security tab.



5. Click add, enter TNDUsers, and click OK.
6. In the Permissions field at the bottom of the Properties dialog, click the Allow box after the Write permission, and click OK.
7. Log off as administrator and log back in as the user that is a member of the TNDUsers group.

## Run Utilities Requiring Administrator Privileges on Windows Vista

By default, Windows Vista gives administrators only standard user privileges through its User Account Control. To install CA NSM and run most CA NSM commands, admin privileges are required; therefore, you may receive an 'Access is Denied' message when trying to run utility commands from a command prompt on Windows Vista even if you are an administrator.

To override User Access Control without disabling it completely and obtain admin level privileges required to run CA NSM utilities, we recommend launching a command prompt with "Run as Administrator" on Windows Vista systems.

### **To run utilities requiring administrator privileges on Windows Vista**

1. Right-click the Command Prompt option in the Start menu and select "Run as administrator".

The User Account Control dialog opens asking you to confirm the action.

2. Click Continue.

The command prompt appears with Administrator: Command Prompt in the title.

You can run all CA NSM utilities requiring administrator privileges in this command prompt.

## Create Additional Users with Administrator Privileges to Run Discovery (Microsoft SQL Server Databases)

Only nsmadmin, or the CA NSM administrator that was used to install the Ingres server on the MDB server, can run Discovery after CA NSM is installed. You may want to give other administrative users authority to run Discovery.

### To create a user *with* administrator privileges

1. On the MDB server, create a Microsoft SQL Server user by running the following command:

```
addntgroup -a "TNDUsers" -s repository_name -u "userid" -p "password" -b mdb -g uniadmin
```

The Microsoft SQL Server user is created as a member of the uniadmin role.

2. Run the following command using the nsmadmin user and password:

```
modp -r repository_name -u nsmadmin -n nsmadmin_password
```

**Note:** You only need to run the modp command if Discovery is run on a new remote MDB, that is, a different MDB than the one used during installation.

The user ID you created has the authority to run Discovery.

## How You Create Additional Users Without Administrator Privileges (SQL Server Databases)

Only nsmadmin and the install user (usually Administrator) can run Discovery after CA NSM is installed. You may want to give other users authority to run Discovery without giving them administrator privileges.

To create a user *without* administrator privileges, follow these steps:

1. Manually create a Windows user and add it to the TNDUsers group.
2. Manually create a Microsoft SQL Server user with SQL Enterprise Manager with uniadmin as its default role.
3. Modify the security permissions of the Program Files\CA\SharedComponents\CCS\Discovery folder to allow users of the TNDUsers group to modify, read and execute, list folder contents, and to have read and write access.
4. Run the modp command using the nsmadmin user and password.

**Note:** For more information about the modp command, see the online *CA Reference*.

## Create Additional Users with Administrator Privileges to Run Discovery (Ingres Databases)

Only nsmadmin, or the CA NSM administrator that was used to install the Ingres server on the MDB server, can run Discovery after CA NSM is installed. You may want to give other administrative users authority to run Discovery.

**To create a user *with* administrator privileges**

1. On the MDB server, create a Windows and Ingres user by running the following command:

```
addntgroup -a "TNDUsers" -s repository_name -u "userid" -p "password" -b TNGDB  
-g uniadmin
```

The Windows user is created as a member of the TNDUsers group and the Ingres user is created as a member of the uniadmin group.

2. Run the following command using the nsmadmin user and password:

```
modp -r repository_name -u nsmadmin -n nsmadmin_password
```

**Note:** You only need to run the modp command if Discovery is run on a new remote MDB, that is, a different MDB than the one used during installation.

The user ID you created has the authority to run Discovery.

**Note:** For more information about the addntgroup and modp commands, see the online *CA Reference*.

## How You Create Additional Users Without Administrator Privileges (Ingres Databases)

Only nsmadmin and the install user (usually Administrator) can run Discovery after CA NSM is installed. You may want to give other users authority to run Discovery without giving them administrator privileges.

To create a user *without* administrator privileges, follow these steps:

1. Manually create a Windows user and add it to the TNDUsers group.
2. Manually create an Ingres user VisualDBA with uniadmin as its default group.
3. Modify the security permissions of the Program Files\CA\SharedComponents\CCS\Discovery folder to allow users of the TNDUsers group to modify, read and execute, list folder contents, and to have read and write access.
4. Run the modp command using the nsmadmin user and password.

**Note:** For more information about the modp command, see the online *CA Reference*.

## WorldView Security

The topics that follow explain WorldView security considerations.

## Set up Read-Only Windows-Authenticated Users (Microsoft SQL Server Databases)

You can set up read-only users for the WorldView tables in the MDB. A read-only user is not permitted to perform any operations in Unicenter MCC or the WorldView Classic GUI that may update the WorldView MDB data.

**Note:** Using this procedure to apply read-only access affects only WorldView data in the MDB. It does not affect DSM, Enterprise Management, or other data providers.

### To set up read-only Windows-authenticated users

1. Create a Windows group called TNDReadOnly with only two operating system rights: Logon as a Batch Job, and Replace a Process Level Token.  
**Note:** Use the Local Security Policy GUI to set up the operating systems rights.
2. Define the TNDReadOnly group to Microsoft SQL Server using the Enterprise Manager, and assign the wvuser role to this group.
3. Add any Windows users that you want to have read-only permissions for Unicenter MCC to the TNDReadOnly group.

**Important!** Do not add these Windows users to the TNDUsers group.

The user has read-only permission for WorldView data.

## Set Up Read-Only Microsoft SQL Users for WorldView

You can set up read-only users for the WorldView tables in the MDB. A read-only user is not permitted to perform any operations in Unicenter MCC or the WorldView Classic GUI that may update the WorldView MDB data.

**Note:** Using this procedure to apply read-only access affects only WorldView data in the MDB. It does not affect DSM, Enterprise Management, or other data providers.

### To set up read-only Microsoft SQL Server users

1. Create a Microsoft SQL Server user that will be the WorldView read-only user, for example, wvreadonly.
2. Assign the Microsoft SQL Server user to the database role of wvuser.

The user has read-only permission for WorldView data.

## Set Up Read-Only Users for WorldView (Ingres Databases)

You can set up read-only users in Ingres for the WorldView tables in the MDB when Data Scoping is not active. A read-only user is not permitted to perform any operations in the Management Command Center and the WorldView Classic GUI that may update the WorldView MDB data.

**Note:** Using this procedure to apply read-only access affects only WorldView data in the MDB. It does not apply the same access to DSM, Enterprise Management, or other data providers.

### To set up read-only users for WorldView

1. Create an operating system user that will be the WorldView read-only user, for example, wvreadonly.
2. (Windows only) Assign the operating system user to the TNDUsers operating system group if you want the user to have access to WorldView data using the Management Command Center.
3. Add an Ingres user of the same name (that is, wvreadonly) to Ingres and assign it to the default group of wvuser.

**Note:** On UNIX/Linux, you can use the `add_ingres_user` script in the `$CAIGLBL0000/wv/script` directory to do this automatically.

The user has read-only permission for WorldView data.

**Note:** Do not use the database security permissions of a particular user for implementing read-only users when Data Scoping is active.

## Connect Remotely to Another MDB Using WorldView Classic (Windows)

To connect remotely to another MDB using WorldView Classic, you must connect to a logical repository. If a logical repository does not exist, you must first define one. See Define a Logical Repository.

### To connect remotely to another MDB

1. Click Start, Programs, CA, Unicenter, NSM, WorldView, and select the name of the component you want to start (2D Map, Object Browser, Class Browser, and so forth.)

The Select Repository dialog appears.

2. Select the name of the logical repository you want to connect to.

**Note:** If the name does not appear in the drop-down, click Find and select the name, or type the name of the logical repository.

You are connected to the logical repository, and the WorldView Classic GUI component opens.

**Note:** When you start the 2D Map using the `catng2d` command, use the `/R` parameter to specify the logical repository. Do not use the `/U` and `/P` parameters for Ingres connections if you are using an Ingres database.

## Connect to a Remote Repository

You may be responsible for managing multiple installations of CA NSM and may need to connect to a remote MDB to run CA NSM applications that update the remote MDB.

### To connect to a remote repository

1. From the CA NSM client computer, [define a logical repository](#) (see page 55) for the MDB.

**Note:** When creating the logical repository you must know the Administrator account for the remote server that was defined, in addition to the password.

2. (Windows only) If you installed management components, run the `modp` command using the Administrator account and password for the remote server and the name of the logical repository that you defined:

```
modp -r repository_name -u userid -n password
```

### Example: Connect to a Remote Repository

Unixp is a CA NSM client computer, and uswv01 is the name of the WorldView server where the MDB resides. On unixp, define a logical repository named uswv01a to associate with the MDB on uswv01 using the nsmadmin user ID and password for uswv01. If unixp contains CA NSM management components, run `modp` to define the nsmadmin user ID and password for uswv01. You can now connect to uswv01a and run WorldView and Discovery applications (Discovery is a management component) from unixp and the data is stored in the MDB on uswv01.

## Define a Logical Repository (Windows)

Before you can connect remotely to an MDB, you must define a logical repository.

### To define a logical repository

1. Click Start, Programs, CA, Unicenter, NSM, WorldView, Define Logical Repository.

The CA NSM Repository Creation wizard appears.

**Note:** You can also run the `iirepdef` command or click Define on the Select Repository dialog to start the CA NSM Repository Creation wizard.

2. Enter a logical name to associate with the MDB on the server to which you want to connect, and click Next.

The Access Type page appears.

3. Select CA Ingres (Ingres Databases) or SQL Server (Microsoft SQL Server databases) and click Next.

The Server Name page appears.

4. Enter information in the following fields, and click Next.

#### **Server Name**

Specifies the name of the MDB server, which must already exist.

#### **Server User (Ingres databases only)**

Specifies the name of the CA NSM user ID for access to the MDB.

#### **Server Password (Ingres databases only)**

Specifies the password for the Server User.

#### **Server Installation Code (Ingres databases only)**

Specifies the name of the Ingres Installation ID used when the MDB was installed. The default is EI.

#### **Instance Name (Microsoft SQL Server databases only)**

Specifies the name of the Microsoft SQL Server instance used when the MDB was installed.

5. Click Define Repository, and click Finish.

You are connected to the MDB, and the logical repository is defined.

**Note:** If the connection to the MDB fails, you receive an error and the wizard reappears. Typically, this signifies that you do not have the proper credentials to connect to the MDB. For the proper credentials, see your CA NSM administrator.

## Define a Logical Repository (UNIX/Linux)

Before you can connect remotely to an MDB, you must define a logical repository.

To define a logical repository, run the `iirepdef` command. For more information about the `iirepdef` command, see the online *CA Reference*.

## Management Command Center Security

The topics that follow explain Management Command Center security.

### Access to the Management Command Center

In a typical installation of CA NSM, the Management Command Center is run remotely from a client computer. Since the MCC is accessing information about Unicenter Manager servers, you must supply a user ID and password for that type of manager before you can access any of the network information using the Management Command Center. The following is a list of things to consider when assigning user IDs to access Management Command Center:

- For the Topology and Tools navigation left hand pane, the TNDUsers group (on the WorldView Manager server) is used to authenticate Windows user IDs when you use the Management Command Center remotely.

The TNDUsers group is added automatically to your Windows computer when you install CA NSM (specifically, the WorldView Manager component) on the MDB server. If you use a Windows user ID in the login security dialog to log on to the Management Command Center, that Windows user ID must be a member of the TNDUsers group on the computer that contains the MDB. You can use the Windows User Manager to add this user ID to the Windows Group TNDUsers.

- Domain and local groups are supported when using the Management Command Center on Windows for the Topology and the Tools navigation left hand pane.

You must enter the domain account in the form "Domain1\User1" in the Management Command Center login security dialog, regardless of where Management Command Center is running. The domain of the computer where the MDB resides must be the same domain or trust the domain of the domain account used for authentication. For example, if you enter "Domain1\User1" in the Management Command Center login security dialog, the computer that contains the MDB must be logged into Domain1, or DomainX, where DomainX trusts Domain1. If you enter the account in the form "User1," then it is considered a local user account defined to the computer where Management Command Center is running. The domain user must be a member of the TNDUsers group, either directly or indirectly through a domain group.



- Different domains with the same user ID are considered the same for Ingres (Ingres databases only).

For example, if a user logged into the client computer as DomainA\joe, the user is authenticated to Ingres as joe, not DomainA\joe because Ingres does not support domain accounts.

- You must always use a local operating system user ID to authenticate to an Ingres database.

This account can be defined to an Ingres server using the `accessdb` utility. This operating system ID must also have access to the MDB. You can do this on the server where the Ingres database resides.

Each Ingres user must be defined to a default user group. For WorldView access that has full authorization, assign the default group of `wvadmin` (or `uniadmin` for all CA NSM tables). For users that should only have only read authorization, assign the default group of `wvuser` (or `uniuser` for all product tables).

- You must always use an SQL Server user ID or Windows-authenticated user ID to authenticate to an SQL Server database.

This account can be defined to SQL Server using the SQL Enterprise Manager utility. This user ID must also have access to the MDB. You can do this on the server where the SQL Server database resides.

Each SQL Server user must be defined to a default user role. For WorldView access that has full authorization, assign the default role of `wvadmin` (or `uniadmin` for all CA NSM tables). For users that should only have only read authorization, assign the default role of `wvuser` (or `uniuser` for all product tables).

## Grant Non-Root Access to the Management Command Center

By default, the Management Command Center runs successfully only under the root user ID on UNIX/Linux.

### To give non-root users access to the Management Command Center

1. Enter the following commands at a command prompt:

```
chmod 777 $JI_SYSTEM/files
chmod 777 $JI_SYSTEM/logs
chmod 666 $JI_SYSTEM/files/catalog.*
chmod 666 $JI_SYSTEM/files/ji.cfg
chmod 666 $JI_SYSTEM/files/jilog.log
chmod 666 $JI_SYSTEM/./wrp2cat.log
```

All non-root users have access to the Management Command Center.

2. (Optional) To allow non-root users to access WorldView data stored in the MDB, enter the following commands:

```
chmod 777 $CAIGLBL0000/wv/config
$CAIGLBL0000/wv/scripts/add_ingres_user non_root_user_name
```

All non-root users have access to WorldView data in the Management Command Center.

## Override the Management Command Center User ID

The Management Command Center maintains separate storage areas for its saved settings and bookmarks, indexed by user ID. When multiple users are running the Management Command Center while logged in to the same account (for example, UNIX/Linux root), you may want to override the user ID on the command line or provide a customized copy of the batch file (camcc on UNIX/Linux; tndbrowser.bat on Windows) for each user.

To override the Management Command Center user ID, append `-profile:xyz` to the shell command, where `xyz` is any unique string to use in place of the default user ID.

When you connect to a provider, the user ID and password are stored so that you do not need to supply this information the next time you use the Management Command Center.

**Note:** If you want to login as someone else, be sure that password caching is deactivated and, delete the password cache file from the `%JI_SYSTEM%\files` directory.

## Deactivate Password Caching

The password cache file is used to save the user ID and password credentials that were used successfully to access a namespace. Therefore, the next time the Management Command Center is started, the saved credentials are used to access the namespace, and if they are still valid, the user does not have to enter them again. The saved credentials are shared by all users of the Management Command Center on the same system, even if they use different profiles. Thus, password caching is not desirable if multiple users are running the Management Command Center on the same system.

To deactivate password caching, enter the following line into the `ji.cfg` file:

```
default.SEC.bypass_cache: 1
```

Passwords are not cached.

## Integrating with eTrust Access Control

CA NSM Security Management no longer provides security for logon authentication or for files. If you need this type of additional security, you may want to integrate CA NSM with eTrust Access Control (eTrust AC). eTrust AC protects enterprise assets including CA NSM assets, file-based assets, and authenticates logons.

You can migrate CA NSM 3.0 or 3.1 security to eTrust AC. When you install CA NSM r11.2, the installation process detects whether CA NSM 3.0 or 3.1 is enabled to use eTrust AC. If CA NSM 3.0 or 3.1 is enabled to use eTrust AC, CA NSM Security Management is not installed. Alternatively, CA NSM r11.2 includes a lightweight security engine for the protection of CA NSM assets only, which is provided if you do not want to install and migrate to eTrust AC.

**Note:** For complete information about CA NSM security migration and integration with eTrust AC, see the *eTrust Access Control Getting Started* guides, the *eTrust Access Control Implementation Guide*, and the *eTrust Access Control Administrator Guides*.

## How Integration and Migration Works

eTrust AC r8 provides automatic integration with CA NSM using options when you install eTrust AC. You must install eTrust AC r8 before you attempt to migrate CA NSM 3.x security. After you migrate CA NSM 3.x security, you can install CA NSM r11.x.

- On Windows, eTrust AC automatically migrates CA NSM 3.x data during eTrust AC installation if the "Migrate Unicenter Security data to eTrust Access Control" option is selected.
- On UNIX and Linux, you must install the Unicenter Integration and Migration package and then run the migration scripts.
- Once the integration and migration process completes successfully, CA NSM login intercepts are disabled (but not removed), and no longer used. However, logon and file protection intercepts are removed when you migrate CA NSM 3.x to CA NSM r11.x.
- After migration on Windows, CA NSM 3.x Security Management is disabled and can no longer be started.
- After migration on UNIX and Linux, a gateway daemon called *sessfgate* is installed and activated. *Sessfgate* processes the CA NSM security API (EMSec) requests sent through the CA NSM security message queue and routes these reformatted and rerouted requests to eTrust AC. eTrust AC return codes are translated to the CA NSM equivalents. This process protects the integrity of existing applications that are currently using the EMSec API.

If the CA Event Notification Facility (CAIENF) is running, *sessfgate* is automatically started or stopped whenever the eTrust AC services are started or stopped. If CAIENF is not running, eTrust AC does not start the *sessfgate* daemon. You must start it manually.

**Note:** After CA NSM Security Management is disabled and eTrust AC is started, the *sessfgate* daemon can accept API requests instead of the Secure Sockets Facility (SSF).

- On Windows, the EMSec API requests channel calls into the CAUSECR.DLL. This DLL is replaced during the integration process. The new DLL receives calls to the EMSec APIs, and then reformats and redirects these requests to equivalent eTrust AC APIs. The return code from the eTrust AC APIs are converted back into their corresponding EMSec API return codes and control is returned to the caller of the EMSec API. This approach protects the integrity of existing applications that are currently using the EMSec API.

- CA NSM 3.x asset types are pre-loaded during eTrust AC r8 installation. After the migration, CA NSM 3.x assets are protected by eTrust AC.
- eTrust AC provides programs that extract data from the CA NSM 3.x Security database and translate it into eTrust AC commands that populate the eTrust AC database. The following data is migrated:
  - CA NSM Security users
  - CA NSM Security user groups
  - CA NSM Security rules
  - CA NSM Security assetgroups

## Rules and Statistics Not Migrated

The following rules and statistics are not migrated from CA NSM 3.x security to eTrust AC:

- CA NSM Data Scoping, Command Scoping, and Keyword Scoping rules (rules that apply to CA NSM asset types with a -DT, -CM or -KW suffix. Rules of this type are ignored during the migration process.
- CA NSM Security rules for any of the following Unicenter Security asset types are obsolete because CA NSM Security is no longer used:
  - CA-USER
  - CA-ACCESS
  - CA-USERGROUP
  - CA-ASSETGROUP
  - CA-ASSETTYPE
  - CA-UPSNODE

Rules that apply to any of these asset types, or any of their derivatives, are ignored during the migration process.

- Creation and modification statistics for all CA NSM objects are lost in the migration process.

## Attributes Not Migrated

Due to CA NSM and eTrust Access Control product differences, the following CA NSM Security attributes cannot be migrated to eTrust AC:

For CA NSM Security users, the following attributes cannot be migrated:

- Statistics:
  - Last login (date and time, node of last login)
  - Password change (date and time, node, user who changed last password, and expiration date of the password)
  - Password violation (date and time, node of last unsuccessful login, and number of unsuccessful logins since last successful login)
  - Access violation (date and time, node of last access violation, and number of access violations)
  - Suspension (date and time of suspension)
  - PWDCHANGE VALUE (RANDOM)-Random password generation.
  - UPSSTATGROUP--UPS station group is not supported by eTrust Access Control.
  - VIOLMODE-Violation mode (FAIL, MONITOR, WARN, QUIET). eTrust Access Control supports FAIL mode only
  - VIOLACTION-Violation action (CANUSER, CANU&LOG, CANU&LOG&SUS). eTrust Access Control supports CANUSER action only.

For CA NSM Security rules, the following attributes cannot be migrated:

- EXPIRES-Rule expiration date is not supported by eTrust AC.

## Protecting and Filtering MDB Data Using Data Scoping

WorldView Data Scoping lets you protect the WorldView MDB object data from unauthorized access. Data Scoping rules let the MDB administrators control a specific user ID's access to a particular WorldView data object or group of WorldView data objects. Data Scoping is intended to provide protection against both malicious and non-malicious access.

Data Scoping also lets you filter large amount of information contained in the MDB into smaller, more pertinent subsets of data. It lets you view the same data differently depending on the current task.

## Data Scoping Rules

To implement Data Scoping, you must define a set of Data Scoping rules. Each set of Data Scoping rules governs data access only for the MDB for which they are defined. If more than one MDB exists on a computer, each MDB has its own set of autonomous rules. Further, if a user of a single computer connects to more than one MDB, each MDB has its own independent set of Data Scoping rules for that user.

Data Scoping rules can control the type of access that is used. You can govern all access to a particular data type, or you can give a specific user ID read access without giving that same user update, create, or delete capabilities.

You can activate and deactivate Data Scoping rules according to the current date at the time of access by defining a Data Scoping rule with an effective date or an expiration date. If Enterprise Management is installed, you can also activate and deactivate a Data Scoping rule by specifying that it use a particular calendar.

By default, users connected to MDB have full access to objects until Data Scoping rules are generated to deny particular types of access. You can define class-level rules or object-level rules. Class-level rules scope data objects by their data classification. Object-level rules let you explicitly filter data objects on an object-by-object basis using the object's instance-level properties.

On Windows, Data Scoping rules are supported for local users, local groups, domain groups, and domain users.

On UNIX/Linux, only local users are supported. All references to MDB refer to the local database.

**Note:** Data Scoping rules do not affect the Discovery process. All discovered devices are added to MDB. If you do not want a particular user to see these discovered devices, you can create rules for those users that deny read access.

## Types of Data Scoping Rules

You can define the following two levels of Data Scoping rules:

### **Class-level Rules**

Class-level rules scope data objects by their data classification. Class-level rules let you restrict access to an entire class of objects, rather than individual objects. Because object-by-object evaluation is not required to support this type of Data Scoping rule, defining class-level rules is the most efficient means of scoping data objects. You can only create rules on classes inherited from the ManagedObject or Association class.

For example, you could define a Data Scoping rule to implement a statement like "User ID JOE cannot access objects of the Payroll class." When a request for access to the Payroll class is evaluated, it determines whether the requesting user ID is JOE and the entire request is immediately allowed or denied.

### **Object-level Rules**

Object-level rules let you explicitly filter data objects on an object-by-object basis using the object's instance-level properties. Object-level rules target a specific object or group of objects within a class.

For example, you could define a Data Scoping rule to implement a statement like "User ID JOE cannot access Payroll objects that were created by user ID MARY." A rule like this requires an object-by-object analysis of each of the Payroll objects accessed by JOE to determine if the createuser instance-level property has a value of MARY.



## How Data Scoping Rules are Inherited

Data Scoping rules are inherited in the following two ways:

- **Class inheritance rules** are rules defined for a particular class that are inherited by all its subclasses. Rules defined for a particular class override rules from any of its parent classes.

If no rule is defined for a particular class, the immediate parent overrides any other rules in the class hierarchy. Thus, if you deny Delete access for the ManagedObject class, but allow Delete access for Host, Windows Server computers and Windows 2000 Server, computers can be deleted but a Crossroads\_Bridge cannot be deleted.

- **Inclusion inheritance rules** are rules defined for a particular class that are propagated to the child objects of that class.

Also, because a Business Process View object has two parents, both of the parents' rules are checked.

If there is a rule for the child object, it overrides the rule for the parent object. Thus, a rule defined for a particular subnet has that rule take effect on all objects within that subnet. For example, if you cannot delete subnet a.b.c.0, you are also not able to delete computer a.b.c.14 or a.b.c.16.

If rules are defined for a parent object and a parent class, the class inheritance rules takes precedence over the inclusion inheritance rules. The inclusion inheritance rules are evaluated only if the class inheritance rules do not apply.

Rule propagation is useful for administrating entire subnets or all objects related to a device. If you deny Delete access for Windows computer ABCD, then any agents, Enterprise Management components, or WBEM object for that device cannot be deleted. You do not need to define separate rules.

## How Data Scoping Order of Precedence Rules Are Applied

Data Scoping rules can be applied to users and local or domain groups, and target a specific class or superclass (if subclasses are present). The order of precedence rules are applied only when there are multiple rules that conflict, that is, one rule denies access and another allows access for the same operation.

The order of rule precedence is as follows:

- Object-level rules take precedence over class-level rules. If object-level rules conflict, the Allow rule takes precedence.
- If class-level rules conflict, the Allow rule always takes precedence.

- This same concept applies to class generation level. Class generation level 0 implies that the rule is tied to the current class and has precedence over all other class generation levels. Class generation level 1 implies that the rule is tied to a class that is a direct superclass of the current class.

Rules for a class always take precedence over superclass rules regardless of whether the rule is a user rule or group rule.

- An Allow object-level rule takes precedence over any other object-level rule regardless of whether that other rule is a user or group rule. An Allow class-level rule takes precedence over any other class-level rule whether that other rule is a user or group rule. An Allow class-level rule **never** takes precedence over any object-level rule.
- If there is no specific class-level rule or object-level for a class or any of its superclasses, the inclusion hierarchy is used for Data Scoping evaluation. That is, the topology of the network is used for evaluation. Rules are evaluated for the parent folder of an object. If the rule applies to the parent, it applies to all children within the folder. If the parent folder has no rule that applies, its grandparent is searched, then its great grandparent, and so forth. Objects that are in multiple folders where one folder has a Deny rule, and the other an Allow, the Allow takes precedence.
- User rules and group rules are treated equally. Any Allow rule in either category takes precedence.

#### **Examples: Data Scoping Order of Rule Preference**

For the following two rules, Rule 2 takes precedence:

- Rule 1: Deny Delete on Windows for User1.
- Rule 2: Allow Delete on Windows for Group1 where User1 is a member of Group1.

For the following two rules, Rule 1 takes precedence:

- Rule1: Deny Delete on Windows where the label equals "Computer1" for User1.
- Rule2: Allow Delete on Windows for User1.

For the following two rules, Rule 2 takes precedence:

- Rule1: Deny Delete on Windows for group1 where User1 is a member of Group1.
- Rule2: Allow Delete on Windows for User1.

For the following two rules, Rule 1 takes precedence:

- Rule1: Deny Delete on Windows for User1.
- Rule2: Allow Delete on ManagedObject for User1.

For the following rule, Computer1 in subnet 192.168.255.0 is denied access for delete:

Rule: Deny Delete on IP\_Subnet where Name="192.168.255.0."

For the following rules, and if Computer1 is in subnet 192.168.255.0 and class Windows, Rule2 takes precedence.

- Rule1: Allow Delete on IP\_Subnet where name="192.168.255.0."
- Rule2: Deny Delete on Windows where Name="Computer1."

For the following two rules, and if Computer1 is on subnet 192.168.255.0, Rule 2 takes precedence:

- Rule1: Deny Delete on IP\_Subnet where Name="192.168.255.0."
- Rule2: Allow Delete on Business Process View BPV1 where Computer1 is in BPV1.

For the following three rules, Rule1 takes precedence over Rule3 for Computer1 in BPV Atlas:

- Rule1: Class = ManagedObject /Deny/create+update+delete/User=user1
- Rule2: BPV/Deny/All/User=User1
- Rule3 BPV/Allow/Read+Update/User=User1/Name=Atlas

Since Rule1 is a class inheritance rule, it takes precedence for objects within the BPV named Atlas over the inclusion inheritance rule Rule3. To allow update access for all objects within the BPV Atlas, Rule1 should be changed to the following:

Rule: Class=ManagedObjectRoot/Deny/Create+Update+Delete/  
Name=ManagedObjectRoot/User=user1

## Rule Performance Issues

Data Scoping rule evaluation can significantly degrade performance because of inclusion inheritance rules. For every object where access is evaluated, all ancestors of that object are also evaluated if no Data Scoping rules apply to that particular object.

The performance degradation is caused by the object-based rules defined in the class hierarchy. This is due to the fact that an object must be constructed by issuing SQL queries as the topology is traversed. Limiting object-based rules at a lower level in the class hierarchy can reduce this database overhead.

For example, if the following rule is defined on the class ManagedObject, and the managed objects are Windows computers, then the rule should be defined on the Workstation class.

Rule1: Deny delete for address=192.168.34.45 or address=192.168.34.46

## How Data Scoping Rules Impact MDB Performance

Data Scoping has minimal impact on MDB performance. No data scoping solution can be completely free of impact on performance; however, the impact on MDB performance varies on a site-by-site basis relative to how you use the Data Scoping feature.

- You are protected against unnecessary performance degradation by both flexible Data Scoping rule syntax and a Data Scoping evaluation cache that holds the most frequently read Data Scoping rules.
- Data Scoping rules are stored in the MDB and evaluation requires additional I/O to retrieve them. To reduce the number of additional I/O requests and their impact on MDB performance, a Data Scoping evaluation cache holds the most frequently used Data Scoping rules in memory. When an incoming MDB request is evaluated for Data Scoping, the evaluator determines the requesting user ID. If the requesting user ID is not already in its cache memory, the evaluator loads the cache with all of the Data Scoping rules that correspond to the user ID.
- The mechanism described previously limits additional Data Scoping I/O to the first MDB request by a user ID. Synchronization of the cache with Data Scoping rule updates is critical and verified on a request-by-request basis.

- If we assume that the user ID is the same user ID that is used to connect to the MDB, the impact of Data Scoping rules on MDB performance can be summarized as follows:
  - When Data Scoping rules are not used, there is no performance impact on the MDB functionality.
  - When a user ID has no Data Scoping rules applied to it, there is no performance degradation for any MDB requests after the first data access.
  - When a user ID has Data Scoping rules that apply to it at a class level, there is no performance degradation for any MDB requests to those classes with no rules applied. Any performance degradation of access is limited to those classes with rules applied and is negligible.
  - When a user ID has Data Scoping rules that define object-level overrides (thus requiring property-by-property analysis), there is no performance degradation of MDB requests to those classes for which there are no rules. The impact on performance is limited to those requests that target a class for which there is a rule.

## Data Scoping Security on Windows

Data Scoping security is designed for non-administrator accounts; that is, for Windows user IDs that are not part of the Windows Administrators Group. While you can create Data Scoping rules for Administrator accounts, and those rules will be enforced, any administrator can delete those rules. Effectively, there is no real Data Scoping security for administrators.

## Data Scoping Security on UNIX/Linux

Data Scoping security on UNIX and Linux computers is designed for non-root accounts only.

## User IDs Required for Data Scoping Rule Evaluations

The user ID that is used for Data Scoping rule evaluation depends on your database, the database configuration, and the type of application you are running.

### User IDs for Data Scoping Evaluation on Windows Platforms (Microsoft SQL Server Databases)

On a local MDB, the user ID that is used to connect to Microsoft SQL Server is used for Data Scoping evaluation.

For a remote MDB, a Microsoft SQL Server user ID or Windows user ID defined on the remote computer is used for Data Scoping evaluation. You can enter the user ID in the Login Security dialog that is accessible from the Management Command Center.

When the Logon Security dialog for remote connections appears, use the following user IDs:

- Microsoft SQL Server user ID
- Windows user ID. Use the User Manager to add this user ID to the Windows Group TNDUsers. TNDUsers is created during CA NSM installation and has all the necessary Windows privileges. You need only to add a Data Scoping user to this group before they can sign on to CA NSM using the Windows user ID.

### User IDs for Data Scoping Evaluation on Windows Platforms (Ingres Databases)

On a local MDB, the currently logged on Windows user ID is used for Data Scoping evaluation.

For a remote MDB, a Windows user ID defined on the remote computer is used for Data Scoping evaluation. You can enter the Windows user ID in the Login Security dialog that is accessible from the Management Command Center.

When the Logon Security dialog for remote connections appears, use the following user IDs:

- Ingres user ID
- Windows user ID. Use the User Manager to add this user ID to the Windows Group TNDUsers. TNDUsers is created during CA NSM installation and has all the necessary Windows privileges. You need only to add a Data Scoping user to this group before they can sign on to CA NSM using the Windows user ID.

### User IDs for Data Scoping Rule Evaluation on UNIX/Linux Platforms

Different user IDs are used for Data Scoping rule evaluation, depending on where your MDB resides:

- For a local MDB, the currently logged on UNIX/Linux user is used for Data Scoping rule evaluation.
- For a remote MDB, a UNIX/Linux user ID defined on the remote computer is used for Data Scoping rule evaluation. You can enter the UNIX/Linux user ID from the logon security dialog from the Management Command Center.

### Data Scoping Rule Evaluation Using Windows Local Groups

You can create Data Scoping rules for local group accounts defined on the local machine. Rules are applied in the following ways:

- If a rule exists for a local group account and the local user account who is authenticated is a member of that local group, the rule applies to that user.
- If there are rules for multiple local groups and the local user who is authenticated is a member of those local groups then, all rules apply.
- If local group rules conflict, an Allow rule takes precedence.
- If rules exist for the specific local user account, these rules are treated the same as any rules that exist for local Groups in which the local user is a member.

### Data Scoping Rule Evaluation Using Windows Domain Groups (Microsoft SQL Server Databases)

Microsoft SQL Server supports Windows domain accounts for authentication.

Data Scoping rules are enforced for domain groups in which the particular user is a member. You can create rules for multiple domains on one MDB using the DataScope Rule Editor. You can create rules when logged into different domains by using the DataScope Rule Editor locally or remotely. Only the rules created for the domain that is used to authenticate Windows to the MDB are applied.

You can then create Data Scoping rules for domain group accounts defined on the domain that is currently logged in. Rules are applied in the following ways:

- If a rule exists for a domain group account and the domain user who is authenticated is a member of that domain group, the rule applies to that user.
- If rules are defined for multiple domain groups and the domain user who is authenticated is a member of those domain groups, then all rules apply.
- If the domain user is a member of a domain group or local group for which a rule exists, or if the domain user is a member of a domain group that is a member of a local group **and** a rule for the local group exists, the rule applies.

Data Scoping rule evaluation takes place as described for a local computer.

## Data Scoping Rule Evaluation Using Windows Domain Groups (Ingres Databases)

Ingres does not support Windows domain accounts for authentication. However, domain support for Data Scoping exists over Ingres in that the domain into which the user is logged on at the client computer is used for Data Scoping rule evaluation.

For example, if a user is logged into a domain on a client computer, that domain user is used for Data Scoping rule evaluation. If a user logged into the client computer as DomainA\joe, the user is authenticated to Ingres as joe (not DomainA\joe because Ingres does not support domain accounts) but DomainA\joe is used for Data Scoping rule evaluation, regardless of whether the Ingres database is local or remote.

If the user is logged into a different client computer, such as DomainB\joe, the user is still authenticated to Ingres as joe but DomainB\joe is used for Data Scoping rule evaluation. Thus, two different client computers connected to the same server are authenticated to Ingres using the same user ID (joe) but two different domain user accounts are used for Data Scoping rule evaluation.

Data Scoping rules are enforced for domain groups in which the particular user is a member. You can create rules for multiple domains on one MDB using the DataScope Rule Editor. You can create rules when logged into different domains by using the DataScope Rule Editor locally or remotely. Only the rules created for the domain that is used to Windows-authenticate to the MDB are applied.

You can then create Data Scoping rules for domain group accounts defined on the domain that is currently logged in. Rules are applied in the following ways:

- If a rule exists for a domain group account and the domain user who is authenticated is a member of that domain group, the rule applies to that user.
- If rules are defined for multiple domain groups and the domain user who is authenticated is a member of those domain groups, then all rules apply.
- If the domain user is a member of a domain group or local group for which a rule exists, or if the domain user is a member of a domain group that is a member of a local group **and** a rule for the local group exists, the rule applies.

Data Scoping rule evaluation takes place as described for a local computer.



## Data Scoping Rule Evaluation in Management Command Center

The user ID and password used for Data Scoping rule evaluation when using the Management Command Center varies.

- When you are running the Management Command Center locally, the user currently logged in is used for Data Scoping rule evaluation.
- When you are running the Management Command Center remotely over RMI, a signon dialog always appears. The user ID and password you enter on the signon dialog is used for Data Scoping rule evaluation.
- Domains and local groups for Data Scoping are supported when using the Management Command Center on Windows.

Enter the domain account in the form "Domain1\User1" in the Management Command Center SignOn dialog, regardless of where the Management Command Center is running.

The domain of the computer where the MDB resides must be the same domain or trust the domain of the domain account used for authentication. For example, if you enter Domain1\User1 in the Management Command Center signon dialog, the computer that contains MDB must be logged into Domain1 or DomainX, where DomainX trusts Domain1. If not, the authentication fails when Data Scoping is active and Data Scoping rules are not evaluated. If you enter the account in the form "User1," then it is considered a local user account defined to the computer where Management Command Center is running.

The domain user must be a member of the TNDUsers group, either directly or indirectly through a domain group. Data Scoping rule evaluation occurs as described for local computers in Data Scoping Rule Evaluation Using Windows Domain Groups. A local user is authenticated as described for local computers in Data Scoping Rule Evaluation Using Windows Local Groups.

**Note:** You should use the TNDUsers group for any Windows user ID that is authenticated using Management Command Center remotely.

Management Command Center running on a UNIX/Linux client can use Windows domain accounts to authenticate when Data Scoping is active because all Data Scoping rule evaluation occurs on the Windows computer that contains the MDB.

When Data Scoping is deactivated, you can use any domain account that is authenticated.

## Data Scoping Rule Evaluation Using Ingres Databases

The Ingres user ID and password that you use to sign on to the MDB is used for Data Scoping rule evaluation.

For example, when you are using the WorldView Classic GUI (Windows), the user ID and password you provide on the Repository Sign On dialog is saved and stored in the VNODE. When you start any additional WorldView component, such as Object Browser or Severity Browser, you are not prompted again for MDB credentials because the credentials saved in the VNODE are used. These saved credentials are used for evaluating any other Data Scoping rules that may apply when using any WorldView component.

When you are running the Management Command Center locally, the user currently logged in is used for Data Scoping rule evaluation. When you are running the Management Command Center remotely over RMI, a signon dialog always appears. The user ID and password you enter on the signon dialog is used for Data Scoping rule evaluation.

## Data Scoping Limitations When the MDB Resides on UNIX/Linux

Certain limitations exist for local and domain groups when the MDB resides on a UNIX or Linux operating system. You can create Data Scoping rules for domain groups from a Windows client connecting to the MDB that resides on UNIX or Linux, and these rules can be enforced properly.

However, these domain group rules cannot be created or enforced from UNIX or Linux clients because these domain groups exist on Windows domain controllers.

You cannot create rules for UNIX or Linux local groups from a Windows client. You can create UNIX or Linux local group rules from a UNIX or Linux server, but these rules cannot be enforced from Windows clients, only from UNIX or Linux clients.

## Data Scoping Limitations on UNIX/Linux When the MDB Resides on Windows

Domain groups and local groups are not supported when CA NSM is installed on UNIX or Linux and the MDB resides on a Windows computer. Only rules for specific Windows users are supported in this case.

## Data Scoping in the 2D Map (Windows)

Data Scoping works differently in the 2D Map than in other CA NSM applications because of the unique design of the 2D Map. Data Scoping rules for all other CA NSM applications are enforced immediately after they are created. Data Scoping in the 2D Map works in the following ways:

### Delta Mode

Objects in the 2D Map are loaded when the folder containing the objects is opened. If Data Scoping rules for read access are created before opening a folder containing objects where Data Scoping rules exist that affect these objects, the Data Scoping rules are enforced. The objects are visible or hidden on the Map depending on the rules.

If rules are created after opening the folder, the Data Scoping rules are not enforced because all the objects for the folder have been read into the map and are not reread until the 2D Map is restarted.

Rules controlling create, update, and delete access are always enforced immediately after the rule is created. You do not have to restart the 2D Map after creating these types of rules.

### Non-Delta Mode

All objects in the 2D Map are loaded when the 2D Map is started. Therefore, any Data Scoping rules you create that affect read access are not enforced until the 2D Map is restarted.

Rules controlling create, update, and delete access are always enforced immediately after the rule is created. You do not have to restart the 2D Map after creating these types of rules.

## Activate Data Scoping on Windows

You must activate Data Scoping using the Data Scoping utility before you can create Data Scoping rules. The Data Scoping utility creates the DataScope classes.

To activate Data Scoping and create the DataScope classes, run the following command:

```
wscpini.exe /rrepositoryname /uuserid /ppassword
```

Data Scoping is activated and classes are created.

**Note:** Only administrators can run this command. If you run this command after you have connected to the MDB, applications already connected will not have Data Scoping rules enforced until the applications are restarted. You also will not be able to create new Data Scoping rules until the applications are restarted.

## Deactivate Data Scoping on Windows

There are times that you may want to deactivate Data Scoping.

To deactivate Data Scoping, run the following command:

```
wscpdel.exe /rrepositoryname /userid /ppassword
```

Data Scoping is deactivated.

**Note:** Only administrators can run this command.

## Activate Data Scoping on UNIX/Linux

You must activate Data Scoping using the Data Scoping utility before you can create new Data Scoping rules. The Data Scoping utility creates the DataScope classes.

To activate Data Scoping and create the DataScope classes, run the following command:

```
$CAIGLBL0000/ww/bin/wscpini.exe -r repository-name
```

where *repository-name* is the name of the MDB on which you want to activate Data Scoping.

Remote MDBs are not supported; therefore, the database name should be that of the local MDB.

**Note:** Only the root user can run this command. If you run this command after you have connected to the MDB, applications already connected will not have Data Scoping rules enforced until the applications are restarted. You also will not be able to create new Data Scoping rules until the applications are restarted.

## Deactivate Data Scoping on UNIX or Linux

There are times that you may want to deactivate Data Scoping.

### To deactivate Data Scoping

1. Run the following command:

```
$CAIGLBL0000/wv/bin/wwscpdel.exe -r repository-name
```

where *repository-name* is the name of the MDB on which you want to deactivate Data Scoping.

Data Scoping is deactivated.

2. Recycle WorldView.

Changes made to the database tables are recorded.

**Note:** Only the root user can run the `wwscpdel` command.

## DataScope Rule Editor

You use the DataScope Rule Editor in the Management Command Center or the WorldView Classic GUI to create Data Scoping rules to protect the MDB object data from unauthorized access, or to filter large amounts of data into smaller subsets of data for a group or user. The administrator must be authorized to create or update rules for the specific class.

**Note:** For specific procedures, see the DataScope Rule Editor online help.

## Implement End-to-End Data Scoping

Data Scoping is designed to be a complete solution. Since the MDB can be accessed with tools that are not part of CA NSM, the possibility of a security breach exists that Data Scoping cannot defend against on its own. However, you can take steps to ensure security no matter how the MDB is accessed. The following steps describe how you, as a MDB administrator, can implement end-to-end Data Scoping in CA NSM:

1. To access MDB, define a database user ID to the database server with the following attributes:
  - No SELECT, UPDATE, INSERT, or DELETE privileges for any table in the MDB.
  - No privileges to run any stored procedure defined for the MDB.
  - No membership in any Windows group or domain that has complete access to the MDB, or no membership in any UNIX/Linux group or any role that has complete access to the MDB.
  - On UNIX/Linux, the database user ID must consist of all lowercase characters.

The user ID is then not able to write any dynamic Ingres or Microsoft SQL Server applications that access the MDB database and cannot access MDB through external tools such as Ingres SQL Microsoft SQL Query Analyzer. The only way to write an application to access the MDB is by using the WorldView Application Programming Interface (API). Since all CA NSM applications access the MDB through the WorldView API, they are assured complete access to the MDB. If Data Scoping is deactivated, this user ID should have full database privileges restored.

2. Create Data Scoping rules to restrict access to users who write applications using the WorldView API, or who use CA NSM tools.

When the rules are saved in the MDB for which access needs to be restricted, Data Scoping enforcement is complete for all CA NSM applications. Applications that are not using CA NSM are denied complete access to the MDB because of the precautions you set up when you created the database user ID.

3. Maintain Data Scoping security.

Data Scoping security is an ongoing process. You can update and delete rules. You can remove DataScope classes to completely deactivate Data Scoping.

On Windows platforms, when you update Data Scoping rules, they are enforced immediately, except for the conditions noted in Data Scoping in the 2D Map (Windows).

On UNIX/Linux platforms, when you update Data Scoping rules, they are enforced immediately.

## Communication Protocol Security

At each successive release of CA NSM, an attempt is made to optimize the communication methodologies, thus reducing the security and administrative concerns. The following primary communication protocols are used with CA NSM:

- Standard SNMP
- CA Common Communication Interface (CAICCI)
- Distributed Intelligence Architecture (DIA)

Each of these methods has different security considerations, such as which ports are required to be open, whether or how they encrypt data, and so forth.

The following sections clarify the usage and considerations for each of these methodologies. For additional information regarding port configuration, see "Utilizing and Configuring Ports."

### Encryption Levels

As data traverses the network, it is important to understand the encryption methodologies in place for security compliance so that you can be assured that data is protected appropriately.

Component	Encryption Level	Comments
CAICCI	SSL 80-bit (See Note 1)	Open SSL
Unicenter Notification Services	SSL (See Note 1)	Open SSL
ORB	48-bit	DES (See Note 2)
SNMPv3	48-bit	DES (See Note 2)
CAM	SSL	SSA 2.0 (See Note 3)
DIA	SSL	Open SSL

**Note 1:** CAICCI comes preconfigured to let you use the strongest encryption possible by downloading the algorithms from the external OpenSSL library. OpenSSL uses two forms of encryption: an asymmetric algorithm to establish a connection, and a symmetric algorithm for the duration of a connection after it has been established. The strongest asymmetric algorithm we recommend using is RSA with a 2048-bit key. The strongest symmetric algorithm we recommend using is AES with a 256-bit key.

**Note 2:** DES encryption is built into the code or product module.

**Note 3:** For more information about configuring CAM to use SSL encryption, see the CAM section in the chapter "Using Ports to Transfer Data."

## Agent to Manager Communication Security

Distributed Intelligence Architecture (DIA) allows for high speed, secure communications to transport data while providing remote node management and inherent failover capabilities. All out-bound communications from all DIA components use secure sockets. The SSL protocol provides connection security that has three basic properties:

- The connection is private. Encryption is used after an initial handshake to define a secret key. Symmetric cryptography is used for subsequent data encryption (DES, RC4, AES).
- The peer's identity can be authenticated using asymmetric or public key cryptography. Authentication is certificate based (RSA, DSS, ADH).
- The connection is reliable. Message transport includes a message integrity check using a keyed MAC. Secure hash functions are used for MAC computations (SHA, MD5).

The cipher suite, which declares the algorithms used for each of these areas, is fully configurable to use any of the combinations available through OpenSSL. In general, we use the strongest ciphers that also provide acceptable performance. The default cipher suites, as delivered, are as follows:

Protocol:	SSLv3 or TLSv1
Key exchange:	RSA
Authorization:	RSA using a 1024-bit key
Encryption:	AES with a 256-bit key
MAC algorithm:	SHA1



If configured to run anonymously (peers are not authenticated), the defaults are as follows:

Protocol:	SSLv3 or TLSv1
Key exchange:	ADH
Authorization:	NONE
Encryption:	AES with a 256-bit key
MAC algorithm:	SHA1

## Common Communications Interface (CAICCI)

CA Common Communications Interface (CAICCI) is used by some components for cross-platform communication. CAICCI is maintained and configured transparently from various user interfaces provided by the products that rely on it. This robust communication facility transfers data reliably across a variety of networks and protocols. See the port configuration considerations in Utilizing and Configuring Ports.

**Note:** Do not shut down CAICCI if other Unicenter NSM services are running. Also, always use the `unishutdown` command to shut down CAICCI.

## CAICCI Secure Sockets Facility (CCISSF)

For CAICCI users who have the highest-level security requirements or must be compliant with certain governmental security certifications, CCISSF provides a solution. CCISSF transmits any data from components or products using CAICCI (such as CA NSM) over a Secure Sockets Layer (SSL) connection, which is a highly secure encryption mechanism.

As installed, CAICCI transmits data from point to point using traditional TCP/IP services, which means the application defines the format and encryption level of the data exchanged. Therefore, products that use CAICCI secure data areas that are sensitive to a product. CCISSF effectively “wraps” all data in an encrypted envelope. The resultant transmission is significantly more resistant to attacks, reverse engineering, and accidental or malicious intrusion.

## OpenSSL

CCISSF uses OpenSSL, which is a consortium open source facility. For more information, see <http://www.openssl.org>. Use of OpenSSL provides standards-based encryption and authentication for the sender and receiver.

In OpenSSL, authentication is achieved through certificates. *Certificates* contain data about the local host that the remote host can use to determine whether the local host is authentic. This mechanism ensures that the communicating partner is who the partner claims to be.

Secure Sockets Layer functionality is provided by dynamically loading available OpenSSL libraries. These libraries must be available on all installed machines. The minimum version requirement of OpenSSL to be used with CCISSF is OpenSSL Version 0.9.7. It is your responsibility to obtain a version of OpenSSL that is consistent with your needs and planned deployment. For your convenience, a version of OpenSSL will be installed with CAICCI.

## Enable CCISSF

CCISSF is disabled by default to provide “out of the box” compatibility with previous versions of CAICCI. Also, not all users will require this enhanced level of security, which comes with some performance costs.

To enable CCISSF, do one (or both) of the following:

- Set the following environment variable in the system environment:

```
CAI_CCI_SECURE=[YES|NO]
```

### **YES**

Specifies that all connections, unless otherwise specified in the remote daemon configuration file, will have CCISSF enabled.

### **NO**

Specifies that the remote daemon will not request SSL for any connections unless they are overridden in the configuration file. However, all incoming secure connections will be accepted using a secure connection. The default is NO.

**Note:** Regardless of the environment variable setting, communications to a remote CAICCI will not use CCISSF unless an entry for that remote node is present in the remote daemon configuration file on the local system.

- Override the environment variable setting in the `ccirmtd.rc` file by setting the following parameter to override the default behavior:

```
SECURE=[YES|NO]
```

#### **YES**

Specifies that a connection attempt to the corresponding remote CAICCI will be required to be made in a secure manner.

#### **NO**

Specifies that an outgoing connection attempt will be made in a non-secure manner unless the corresponding remote CAICCI requires it.

### How to Determine the Effective Security Value

The effective security value is determined by applying rules in this order:

1. The connection security as determined by a product configuration if it exists.
2. The value of the host-specific `SECURE` value in a configuration if it exists.
3. The value of the `CAI_CCI_SECURE` environment variable.

CCISSF will always connect with the highest possible level of security when communicating with another CCISSF-capable CAICCI. The following table describes the behavior:

<b>Source CAICCI Effective SECURE Value</b>	<b>Target CAICCI Effective SECURE Value</b>	<b>Connection Status</b>
Yes	Yes	Secure
Yes	No	Secure
No	Yes	Secure
No	No	Not secure

### Compatibility with Previous Versions of CAICCI

When two remote daemons (one supports CCISSF and the other does not) connect to each other, the connection will never use the Secure Sockets Layer; therefore, the connection will be non-secure or denied completely. The following table defines the behavior:

<b>Effective SECURE Value of the CCISSF-Capable Version</b>	<b>Connection Status</b>
Yes	Denied

Effective SECURE Value of the CCISSF-Capable Version	Connection Status
No	Accepted, but not secure

### Configuring CCISSF

CCISSF depends on OpenSSL for effective communication. To use CCISSF, you must configure it to use OpenSSL.

If OpenSSL is available on the system when CAICCI initializes, CAICCI uses the OpenSSL libraries to provide service for any secure connections. If OpenSSL is not available, CAICCI follows the behavior defined in the following table:

Effective SECURE Values of All Connections	CAICCI Behavior if OpenSSL Is Not Available
No—Default is No and no remote configuration file entries with SECURE=YES	<p>Warning message to the Event Log or syslog indicating that OpenSSL is not present at the time of initialization.</p> <p>All inbound connections will be denied if a secure connection request is made.</p> <p>All outbound connections will be made as a non-secure request.</p>
Yes—Default is Yes or at least one remote configuration file with SECURE=YES	<p>An error message will be issued to the Event Log or syslog indicating the required OpenSSL component is not present and that only non-secure connections will be made.</p> <p>CAICCI will initialize but only connections that are requested to be non-secure will be made. Any connection for which the effective value of Secure is Yes will be disabled.</p>

**Note:** SSL connections are currently supported **only** between CAICCI remote daemons. Communication between hosts that use the QUES layer (transport daemon) cannot use SSL. The QUES implementation is typically used in Windows environments. For those users that want to use CCISSF, you must migrate to the remote daemon implementation.

## Default Certificate

To facilitate installation, CCISSF supplies a default certificate to allow “out of the box” operation. However, use of the default certificate cannot ensure any significant level of authentication since all certificates are identical. For true authentication, we strongly recommend you use customized PEM format certificates in accordance with site standards, and replace the default certificate in the location discussed in this topic.

The default certificate has the following properties:

- Common Name is “Default CCI Certificate.”
- Issuer Common Name is “Default Root Certificate.”
- Serial number is 0xC.
- The certificate becomes invalid after January 2, 2033.

The default certificate’s private key is encrypted with the passphrase “CACCI”. When CCISSF is installed a default certificate called cert.pem is installed on the system. Unless the default configuration is altered (see ccisslcfg Utility) CCISSF will use this default certificate. This default certificate can be replaced with a user-provided certificate of the same name, or the ccisslcfg utility can be used to configure CCISSF to use a user-provided certificate with a different name.

## Location of Default Certificate

This certificate must be placed in \$CAILOCL0000.

The default root certificate chain must be placed in one of the following locations:

- OpenSSL’s default “certs” directory (specified during compilation of OpenSSL)
- Appended or renamed to the OpenSSL default cert.pem file (specified during compilation of OpenSSL)
- Placed in the same directory as the default CAICCI certificate (as previously described) and be named root.pem

## Certificate Revocation Lists (CRLs)

Certificate revocation lists (CRLs) are a common method to track probable rogue certificates. These are certificates that can no longer be trusted because the private key has become “public knowledge,” when using a public key infrastructure.

CCISSF allows for the use of CRLs. To use this feature, place a generated CRL in the default CRL location of `$CAILOCL0000\crl.pem`.

Using CRLs is not useful when using the CCISSF default certificate, since this certificate is the same across all nodes. You can use this feature only when generating your own certificates.

**Note:** The `ccisslcfg` utility can override these default file locations.

## Ccisslcfg Utility--Specify Certificate Location

CCISSF looks in a predetermined location for items such as a host’s certificate or root certificate (see [Default Certificate](#) (see page 85)). However, users may want to keep them in other locations and tell CCISSF of the change. The `ccisslcfg` utility lets you do this. When executed, `ccisslcfg` prompts for the following:

- Host’s certificate and private key file.
- How to get the corresponding passphrase of the private key file, which can be one of the following options:
  - By typing in the passphrase and `ccisslcfg` will store it in an encrypted form for you.
  - By specifying the absolute path of a file that contains the passphrase in unencrypted form.
  - By specifying that you will provide the passphrase in the `password_cb()` callback in the `cauccissl/libccissl` library.
  - By stating that the private key is not protected with a passphrase.

- Whether to use OpenSSL's default root certificate authority locations. CCISSF will use these locations in addition to any locations you specify in the following items below:
  - Any number of root certificates
  - Any number of directories containing root certificates (when specifying a directory, it is assumed that the files inside are all named with the convention of using the hash value of the issuer's subject name. OpenSSL will not be able to correctly look up these files if they are not named with this convention. Read the OpenSSL documentation for more information.)
- The location of any certificate revocation lists (CRLs), which can be any number of files or directories (As stated before, when specifying a directory, we assume the files inside are all named with the convention of using the hash value of the issuer's subject name.)

After `ccisslcfg` prompts you for all these settings, it will write them in encrypted form to the file `%CAILOCL0000%\ccissl.cfg`.

`Ccisslcfg` will overwrite any previous settings you may have set in the past. Because this configuration file is encrypted, only the `ccisslcfg` utility can change these settings. Note that although the contents of this file are encrypted, we recommend that the permissions are set so that only administrators and CCISSF have access to this file.

The presence of this file overrides CCISSF's default behavior with respect to where it looks for certificates. This configuration utility does not need to be used if you plan to use the default CCISSF certificate locations along with providing the `password_cb()` callback in the `cauccissl\libccissl` library.

### Ccicrypt Utility--Set Encryption Properties

CCISSF also includes `ccicrypt`, a general purpose encryption utility. This utility has the following format:

```
Ccicrypt [-help] [-in infile] [-out outfile] [-p password] [-cipher cipher] [-dv dataVector] -encrypt|-decrypt
```

#### **-in**

Specifies a file to read input from. If it is left out, `ccicrypt` uses standard input.

#### **-out**

Specifies a file to direct output to. If it is left out, `ccicrypt` uses standard output.

#### **-p**

Specifies a password for `ccicrypt` to use. If it is left out, `ccicrypt` uses a default internal password.

**-cipher**

Tells ccicrypt what type of encryption algorithm to use. Enter ccicrypt -help for a list of choices and refer to <http://www.openssl.org> for descriptions. If an algorithm is not specified, DES in cipher feedback mode (des-cfb) is used.

**-dv**

Specifies a data vector to ccicrypt. In addition to a password, some encryption algorithm modes (like cipher feedback mode) require additional random data to further randomize the output. The data vector is any string of characters.

**-encrypt or -decrypt**

Specify whether ccicrypt should encrypt or decrypt data. There is no default for this. One option should always be specified.

**-help**

Lists the available types of encryption algorithms.

## Using a Passphrase

A passphrase is used to protect elements of the certificate while it resides on the local file system. CAICCI requires access to your system passphrase to properly function. By default, CCISF will provide SSL with the passphrase used to encrypt the default certificate.

To use a different passphrase, several options exist. First, the password\_cb function can be provided in a shared library (see [Programming a Customized SSL Environment](#) (see page 88)). Additionally, you can use the ccisslcfg utility to provide the passphrase itself or the absolute path of a file that contains the desired passphrase.

## How You Program a Customized SSL Environment

We encourage you to create a customized SSL environment. Customization involves creating new certificates using OpenSSL tools. CCISF has no special restrictions on the data within certificates. Any validly formatted SSL certificate will work. CAICCI lets you retrieve data in the following fields from a certificate:

- Serial Number
- Common Name
- Locality
- State or Province
- Country
- Organization
- Organizational Unit



- Email Address
- Not Valid Before Date
- Not Valid After Date
- Issuer Common Name
- Issuer Locality
- Issuer State or Province
- Issuer Country
- Issuer Organization
- Issuer Organizational Name
- Issuer Email Address

Additionally, the following popular certificate extensions can be retrieved:

- Basic Constraints
- Key Usage
- Subject Alternate Name
- Issuer Alternate Name

Additional fields can be defined in the provided CAICCI header file but are not supported by CAICCI at this time.

## Default Functions

Along with the default certificates, CCISSF also provides two default functions, `password_cb` and `verifyCert`, to provide the private key's password and authenticate the remote host respectively. To facilitate the customized environment, we provide an API. This interface acts as a convenient way to access the underlying Open SSL environment.

### password\_cb Function--Provide Default Passphrase

The password\_cb function is called whenever the user certificate's private key must be used. The default functionality is to provide the default passphrase used to encrypt the default private key.

This function has the following format:

```
int password_cb(char *buf,  
               int num,  
               int rwflag,  
               void *userdata);
```

#### **buf**

Specifies the SSL-provided buffer of length *num* that points to the null-terminated password string upon exit of the function.

#### **num**

Specifies the length of buffer pointed to by *buf* (includes space for terminating character).

#### **rwflag**

Specifies the flag that indicates whether the function was called for encryption (*rwflag* is nonzero) or decryption (*rwflag* = 0).

#### **userdata**

Reserved for future use (will always be NULL).

This function returns the length of the password string pointed to by *buf*.

### How You Provide Custom Functionality

To provide custom functionality, supply a DLL. This library should export at least one of the symbols password\_cb and verifyCert, and be named cauccissl.dll.

If this library is present, CAICCI checks for the exported symbols password\_cb and verifyCert. For each one that is exported in the library, CAICCI loads and calls that function. If either symbol is not exported, CAICCI uses the corresponding CCISSE default function (see Default Functions).

### User-Exported Symbols

#### **password\_cb**

Function to supply private key passphrase.

#### **verifyCert**

Function to check the authenticity of a remote certificate.

## User Available Functions

You have access to the following functions when writing customized authentication. Function pointers to the functions are supplied inside the data structure supplied to `verifyCert`.

### `get_cert_info` Function--Retrieve Certificate Information

When writing a customized authentication function, you can use `get_cert_info` to retrieve information from a certificate on the local computer or the foreign computer. The programmer must supply only the ID of the data requested.

This function has the following format:

```
int          get_cert_info(psCACE,
                          CCISSL_CERT_ID,
                          CCISSL_CERT_DATA_ID,
                          char** cert_data,
                          int* cert_data_len);
```

#### **psCACE**

Specifies a pointer to user structure (see `CCISSL_CLIENT_Auth_Callback_Env` following).

#### **CCISSL\_CERT\_ID**

Specifies an enumerated type indicating which certificate (local or remote) to look at (see header file at end).

#### **CCISSL\_CERT\_DATA\_ID**

Specifies an enumerated type indicating which field of the certificate to return (see header file at end).

#### **cert\_data**

Specifies a pointer to a `char**`. This will hold the data requested upon successful return. The user need not `malloc` or `free` this space. CAICCI will handle all memory management.

#### **cert\_data\_len**

Specifies a pointer to an `int*`. This will contain the length of `cert_data` upon successful return.

This function returns -1 on error, or if the data requested does not exist. On success it returns the index into the array of data where the requested information exists.

### enum\_cert\_info Function--Return Certificate Data by Array

Upon first call to this function, certificate data held in the first element of the local\_cert or remote\_cert array is returned. Each successive call returns certificate data held in the next element of the array. When all data has been returned, the next call will return data in the first element of the array and the process repeats.

This function has the following format:

```
int    enum_cert_info(psCACE,
                    CCISSL_CERT_ID,
                    CCISSL_CERT_DATA_ID*,
                    char** cert_data,
                    int* cert_data_len);
```

#### **psCACE**

Specifies a pointer to user structure (see CCISSL\_CLIENT\_Auth\_Callback\_Env following).

#### **CCISSL\_CERT\_ID**

Specifies an enumerated type indicating which certificate (local or remote) to look at (see header file at end).

#### **CCISSL\_CERT\_DATA\_ID\***

Specifies a pointer to CCISSL\_CERT\_DATA\_ID that contains the data ID of the data pointed to by cert\_data upon return of the function.

#### **cert\_data**

Specifies a pointer to a char\*\*. This will hold the next piece of data in the array. The user need not malloc or free this space. CAICCI will handle all memory management.

#### **cert\_data\_len**

Specifies a pointer to an int\*. This will contain the length of cert\_data upon successful return.

This function returns -1 on error; 0 if data returned is the last in the array. If not, the index of the next piece of data is returned.

## output\_cert\_info Function--Return Certificate Information to a File

The `output_cert_info` function outputs a string to the destination indicated by `CCISSL_OUTPUT_CERT_INFO_TYPE`. All output will have the string, "CCI\_1024:" pre-pended to it.

This function has the following format:

```
void    output_cert_info(psCACE,
                        CCISSL_OUTPUT_CERT_INFO_TYPE,
                        char* format, ...);
```

### **psCACE**

Specifies a pointer to user structure (see `CCISSL_CLIENT_Auth_Callback_Env`).

### **CCISSL\_OUTPUT\_CERT\_INFO\_TYPE**

Specifies an enumerated type indicating the destination of the output (see header file described in next section). `CCISSL_OUTPUT_TO_LOG` specifies the event log.

### **format**

Specifies a string to be output.

...

Specifies variables to be substituted into format.

## CCISSL\_CLIENT\_Auth\_Callback\_Env

The CAICCI structure (`CCISSL_CLIENT_Auth_Callback_Env`) is specified at the end of this document. This structure stores values that are taken from the certificates. The fields of the structure are as follows:

### **client\_callback\_handle**

Specifies a value reserved for future use and always set to zero.

### **local\_hostname**

Specifies a pointer to string representing local host name.

### **local\_ipaddr**

Specifies a character array representation of local host's IP address associated with remote daemon.

### **local\_portno**

Specifies the port number of the local side of the SSL connection.

### **local\_CCI\_sysid**

Specifies a pointer to string representing the system ID CAICCI has assigned to the local host.

**remote\_hostname**

Specifies a pointer to string representing the remote host name.

**remote\_ipaddr**

Specifies a character array representation of remote host's IP address associated with its remote daemon.

**remote\_portno**

Specifies the port number of the remote side of the SSL connection.

**remote\_CCI\_sysid**

Specifies a pointer to string representing the system ID CAICCI has assigned to the remote host.

**remote\_appname**

Specifies a pointer reserved for future use and set to NULL.

**remote\_taskid**

Specifies a pointer reserved for future use and set to NULL.

**remote\_userid**

Specifies a pointer reserved for future use and set to NULL.

**local\_cert**

Specifies a pointer to an array holding information from the local machine's certificate.

**local\_cert\_elem\_ct**

Specifies the count of how many elements are in the array pointed to by local\_cert.

**remote\_cert**

Specifies a pointer to an array holding information from the remote machine's certificate.

**remote\_cert\_elem\_ct**

Specifies the count of how many elements are in the array pointed to by remote\_cert.

**local\_cert\_elem\_loc**

Specifies the current index into array pointed to by local\_cert.

**remote\_cert\_elem\_loc**

Specifies the current index into array pointed to by remote\_cert.

**get\_cert\_info**

Specifies the pointer to the function that returns data from a certificate based on ID (see Header Information for User Customization).

**enum\_cert\_info**

Specifies the pointer to the function that sequentially returns data from a certificate (see Header Information for User Customization).

**output\_cert\_info**

Specifies the pointer to the function that allows the user to output a data string (see Header Information for User Customization).

**Header Information for User Customization**

Use the following structures if you want to write customized authentication. Copy these structures and user-supplied code into a header file.

```
typedef enum CCISSL_CERT_ID_T {
    CCISSL_REMOTE_CERT_INFO,
    CCISSL_LOCAL_CERT_INFO
} CCISSL_CERT_ID;
```

```
typedef enum CCISSL_CERT_DATA_ID_T {
    CCISSL_CERT_BODY_DER,
    CCISSL_CERT_BODY_BASE64,
    CCISSL_CERT_SERIAL_NUMBER,
    CCISSL_CERT_COMMON_NAME,
    CCISSL_CERT_LOCALITY,
    CCISSL_CERT_STATE_OR_PROVINCE,
    CCISSL_CERT_COUNTRY,
    CCISSL_CERT_ORG,
    CCISSL_CERT_ORG_UNIT,
    CCISSL_CERT_DN_PRINTABLE,
    CCISSL_CERT_DN_DER,
    CCISSL_CERT_POSTAL_CODE,
    CCISSL_CERT_EMAIL,
    CCISSL_CERT_NOT_BEFORE,
    CCISSL_CERT_NOT_AFTER,
    CCISSL_CERT_EXTENSION_BASIC_CONSTRAINTS,
    CCISSL_CERT_EXTENSION_KEY_USAGE,
    CCISSL_CERT_EXTENSION_SUBJECT_ALT_NAME,
    CCISSL_CERT_EXTENSION_ISSUER_ALT_NAME,
```

```
    CCISSL_CERT_ISSUER_COMMON_NAME,  
    CCISSL_CERT_ISSUER_LOCALITY,  
    CCISSL_CERT_ISSUER_STATE_OR_PROVINCE,  
    CCISSL_CERT_ISSUER_COUNTRY,  
    CCISSL_CERT_ISSUER_ORG,  
    CCISSL_CERT_ISSUER_ORG_UNIT,  
    CCISSL_CERT_ISSUER_DN_PRINTABLE,  
    CCISSL_CERT_ISSUER_DN_DER,  
    CCISSL_CERT_ISSUER_POSTAL_CODE,  
    CCISSL_CERT_ISSUER_EMAIL  
} CCISSL_CERT_DATA_ID;  
  
typedef enum CCISSL_OUTPUT_CERT_INFO_TYPE_T {  
    CCISSL_OUTPUT_TO_STDOUT,  
    CCISSL_OUTPUT_TO_LOG  
} CCISSL_OUTPUT_CERT_INFO_TYPE;  
  
typedef struct Certificate_Element {  
    char *data;  
    int length;  
    CCISSL_CERT_DATA_ID id;  
} certElem;
```



```
typedef struct CCISSL_Client_Auth_Callback_Env sCACE;
typedef sCACE* psCACE;

struct CCISSL_Client_Auth_Callback_Env {
    int            client_callback_handle;
    char           *local_hostname;
    int            local_ipaddr;
    int            local_portno;
    char           *local_CCI_sysid;

    char           *remote_hostname;
    int            remote_ipaddr;
    int            remote_portno;
    char           *remote_CCI_sysid;
    char           *remote_appname;
    char           *remote_taskid;
    char           *remote_userid;

    certElem*local_cert;
    int            local_cert_elem_ct;

    certElem*remote_cert;
    int            remote_cert_elem_ct;

    int            local_cert_elem_loc;
    int            remote_cert_elem_loc;

    int            (*get_cert_info)(psCACE,
                                   CCISSL_CERT_ID,
                                   CCISSL_CERT_DATA_ID,
                                   char** cert_data,
                                   int* cert_data_len);

    int            (*enum_cert_info)(psCACE,
                                   CCISSL_CERT_ID,
                                   CCISSL_CERT_DATA_ID*,
                                   char** cert_data,
                                   int* cert_data_len);

    void           (*output_cert_info)(psCACE,
                                   CCISSL_OUTPUT_CERT_INFO_TYPE,
                                   char* format, ...);
};
```



# Chapter 3: Discovering Your Enterprise

---

This section contains the following topics:

[Discovery](#) (see page 99)

[How You Can Combine Running Classic and Continuous Discovery](#) (see page 101)

[Classic Discovery Multi-Homed Device Support](#) (see page 102)

[Discovery Classification Engine](#) (see page 102)

[Discovery Timestamp](#) (see page 102)

[How Subnet Filters Work](#) (see page 103)

[How Timeout Values Affect Discovery](#) (see page 103)

[Discovery Object Creation Rules](#) (see page 104)

[Device Not Discovered](#) (see page 114)

[Discovering Your Network Devices Continuously in Real-Time Mode](#) (see page 115)

[Discovering Your Network Devices on Demand Using Classic Discovery](#) (see page 129)

[Discovering IPv6 Network Devices using Common Discovery](#) (see page 138)

## Discovery

Discovery is the process by which devices on the network are found and classified and then placed in the MDB as managed objects. Discovery discovers and classifies devices on Internet Protocol (IP). Classic Discovery also supports two non-IP plugins that are fully integrated:

- Storage Area Network (SAN)--SAN Discovery can be started using the command line or the SAN pages in the Discovery Classic GUI or the Unicenter MCC. It discovers devices that are on a SAN and IP-enabled, as well as access points in the SAN. It automatically creates SAN Business Process Views in the 2D Map.
- Internetwork Packet Exchange Protocol (IPX) Discovery--IPX Discovery can be started using its own command line (ipxbe) or by using the Discovery Classic GUI. It discovers devices located on an IPX network.

Discovery also determines if a device provides Web-Based Enterprise Management (WBEM) data, and if so, creates a WBEM object in the device's Unispace. The Agent Technology WorldView Gateway service locates agents running on the network objects.

**Note:** An MDB must exist before you can run Discovery to discover your network devices and populate the MDB.

Once defined, you can view, monitor, and manage these objects and their Management Information Base (MIB) through the 2D Map, ObjectView, and the Topology Browser. You can manage the entities they represent using Event Management, Agent Technology, and third-party manager applications.

When you install your product, you can decide which type of Discovery you want to use:

- Classic Discovery is an on demand process that lets you decide which subnets you want to discover and when. You can also configure Classic Discovery to run in regular intervals, which can be used as an alternative to Continuous Discovery and ensures that your discovered environment in the MDB is always current. You can start a Classic Discovery from the Discovery Classic GUI, the Unicenter MCC, the Unicenter Browser Interface, or the command line.
- Continuous Discovery is event-driven and ongoing. It employs a manager and agents that continuously scan your network in real-time mode for new devices or changes in IP addressing of existing IP devices. You can configure Continuous Discovery for optimal load balancing between the Discovery agents and the Discovery Manager. If you choose this method of discovery, you must install the Discovery agents and the Discovery Manager.
- Common Discovery is a new tool used for discovery across multiple CA products. In CA NSM r11.2, you can use Common Discovery to discover devices on IPv6 networks. WorldView uses the Common Discovery Import service to poll Common Discovery and populate the MDB with IPv6 network entities. If you choose this method of discovery, you must install the WorldView Manager, MCC, and WorldView Provider. You also must install the Common Discovery component on all servers on which you want to discover IPv6 entities.

**Note:** We do not recommend that you run Continuous Discovery and Classic Discovery concurrently on your network. Timing issues could result in the duplication of objects and an unnecessarily heavy load on the network and the MDB. You can however, run a combination of Classic and Continuous Discovery. For more information, see [How You Can Combine Running Classic and Continuous Discovery](#) (see page 101).

## How You Can Combine Running Classic and Continuous Discovery

We recommend that you run a combination of Classic and Continuous Discovery when you want to discover subnets. However, Classic and Continuous Discovery work differently depending on what options you select for both methods. You need to be aware of these differences to avoid creating duplicate devices in the MDB.

- Classic Discovery and Continuous Discovery name devices differently:
  - Classic Discovery supports naming a device using its sysname (the MIB-II value for a device that supports SNMP), which is the default if no DNS name is available.
  - Sysnames are not supported by Continuous Discovery, except for routers that do not have valid DNS names for their IP interface cards.

To avoid discovering duplicate devices, in Classic Discovery, set the `dscrbe -j` option to IP to use the IP address if the DNS name cannot be found. Using IP addresses to name discovered devices ensures that objects are named using the same method and that no duplicates result. Set this option only if DNS is not enabled in your environment.

**Note:** If you are using the Discovery Classic GUI to run Discovery, select "Use IP address instead of Sysname."

- When you run a full subnet discovery using Classic Discovery, stop the Continuous Discovery services.
- Continuous Discovery discovers only subnets on which an agent is deployed. To discover non-agent subnets because you want to automatically monitor them using Continuous Discovery, you can run the Classic Discovery `dscrbe` command to discover a router and all of the subnets it supports, or you can write a script using the `dscrbe -7` option to discover all of the gateways on the desired subnets.

## Classic Discovery Multi-Homed Device Support

When a device has more than one physical interface to one or more networks, it is known as a *multi-homed device*. You can discover and manage multi-homed devices.

When a multi-homed device is discovered, only one instance of the object is created in the MDB. However, objects that represent each of the device's IP addresses appear on the 2D Map and in the Topology Browser.

**Note:** The only multi-homed devices that Continuous Discovery supports are routers.

**Note:** By default, interface objects are not created for non-router and single NIC devices. To enable this, set the `InterfaceForAllDevices` registry key under `HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Discovery` to true. This key is set to false by default.

## Discovery Classification Engine

After Discovery discovers devices on your network, the classification engine then classifies these devices according to how you have configured the classification engine. Continuous and Classic Discovery use the same classification engine. The classification engine configuration files let you customize the discovery rules to your environment.

Classification means that a class and subclass is defined for each discovered object, it is added to the MDB, and you can manage the object using CA NSM.

## Discovery Timestamp

Discovery maintains a "last seen on the network" timestamp for each device in your network.

This timestamp can help you determine if a device should not be monitored anymore and if it can be removed from the system. Using the information in this timestamp, you can define usage policies. For example, you may conclude that a device that has not been seen by the Discovery process in more than 45 days is no longer valid. This information is stored in the MDB so that you can run queries against this value for inventory reports.

The "last seen on the network" timestamp updates a property that contains the length of time that a device was not seen by the Discovery process. This time can be determined by when a device was last successfully accessed or when network traffic was last seen from the device.

Whenever Discovery runs a ping sweep and successfully addresses an object, it saves this information. The Classic Discovery process updates the MDB immediately, and the Continuous Discovery process holds this information in the Discovery agent caches until the Discovery Manager requests this information. You can configure the Discovery Manager to poll the agents for this information on a regular basis.

## How Subnet Filters Work

Using a subnet filter on large networks with multiple large subnets is advantageous because you can limit your search to certain subnets within the network, which can mean a shorter Discovery process.

Use a subnet filter to do the following tasks:

- Limit the scope of Discovery by confining it to a certain range of subnets and devices. For example, if you use the subnet filter `172.24.*.*`, for example, only the subnets from `172.24.1.0` to `172.224.255.0` are searched. If there is a subnet called `172.119.1.0`, that subnet is not searched because it does not fall in the range specified by the subnet filter.
- Enter a range of as many as ten filters. The filter statement uses a comma separated format of `a1.b1.c1.d1,a2.b2.c2.d2,...a10.b10.c10.d10`. Only those subnets passing through filter1 (`a1.b1.c1.d1`) or filter2 (`a2.b2.c2.d2`) or filter $n$  (up to 10) will be searched and created as `TNG/IP_Subnet`.
- Use the default subnet filter of `*.*.*.*`, which does not limit the scope of the Discovery process. After the selected subnets are searched by Discovery, they are placed on a list in the right pane of the Discovery Subnet Management dialog.

You can also use the `dscvrbe` command to limit Discovery to specific subnets to discover, a range of IP addresses to discover, a range of subnets to discover, or subnets to exclude subnets from the Discovery process. You define the subnets in a text file and then specify this text file using the `-8` parameter of the `dscvrbe` command. This functionality is available only when you use the `dscvrbe` command.

## How Timeout Values Affect Discovery

The values you specify for SNMP timeout and Ping timeout greatly affects how successful and how long your Discovery takes to run. If you set higher timeout values, Discovery takes longer to run, but has plenty of time to communicate with the devices and obtain the needed information. If you use lower timeout values, Discovery runs faster, but devices may not be classified correctly, or even discovered at all.

If you are discovering routers, we recommend that you use higher SNMP timeout values. You can specify the timeout value in any of the following places, depending on how you are running Discovery:

- Using the command line by specifying the `-W` parameter on the `dscvrbe` command
- Using the Management Command Center Discovery or Advanced Discovery Wizard Timeouts page
- Using the Discovery Classic GUI Timeouts box on the Discovery page of the Discovery Setup dialog

## Discovery Object Creation Rules

You can define rules to help you create objects in the MDB. These rules are used by the classification engine that is used by Classic Discovery and Continuous Discovery to help you expand and customize the classification and discovery of objects.

You can write object creation rules to perform the following tasks:

- Reclassify objects using the `sysObjectID` in the MIB-II.
- Reclassify objects using additional ports, such as FTP, Telnet, SMTP, and HTTP, to gather information.
- Reclassify objects by reading the agent MIB for a descriptor to define the class.

## Types of Discovery Methods

You can write object creation rules that support the following methods of classification. You can combine these methods using logical ANDs and ORs to classify an object as exactly as possible.

You define these rules in configuration files, `methods.xml` and `classifyrule.xml`, that are used by the Discovery process to classify the object before it is added to the MDB. Classification rule files are in XML format and reside in the `\Config` subdirectory on the Discovery agent.

By tuning the priority and timeout properties, you can configure the system for optimal classification performance. You should give the most successful rule in your environment the highest priority. By default, the highest priority rule is Generic SNMP. However, if you have very few native SNMP installs in an environment, you should give a different rule the highest priority. For example, if you have many CA NSM agents installed, the `SNMPAgentOID` rule results in the best performance.



**SNMP**

Uses a certain port and community string for classification. You can customize your rule files by adding a new general method to the methods.xml file or by changing the existing SNMPGeneric string. All pattern matches for the results of SNMP queries are specified in the classifyrule.xml file. Review the classifyrule.xml file for more information about how to classify by evaluating SNMP query results.

**Telnet reply pattern match**

Attempts to establish a Telnet session, and returns the Telnet login screen if successful. In the classification methods, this screen can then be matched with a default pattern. The Telnet method could also be described as “screen scraping” of the Telnet login screen. Default classification rules are supplied for all major operating system vendors. In many environments, these login screens are standardized. You can modify the Telnet classification rules by entering your own pattern matches if you have specialized login screens. Telnet methods specify a state computer that usually consists of establishing the connection and then waiting for the amount specified in the timeout parameter. After the timeout is reached, the connection can be closed.

**UDP/TCP port scanning (socket)**

Socket type methods scan ports of a computer to retrieve a port map that can be used to identify what type of device was discovered on the network. The desired port combination can be defined in the classifyrule.xml file (see this file for examples). In the port combination, you can specify whether a port should be found at all. For example, the absence of a Telnet port may signify that the device could be a Windows computer. You can now combine this rule with the NetBios port scan (SocketWindows\_NetBios method) to describe the port layout of the computer so that the computer can be classified as correctly as possible. You can configure port scans for TCP/IP or UDP. You can specify pattern matches in the classification rule in the classifyrule.xml file if you know the byte pattern.

**MAC address patterns**

Specifies the first six bytes of a MAC address in the filter of a classification rule.

**HTTP response pattern match**

Queries a computer using the HTTP protocol and returns the response. The response is matched with a byte pattern in the classifyrule.xml file. Default methods are provided by Discovery.

**SMTP**

Attempts to establish an SMTP session with a mail server. The SMTP method is very similar to the Telnet and FTP methods. You can customize this method to fit different types of mail servers. The default method supplied by the default Discovery configuration files works for Microsoft Exchange Mail servers.

## FTP

Attempts to establish an FTP session with the computer and returns the FTP login screen. The FTP method is very similar to the Telnet method.

## How You Modify or Write Classification Rules

You can modify the classification rules in the classifyrule.xml file provided with the CA NSM, or you can write additional Discovery classification rules to refine the classification of your network devices.

To write Discovery classification rules, follow this process:

1. Modify existing rules or add new rules to the classifyrule.xml file in the *discovery\_install\config* directory.
2. For Continuous Discovery, move classifyrule.xml and methods.xml to the \config folder on each Discovery agent for which the rules apply. Then restart the agent.

**Note:** Each Discovery agent can have a different set of rule files. You may want to do this if you want to each agent to classify devices differently.

For Classic Discovery, run ruletodbconverter.exe on the classifyrule.xml file in the *discovery\_install\config* directory on the computer where the MDB resides. Then run Discovery.

**Note:** For Classic Discovery, you do not need to move the rule files from the *discovery\_install\config* directory. Only one set of rule files is required.

## How to Enable Classification of New Classes

After you write new class rules, you can enable classification of these rules using the following process:

1. Create a new class in WorldView using ClassWizard or TRIX with an associated sysObjID.
2. Open a command prompt, and navigate to DISC\_INSTALL\_PATH\bin.
3. Run the updateclassrules utility.
4. Run the ruletodbconverter command.

## methods.xml file--Configure Classification Methods

The methods.xml file defines the methods that are used by the Discovery classification engine to classify discovered devices. The methods are based on existing protocol plugins supported by Discovery. The methods.xml file is located in the *discovery\_installdir\config* folder, and contains the following methods:

**SNMPGeneric**

Uses the common MIB-II sysobjid entries to classify a computer. SNMP must be installed on the computer that is to be discovered. Contains the following parameters:

**Port**

**Default:** 161

**Community**

**Default:** public

**Timeout (in milliseconds)**

**Default:** 2000

**SNMP\_AgentOID**

Finds Agent Technology common services (if aws\_admin was configured to respond to SNMP requests). Contains the following parameters:

**Port**

**Default:** 6665

**Community**

**Default:** admin

**Timeout (in milliseconds)**

**Default:** 4000

**SNMP\_SysEdgeAgentOID**

Finds active CA SystemEDGE agents and evaluates their operating system information. Contains the following parameters:

**Port**

**Default:** 1691

**Community**

**Default:** public

**Timeout (in milliseconds)**

**Default:** 4000

### **SNMPSuspect\_AP**

Finds special wireless access points in an environment. If there are none, remove this method from the configuration file for better performance. All references to a removed method must be deleted from the classifyrule.xml file. Contact Technical Support for help with this type of rule modification.

Contains the following parameters:

#### **Port**

**Default:** 161

#### **Community**

**Default:** public

#### **Timeout (in milliseconds)**

**Default:** 2000

### **SocketWindows\_DS**

Scans for the Windows domain server port. Contains the following parameters:

#### **Port**

**Default:** 445

#### **InitDataLength**

**Default:** 100

#### **Timeout (in milliseconds)**

**Default:** 2000

#### **TCP**

**Default:** True

### **SocketWindows\_NetBios**

Scans for the Windows NetBios port. Contains the following parameters:

#### **Port**

**Default:** 139

#### **InitDataLength**

**Default:** 100

#### **Timeout (in milliseconds)**

**Default:** 2000

#### **TCP**

**Default:** True

**SocketUnix\_RPC**

Scans for the RPC port, which is a common port on Sun Solaris computers. Contains the following parameters:

**Port**

**Default:** 111

**InitDataLength**

**Default:** 100

**Timeout (in milliseconds)**

**Default:** 2000

**TCP**

**Default:** True

**SocketSuspect\_AP**

Scans ports for suspect access points. Contains the following parameters:

**Port**

**Default:** 192

**InitDataLength**

**Default:** 116

**Timeout (in milliseconds)**

**Default:** 2000

**TCP**

**Default:** False

**HTTPGeneric**

Sends a generic HTTP request to a device. User ID and password are not specified. If the device has a web service that returns a login screen or an error screen, a pattern in that response can be matched with classification rules in classifyrule.xml. Contains the following parameters:

**Port**

**Default:** 80

**Timeout**

**Default:** 1000

**UserID**

**Default:** " "

**Password**

**Default:** " "None

### **HTTPAuthenticate**

Sends an HTTP authentication request to a device. Contains the following parameters:

#### **Port**

**Default:** None

#### **Timeout**

**Default:** 1000

### **TelnetWithSend**

Attempts to establish a Telnet connection and sends bogus data. Some specialized devices will not acknowledge Telnet commands without a subsequent send. Contains the following parameter:

#### **Timeout**

**Default:** 1000

### **TelnetGeneric**

Attempts to establish a Telnet connection and returns the Telnet login screen if successful. Contains the following parameters:

#### **Port**

**Default:** 23

#### **Timeout**

**Default:** 5000

### **FTPGeneric**

Attempts to establish an FTP session with the computer and returns the FTP login screen. This method is very similar to the Telnet method. Contains the following parameters:

#### **Port**

**Default:** 21

#### **Timeout**

**Default:** 1000

**SMTPGeneric**

Attempts to establish an SMTP session with a mail server. This method is very similar to the Telnet and FTP methods. You can customize this method to fit different types of mail servers. The default method supplied by the default Discovery configuration files works for Microsoft Exchange Mail servers. Contains the following parameters:

**Port**

**Default:** 25

**Timeout**

**Default:** 1000

**ClassHint**

(Used only for Continuous Discovery and should not be modified.) Re-uses previously discovered data, and uses some limited SNMP queries in the first discovery phase to find system information such as host names, router flags, and multiple IP addresses. If the ClassHint method is specified, it reuses previously gathered information for classification purposes. Contains the following parameters:

**Port**

**Default:** None

**Timeout**

**Default:** None

## classifyrule.xml--Configure Classification Rules

The classifyrule.xml file defines the classification rules that associate class names with methods. A device can be classified by any of the classification rules. Each classification rule consists of a sequence of methods that is used to classify the device.

Rules are organized by device class and can include the following information:

- Class relationships--combination classes, and child classes (subclasses)
- Filters--applied to method results and are specific to the methods they support
- Method instances--parameters that are listed for the discovery methods in methods.xml that are modified in the class rules because they differ for specific classes

For example, you may want to set the SNMP port for all UNIX computers to 161, but on Windows computers, you want to set the SNMP port differently. To do this, you can define a new global method for UNIX and one for Windows in the methods.xml file, or you can change the port parameter for the Windows classes in the classifyrule.xml file.

- Priority parameter--designates the precedence of rules to apply when classifying a device. The lower the priority value, the higher the rule is in the hierarchy. For example, a rule with a priority value of 1 would be applied first and take precedence over a rule with a priority value of 2.
- A rule may also contain a prerequisite device name, which when combined with the method, is used to classify the device.

The classifyrule.xml file is located in the *discovery\_installdir*\config folder on the Discovery agent computer. You can define as many classification rules as needed in the classifyrule.xml file.

### Examples: Simple SNMP and Subclasses

Here is an example of a simple SNMP rule:

```
<Device Class="WindowsNT" ClassScheme="Operating System">
  <ClassificationRule Enabled="1" Priority="1">
    <Method Name="SNMPGeneric">
      <Filter Type="RegExp">((SysOID REGEX "1.3.6.1.4.1.311.1.1.3.1")||
        (SysOID REGEX "1.3.6.1.4.1.311.1.1.3.1.1"))&&(SysDescr REGEX "Windows NT
Version 4.0")
      </Filter>
    </Method>
  </ClassificationRule>
</Device>
```



In the previous example, an object is classified as a Windows NT computer if the method `SNMPGeneric` (which is defined in the `methods.xml` file) returns with the value `1.3.6.1.4.1.311.1.1.3.1` in the `sysobjid` field of MIB-2 and the system description field in MIB-2 returns a string that contains "Windows NT Version 4.0". The `ClassScheme` is a reference for the classification hierarchies available in MDB. For more information, see the MDB schema description for the table `ca_class_hierarchy`.

Here is an example of a rule that contains subclasses:

```
<Device Class="Unix" ClassScheme="Operating System">
  <Relation DeviceName="RISC6000" Type="child"/>
  <Relation DeviceName="Solaris" Type="child"/>
  <Relation DeviceName="HPUnix" Type="child"/>
  <Relation DeviceName="DG_UX" Type="child"/>
  <Relation DeviceName="Linux" Type="child"/>
  <Relation DeviceName="NCRUnix" Type="child"/>
  <Relation DeviceName="UnixWare" Type="child"/>
  <Relation DeviceName="SCOUnix" Type="child"/>
  <Relation DeviceName="Silicon" Type="child"/>
  <Relation DeviceName="SiemenUX" Type="child"/>
  <Relation DeviceName="FUJIUxp" Type="child"/>
  <Relation DeviceName="Sequent_Server" Type="child"/>
  <Relation DeviceName="OpenVMS" Type="child"/>
  <Relation DeviceName="ICLUnix" Type="child"/>
  <ClassificationRule Enabled="1" Priority="3">
    <Method Name="SocketUnix_RPC">
      <Filter Type="RegExp">00</Filter>
    </Method>
  </ClassificationRule>
</Device>
```

In the previous example, all child classes that are allowed for a parent class are listed in the classification rule. Listing child classes tells Discovery to execute additional rules, and enables Discovery to determine when a final, best fit rule can be determined. Specifying this type of rule applies mostly to Continuous Discovery because classification is an ongoing effort. In the MDB, the class hierarchy that works in the same way. Many products' rules are evaluated to find the best classification for any common object in MDB. For more information, see the MDB schema description.

## Device Not Discovered

**Symptom:**

A device was not discovered.

**Solution:**

Check the following:

- The SNMP get community name may be incorrect. Use ObjectView to verify.
- The SNMP agent on the device is not running. Start the SNMP agent on the device.
- Check for a conflicting IP address in the system. Remove the conflicting device from the system.
- The SNMP agent is accepting SNMP traffic from only a particular set of IP addresses. Check the Access Control List found in that device's SNMP security setup utility. This is normally used as a form of security to allow only the Network Management software, running on particular IPs, to access and change the SNMP information. The actual utility and method for changing security is device-dependent and vendor-supplied.
- If a particular device was discovered but placed in the UNCLASSIFIED\_TCP class, then one of two situations has occurred:
  - No classification rule could be successfully executed for the device. In this case you should create a new rule or class that lets Discovery correctly identify the device. To add a new rule edit the classification configuration files located in the Discovery\Config subfolder. If no appropriate class exists for the device, create a new class using the WorldView Class Editor. After the rules and classes are added, run the RECLASS command to reclassify the UNCLASSIFIED\_TCP objects to the newly created classes or run Classic Discovery again. The RECLASS utility will reclassify only the objects that match the newly created class's sysObjectID. When you run RECLASS, you will be prompted to specify the logical repository, its database ID, and password, if any.
  - Discovery supports the DHCP environment. If your site is using DHCP, you can check the DHCP checkbox in the Setup Dialog or in the Wizard. You also need to specify the DHCP Scope from the Wizard or by creating objects in the DHCP\_Scope class from the Object Browser. The wild card character (\*) can be used to specify the starting and ending IP address. You need to provide both the starting and ending address for each DHCP\_Scope object. You can create one object that has starting and ending addresses of \*.\*.\*.\* to indicate the whole IP network is using DHCP.

**Note:** If your network is very dynamic, you should check the DHCP/Dynamic Network check box.

- If you are running Discovery from a Windows XP system, the port scan methods (SocketWindows\_DS, SocketWindows\_NetBios, SocketUnix\_RPC, SocketSuspect\_AP) are taking a much longer time than on other operating systems. You may need to increase the timeout values for all the socket type methods in the methods.xml files until the device is correctly classified as UNIX or Windows.

## Discovering Your Network Devices Continuously in Real-Time Mode

Continuous Discovery is event-driven and ongoing. Discovery agents continuously scan your network for the addition of new devices or changes in IP addresses for existing devices. To use the Continuous Discovery feature, you must install the Discovery Manager and Discovery agents. The Discovery agents report to the Discovery Manager, which then updates the MDB.

You can configure the Continuous Discovery agents and the Discovery Managers to best suit your environment by setting instance-level properties, such as the configuring the following options:

- Optimal load balancing between the Discovery Agents and the Discovery Manager
- Agent monitoring of subnets other than the local one
- DHCP listening on the Discovery agents or the Discovery Manager

### Continuous Discovery Architecture

Continuous Discovery consists of the following components:

- **Discovery agents** act as data collectors and, by default, monitor only the local subnet.
- **The Discovery Manager** connects to the MDB through a WorldView client or is installed directly on the computer where the MDB resides. The latter is recommended because it greatly reduces network traffic.

The Discovery Manager controls multiple Agents. The Discovery agents discover and classify devices, while the Discovery Manager consolidates the discovered information from the agents and interacts with the MDB through WorldView. The Discovery Manager also distributes the workload of Discovery agents. A Discovery agent's workload is simply the list of subnets the agent monitors. Discovery and subsequent classification is restricted to this list of subnets.

After a device is discovered by a Discovery agent, it is classified. Classification is rule-driven, which facilitates quick additions and updates of classification rules without needing new code modules, for example, libraries. Classified devices are sent to the Discovery Manager, which updates the WorldView repository in the MDB.

The Discovery Manager updates the WorldView repository with any information received from Discovery agents, and also registers with the WorldView repository for any notifications regarding the network entities. The Discovery Manager and Discovery agents "handshake" with each other at the start of Continuous Discovery. During this "handshake," the Discovery Manager distributes subnets and their constituent devices already available in the WorldView repository to the agents. If multiple agents exist, the Manager distributes subnets to the agents, that is, performs a "load balance" to ensure that all agents have an optimal load.

The Discovery Manager and Discovery agents maintain their device information in caches. They communicate with each other using a messaging service.

Continuous Discovery is implemented as services or daemons, which enables real-time discovery. The discovery mechanism employs various methods, such as DHCP, ICMP, SNMP, ARP cache of router scans and sniffing network packets using CTA. Additional discovery methods can be accommodated using the plugin interface, where components can be "plugged" into the framework.

## How Continuous Discovery Monitors Your Network

The Continuous Discovery agent uses the same network protocol query-based mechanisms as the Classic Discovery component. Continuous Discovery uses the following monitoring options to determine what changed in a network environment:

- **Ping**--The Discovery Agent actively pings the IP network at initial discovery and then at regular time intervals (the default is every 24 hours and is configurable) to determine if there are any new devices in the network.
- **ARPcache monitoring**--The Discovery Agent queries its gateway router's ARP cache tables using SNMP to retrieve MAC address/IP address pairs. It uses these pairs to correlate a device's MAC address with its IP address. This correlation is required because some discovery methods, such as ping, only retrieve a device's IP address. For the local gateway router, Continuous Discovery queries using only the "public" community string. You do not need to specify the gateway because it is determined by Continuous Discovery. The only dependency is the community string, so if the gateway uses any other name, the ARP query fails.

- **Common Traffic Analyzer (CTA) network sniffing engine**--CTA is a shared component of CA's sonar technology and is installed by Continuous Discovery if it not already available on the system. Every Discovery agent is installed with the CTA plugin enabled. CTA lets the agent "sniff" traffic from devices on the network to determine MAC address/IP address changes, discover new devices, and attempt to classify the devices.

You can disable the CTA plugin.

- **Dynamic Host Configuration Protocol (DHCP) traffic monitoring (either agent or manager based)**--By default, the Continuous Discovery Manager listens to DHCP traffic on the network for discovery of new devices or changes in IP address/MAC address pairs. To fully utilize this method, configure the local router to redirect DHCP requests to the Discovery Manager host. This configuration lets the Discovery Manager discover devices using DHCP other than on the local subnet. You can integrate Continuous Discovery with third-party DHCP servers.

## How Continuous Discovery Discovers and Monitors Subnets

Continuous Discovery will discover and monitor a subnet in the following situations:

- A Continuous Discovery Manager or agent is located on the subnet. This is the only case in which Continuous Discovery creates a subnet object in the MDB automatically.
- Enable Workload Balancing in the Discovery Manager (workload balancing is enabled by default). Continuous Discovery finds the subnets in the MDB and assigns them to an agent based on a random distribution formula that depends on the agents currently available. The subnet must first exist in the MDB for the Discovery Manager to find and assign the subnet.
- Manually configure the Discovery agents by setting agent properties to discover the subnet and add it to the MDB.

## Continuous Discovery Default Configuration

The default Continuous Discovery configurations is as follows:

- The DHCP server is configured to forward DHCP requests to the Discovery Manager, and the Discovery Manager listens for DHCP events.
- Workload balancing is enabled on the Discovery Manager.
- A local Discovery agent is deployed on the Discovery Manager computer and the Discovery Manager depends on the local Discovery agent.
- The CTA network sniffing engine is enabled on the Discovery agent.
- After installation and startup, the local subnet is automatically discovered.

## DHCP Engine Configuration

The DHCP Discovery Engine listens for DHCP traffic to discover new devices or to reclassify them dynamically. You can configure the Discovery DHCP Engine in a distributed or centralized mode.

In distributed mode, the Discovery agents listen for DHCP requests and update the Discovery Manager if applicable.

In centralized mode, the Discovery Manager listens to the DHCP requests and updates the MDB directly.

The mode you employ depends on your network DHCP configuration, such as how many DHCP servers you have, where they are located, how many Discovery Managers you install, and how many Discovery agents are deployed.

## Set the Admin Status Property for an Object Using Continuous Discovery

Using Continuous Discovery, you may want to discover devices but set their administrative status to Unmanaged. You do this by setting the DeviceDefaultAdminStatus property for a Discovery Manager. The default setting for this property is Managed.

**Note:** In Classic Discovery, you can also set this flag using the -24 parameter on the dscrbe command.

### **To discover all devices in a subnet and set their administrative status to Unmanaged**

1. Open the Management Command Center, choose Class Specification from the left pane drop-down menu, expand the tree until you see ManagedObject, expand ManagedObject until you see ContinuousDiscoveryManager.  
The ContinuousDiscoveryManager object appears.
2. Right-click ContinuousDiscoveryManager and choose Add Viewer, Instances.  
All instances of Discovery Managers appear in the left-pane.
3. Right-click a Discovery Manager instance and choose Open Viewer, Properties.  
The Properties notebook for the Discovery Manager instance appears.
4. Click the RunTime tab, double-click the DeviceDefaultAdminStatus property, and set the property to 1 (Unmanaged).

The DeviceDefaultAdminStatus property for the Discovery Manager is set to 1 (Unmanaged), and the Admin Status property for all devices that are discovered by agents that this manager monitors is set to 1 (Unmanaged).

**Note:** You do not need to restart the Discovery Manager for the property change to take effect.

## Exclude Classes from Discovery

You can exclude certain classes or subclasses of objects from being classified by Discovery agents. For example, this is useful if you are an administrator who is responsible only for network devices such as routers and switches. You can set up class filters using the Object Browser in the WorldView Classic GUI to work with devices that are only in the router or switch classes.

### To set up filters to exclude classes from Discovery

1. Choose Start, Programs, CA, Unicenter, NSM, WorldView, Object Browser.  
The Object Browser appears.

2. In the left pane, expand TNGRoot, Reference, CaMtxReference, and click CaMtxClassFilterEntry.

The CaMtxClassFilterEntry object opens in the right pane. This object contains a list of all the names of classes that are handled by the Discovery Engine.

3. To exclude a class or subclass, find the entry that corresponds to the class, select the entry, choose Object Delete from the main menu, and click OK on the confirmation window that appears.

The class or subclass is deleted from the filter.

Only objects in the classes and subclasses that appear in the filter are used during classification by Discovery agents. Instances of classes that are excluded show up as unclassified TCP/IP devices.

4. Run the updateclassrules utility.

**Note:** Be sure that the classdefinition.xml and classifyrule.xml files are writable before you run updateclassrules.

The classdefinition.xml and classifyrule.xml files in the Discovery\Config folder are updated.

5. (Optional) If the Discovery agent is not local, copy the classdefinition.xml and classifyrule.xml files to the corresponding folder on the Discovery agent computer.
6. Restart the Discovery Manager and Discovery agent.

## How You Set Up SNMP Community Strings for Continuous Discovery

It is a best security practice to change the community strings for SNMP in a network environment. You can set up SNMP community strings for Continuous Discovery in several ways:

- If you want to globally apply the additional community strings all subnets, which may be divided among more than one agent, add the community strings to the Auth table in WorldView on the computer where the WorldView Manager resides. The WorldView Manager will then send these strings to each of the agents managed by it during its handshake.
- If you want to specify community strings according to the subnet to be discovered, you can specify the additional community strings as a parameter in the methods.xml file that is located on the agent computer that manages that subnet. Create a corresponding SNMP-based method and specify the community name in the method. For example, if you want to use the community string "admin" instead of the default (usually "public"), then the SNMPGeneric method in methods.xml would look like this:

```
<Method Name="SNMPGeneric" Type="SNMP" Priority="1">
  <Params Timeout="2000" Community="admin"/>
  <StateMachine>
    <State Name="get" OutputToVar="SysOID">1.3.6.1.2.1.1.2.0</State>
    <State Name="get" OutputToVar="SysDescr">1.3.6.1.2.1.1.1.0</State>
    <State Name="get"
      OutputToVar="IPForwarding">1.3.6.1.2.1.4.1.0</State>
  </StateMachine>
```

If you create a new method, you must also update the classifyrule.xml to use the new method in its classification rules for the corresponding classes. Optionally, you can also edit classifyrule.xml to add the community name as a parameter within a method specification under the classification rules for a particular device. You may want to do this if you want to have all devices of a particular class to be queried through SNMP using a specific community string, but the remaining classes to be queried using the default set of community strings. To do this, in classifyrule.xml, add the community string as a parameter under the <Method> XML element as follows:

```
<Params Community="admin"/>
```



## Discovery Managers

A Discovery Manager provides the necessary communication between WorldView, the MDB, and the Discovery agents. It performs the following functions:

- **Cache Management**--The Discovery Manager stores in memory a view of all discovered objects. The information in the cache is updated as messages and events from the MDB or Discovery agents are received.
- **Discovery Agent Management**--The Discovery Manager is responsible for discovering the Discovery Agents and coordinating the functions that an agent performs.
- **Central DHCP Discovery Engine**--A DHCP Discovery Engine is built into the Discovery Manager because in certain DHCP scenarios, distributed monitoring of DHCP is not feasible. The DHCP Discovery Engine listens for DHCP requests to discover new devices or to reclassify them dynamically.

### Set Properties for Continuous Discovery Managers

You can set runtime (instance-level) properties for the Discovery Managers. All Discovery agents are instantiated from the ContinuousDiscoveryManager class.

#### **To set Discovery Manager instance-level properties using the Management Command Center Properties notebook**

1. In Management Command Center Topology view, expand ManagedObjectRoot, and keep expanding the topology until you see the computer that is running the Discovery Manager, then click the Discovery Manager object that you want to configure.

The manager appears in the right pane in Explorer view.

2. In the right pane, choose Properties from the drop-down list, or right-click the Discovery Manager and choose Open Viewer, Properties from the context menu.

The Properties notebook appears in the right pane.

3. Click the WLB tab and set the following property:

##### **ENABLE\_WLB**

Specifies whether Workload Balancing is enabled. The default value is true.

4. Click the DHCP tab and set the following property:

##### **DHCP\_ENABLE**

Specifies whether the DHCP Discovery mechanism is enabled. The default value is true.

5. Click the RunTime tab and set the following properties:

**DeviceDefaultAdminStatus**

Specifies the default value of admin\_status property for newly discovered devices.

**SubnetFilter**

Specifies the string expression specifying the list of subnets that are to be discovered.

**Note:** For more information about subnet filters, see [Configure a Discovery Agent to Manage Additional Subnets](#) (see page 123).

6. Click the EventMgmt tab and set the following properties:

**Enable\_Discovery\_Events**

Specifies whether all Discovery events (new device events, new subnet events, address change events, and handshake events) are reported to the Event Console. The default is true. If set to false, no Discovery events are sent to the Event Console.

**Enable\_New\_Device\_Events**

Specifies whether new device events are reported to the Event Console. The default is true.

**Enable\_New\_Subnets\_Events**

Specifies whether new subnet events are reported to the Event Console. The default is true.

**Enable\_Address\_Change\_Events**

Specifies whether address change events are reported to the Event Console. The default is true.

**Enable\_Handshake\_Events**

Specifies whether handshaking events between the Continuous Discovery Manager and the Continuous Discovery Agents that report to it are reported to the Event Console. The default is true.

7. Close the Properties notebook.

Properties are saved for the Discovery Manager.

## Discovery Events Reported to the Event Console

By default, the Continuous Discovery Manager sends messages to the Event Console so that you are notified about important Discovery events that occur in your network. The Continuous Discovery Manager reports the following types of events to the Event Console:

- New device events are reported for newly discovered and classified devices. Because Discovery is periodically scheduled to take place on agents, devices discovered in the initial Discovery process are not reported to the Event Console. The Continuous Discovery Manager starts reporting the newly discovered and classified devices after the initial Discovery process is complete, which prevents flooding the Event Console with too many messages.
- New subnet events are reported when new subnets are discovered.
- Address change events are reported when IP addresses change for devices.
- Handshake events are reported to indicate the progress of handshaking between the Continuous Discovery Manager and the Continuous Discovery Agents that report to it.

You can disable sending messages to the Event Console by setting Event Management properties for the Continuous Discovery Manager. For more information, see [Set Properties for Continuous Discovery Managers](#) (see page 121).

## Discovery Agents

A Discovery agent consists of the following components:

- Network sniffing technology (enabled by default)
- DHCP request listener (disabled by default)
- Ping discovery engine (always enabled)
- Classification engine (always enabled)

### Configure a Discovery Agent to Manage Additional Subnets

By default, a Continuous Discovery agent discovers and monitors any newly discovered subnet (added to the MDB by Classic Discovery). You can configure a Discovery Agent to discover and monitor additional subnets by assigning subnets to each agent manually. Workload balancing can be enabled or disabled on the Discovery Manager when you add additional subnets using this method.

**Note:** You can also discover only a subnet using the `dscrbe` command and workload balancing. After the subnet is added to the MDB, workload balancing assigns the subnet to an available agent and the subnet is then monitored by the agent. For more information about the `dscrbe` command, see the online *CA Reference*.

### **To manually add subnets to a Discovery Agent**

1. In the Management Command Center, right-click the agent in the left pane tree, and choose Add Viewer, Properties.

The Properties notebook for the agent appears.

2. Click the RunTime tab.

The RunTime page appears.

3. In the SubnetToManage field, add the additional subnets you want the agent to monitor.

Separate subnets with commas. You can use wildcards to specify subnets. You can also define a range of subnets by separating the range with a hyphen (-). Only one range of IP addresses per subnet is permitted.

Additional subnets to manage are defined.

4. Close the Properties notebook for the Discovery Agent, right-click the Discovery Manager in the left pane, and choose Add Viewer, Properties from the context menu.

The Properties notebook appears in the right pane.

5. Click the WLB tab, set the ENABLE\_WLB property to true, and close the Properties notebook for the Discovery Manager.

Workload Balancing is enabled.

6. Stop and restart the Discovery Agent and the Discovery Manager.

The additional subnets will now be discovered and monitored.

### Example: Valid Subnet Filters

Valid subnet filters are as follows:

`xxx.xxx.xxx.xxx` where `xxx` is a valid number between 1-254.

`*.*.*.*` specifies that any subnet should be monitored by the agent.

`xxx.xxx.xxx.*` specifies that all subnets of `xxx.xxx.xxx` should be monitored by the agent.

`xxx.xxx.xxx.xxx - xxx.xxx.xxx.yyy` specifies a range of IP addresses in a subnet that the agent should monitor.

To specify that a range of IP addresses for three subnets should be monitored, use an entry similar to the following entry:

`172.16.333.0 - 172.16.333.128, 172.16.334.1 - 172.16.334.128, 172.16.335.1 - 172.16.335.128`

## Set Properties for Continuous Discovery Agents

You can set runtime (instance-level) properties for the Discovery agents. All Discovery agents are instantiated from the `ContinuousDiscoveryAgent` class.

### To set Discovery agent instance-level properties using the Management Command Center Properties notebook

1. In Management Command Center Topology view, expand `ManagedObjectRoot` and click the Discovery Agent object that you want to configure.

The agent appears in the right pane in Explorer view.

2. In the right pane, choose Properties from the drop-down list, or right-click the agent and choose Open Viewer, Properties from the context menu.

The Properties notebook appears in the right pane.

3. Click the Classifier tab and set the following properties:

#### **CL\_THREADPOOL\_MAX\_THREAD**

Specifies the maximum number of threads are allowed to run in the agent classifier thread pool. The default value is 16.

#### **RECLASS\_POLLINTERVAL**

Specifies the time period in milliseconds over which the reclassification of unclassified objects in the agent service will occur. The default value is 24 hours.

- Click the DHCP tab and set the following properties:

**DHCP\_ENABLE**

Specifies whether DHCP Discovery is enabled. The default value is false.

**DHCP\_POLL\_INTERVAL**

Specifies the time period in milliseconds over which the agent will poll the DHCP Discovery component for newly discovered devices.

- Click the ICMPBroadC tab and set the following property:

**ICMP\_ENABLED**

Specifies whether ICMP Discovery is enabled. The default value is true.

- Click the SNMPBroadC tab and set the following property:

**SNMP\_ENABLED**

Specifies whether SNMP broadcast Discovery is enabled. The default value is true.

- Click the STANDARD tab and set the following properties:

**ICMPTIMEOUT**

Specifies the timeout value in milliseconds when pinging a device and waiting for a response.

**ICMPRETRY**

Specifies the number of pings sent to each device during Discovery.

**SNMPTIMEOUT**

Specifies the timeout value in milliseconds when pinging an SNMP device and waiting for a response.

**SNMPRETRY**

Specifies the number of SNMP queries sent to each device during Discovery.

**DISCOVERY\_INTERVAL**

Specifies the time period in hours over which the Discovery engine is polled to receive and update newly discovered devices.

- Click the RunTime tab and set the following properties:

**SubNetToManage, SubNetToManage1, SubNetToManage2, SubNetToManage3, SubNetToManage4**

Specifies a list of subnets the agent is to discover. The syntax is the same as for subnet filters, except for IP address ranges.

- Close the Properties notebook.

Properties are saved for the Discovery agent.

## Change Continuous Discovery Agent Polling Interval

By default, the Continuous Discovery agent has a one hour polling cycle, that is, the status of the agent is updated once per hour.

If you need to change this polling interval, edit the following registry key and restart the Continuous Discovery services:

```
hkey_local_machine\software\computerassociates\discovery\SharedComponents\  
AgentService\AgentHandler\poll_interval
```

## Change Continuous Discovery Agent Manager

You can change the Continuous Discovery Manager that a Continuous Discovery Agent points to by editing a registry entry on the agent machine.

### To change Continuous Discovery Agent Manager

1. Enter the new manager name in the following registry entry on the agent machine:  
HKEY\_LOCAL\_MACHINE\SOFTWARE\ComputerAssociates\Discovery\SharedComponents\AgentEngine\AgentManager.

**Note:** Do not edit the OldManager entry for this purpose.

2. Restart the Continuous Discovery Agent service.

**Note:** You do not need to restart the Continuous Discovery Manager service.

Once the service is restarted, the previous manager machine displays the new manager as the agent manager in the 2D Map on the run time tab of the agent icon, and the agent icon enters the unknown state.

If the manager service of the newly assigned manager is running, the agent's icon appears on the new manager machine's 2D Map in the normal state with the newly assigned manager displayed as the agent manager.

**Note:** You cannot use the agent manager field in the agent notebook to change an agent's manager. You must edit the registry key to make this change.

## Discovery and Firewalls

Discovery depends on SNMP for some of its classifications. Firewalls that have the default SNMP ports disabled will prevent some computers from being classified properly if they are located behind such a setup. Possible workarounds would be to use Continuous Discovery by placing a Continuous Discovery agent behind the firewall and opening the CAM port for communication across the firewall.

Discovery also uses the default ICMP ports on the agent computer to find devices. Other discovery mechanisms such as HTTP and port scans also require that ports be open if run from behind a firewall, although we do not recommend this approach.

## Continuous Discovery Rapidly Consumes Memory

### **Symptom:**

Continuous Discovery consumes a lot of memory and creates incorrect subnets.

### **Solution:**

Continuous Discovery retrieves the local subnet from the agent computer. If the agent is running on a computer with an incorrect subnet configuration, Continuous Discovery calculates the IP addresses of subnets and devices incorrectly, which results in incorrect data in the Continuous Discovery agent, causing it to rapidly consume memory.

Verify that the subnet mask is set correctly. The default subnet mask that Classic Discovery uses is 255.255.255.0. If the subnet mask is different, specify the `-N subnet-mask` parameter in the `dscvrbe` command, or set the subnet mask on the Discovery Wizard or Advanced Discovery Subnets Page in Unicenter MCC. Classic Discovery also retrieves subnet masks from a router if SNMP is enabled. Be sure that subnet masks are set correctly for all router interface cards that were discovered.



## Discovering Your Network Devices on Demand Using Classic Discovery

Classic Discovery is the Discovery process that you can set up and run on demand to find and classify network devices and then place them in the MDB as managed objects. Classic Discovery lets you discover and classify devices on IP and IPX networks.

Classic Discovery lets you decide which subnets you want to discover and when. You can start a Classic Discovery from the Discovery GUI, the Management Command Center, the Unicenter Browser Interface, or the command line (dscvrbe).

**Note:** You use Classic Discovery if you did not install the Discovery agents and the Discovery Manager, which continuously discover your network. If you are using the Continuous Discovery method, you do not need to run a Classic Discovery.

The IP Discovery process consists of the following main functions:

- **Ping**--IP Discovery identifies whether a network device exists and is able to communicate. The Internet Control Message Protocol (ICMP) uses the ping utility to send requests to the designated computer at periodic intervals and waits for a response.
- **Simple Network Management Protocol (SNMP)**--After receiving a response and confirmation that a network device is valid, IP Discovery issues an SNMP request to the network device. This request asks for specific Management Information Base (MIB) information, which is used to classify and gather information about the network device.

Object descriptions and relationships based on the information in the device's SNMP Management Information Base (MIB) is then used by IP Discovery to create a managed object for this network device in the MDB. SNMP MIB agents typically are resident in network device firmware and are provided by each device's vendor.

Discovery also determines if a device provides Web-Based Enterprise Management (WBEM) data, and if so, creates a WBEM object in the device's Unispace. The Agent Technology WorldView Gateway service locates agents running on the network objects.

**Note:** An MDB must exist before you can run Discovery to discover your network devices and populate the MDB.

Once defined, you can view, monitor, and manage these objects and their Management Information Base (MIB) through the 2D Map, ObjectView, and the Topology Browser. You can manage the entities they represent using Event Management, Manager/Agent Technology, and third-party manager applications.

## Discovery Methods

You can use any of the following Classic Discovery methods to discover your network:

### **ARP Cache**

Starts at the gateway address (the address of the nearest router to the computer running Discovery) for the current subnet and uses the ARP (Address Resolution Protocol) Cache of that device to determine information about the devices. The ARP Cache contains the IP-to-MAC (physical network) address mappings.

Discovery retrieves the gateway address from the computer on which it is running and gets the IP list from the ARP Cache on that router. It then discovers the subnets nearest that router and for each subnet it discovers, queries its gateway, doing the same thing over and over again.

For each device found in the ARP Cache, an SNMP request is initiated. If the device does not respond, it is assumed to be a non-SNMP device, just the IP address is retrieved, and the object is created as an Unclassified\_TCP object.

### **Ping Sweep**

Pings all of the devices on the network based on the subnet mask, finds IP devices, and then retrieves SNMP information. If no SNMP information is retrieved, just the IP address is retrieved, and the object is created as an Unclassified\_TCP device. This is the slowest but most thorough method.

### **Fast ARP**

Similar to ARP Cache, Fast ARP saves time by checking only the ARP Cache of routers. Fast ARP is the best method for updating the MDB when you do not want to use the more intensive searches provided by Ping Sweep and ARP Cache. This is the fastest way to discover your network.

### **DNS Search**

Limits the discovery of devices to those that are defined in the domain name server (DNS). The IP address of each of these devices is combined with the defined subnet mask to determine whether or not to discover the device. (In contrast, the Ping Sweep option tries to discover all active devices numerically, without regard to their definition in the DNS).

Each Discovery method has advantages and disadvantages. The Ping Sweep method provides more comprehensive quantitative information—in the form of the number of devices—because every device on the network is pinged. Even devices not recognized by the router, which may not be discovered through the ARP Cache method, can be discovered using Ping Sweep.

On the other hand, ARP Cache provides the MAC and IP address information on all the devices that are found in the ARP Cache of the router. Ping Sweep, however, generates additional network traffic and is thus more time consuming than ARP Cache and Fast ARP. Sometimes, to discover every device in the network, a combination of Ping Sweep and ARP Cache is required.

We recommend that when you first install your product that you run a Ping Sweep Discovery so that a comprehensive search of your network is done. Periodically, it is a good idea to run an ARP Cache Discovery to check your network for devices added after the initial Discovery was done.

## How Agent Discovery Works

Agent Discovery occurs automatically after you install Agent Technology and start the Agent Technology services. The Agent Technology WorldView Gateway service locates agents running on the network objects discovered during Discovery, IPX Discovery, and SAN Discovery.

After you install Agent Technology and start the Gateway services, you must close WorldView so that the Distributed State Machine Gateway can get information about the nodes that WorldView has discovered. When you restart the 2D Map, managed objects appear for all of the agents and agent objects that have been discovered.

## How IPX Discovery Works

IPX Discovery searches for Novell NetWare servers and network segments that use the IPX protocol. Objects representing the NetWare servers and their associated segments are stored in the MDB. The objects appear on the 2D Map under the IPX Network icon.

- For each server the Discovery process finds, an `IPX_Host` object is created and stored in the CA MDB. The physical MAC address and its segment from the primary physical interface identify the server. In addition, the server has the following properties: a server name, an IPX/SPX version, a virtual MAC address, and a virtual segment address. The virtual information is retrieved from the first virtual IPX interface in the server. Each server object is included in a segment object.
- The segment object name refers to the segment of the first nonvirtual interface found in the server. Each server can then have one or more physical interfaces (`IPX_Generic_Interface`). Each physical interface entry in the `IPX_Generic_Interface` class has properties pertaining to the LAN card itself. The `IPX_Generic_Interface` properties are as follows:
  - MAC Address and segment; both of which are physical. LAN cards do not have virtual segments or MAC addresses. These are found only on the server.
  - IRQ and DMA channels for the interface.
  - A text description by the manufacturer.
- If IPX Discovery finds a NetWare server already in the CA MDB, it ignores it and moves on to the next server with no interruption. If an existing server is found with new interfaces installed that were not previously discovered, the new interfaces are added to the CA MDB.
- IPX Discovery can run concurrently with Auto Discovery, which uses SNMP and TCP/IP protocols. When Auto Discovery and IPX Discovery find two or more objects with matching MAC addresses and different interface types, a `Multi_Protocol_Host` object can be created from them using the utility `multi_if`. Use `multi_if` after running Discovery and IPX Discovery to create relationships between two servers with different protocols that share the same MAC address.

## How SAN Discovery Works

SAN Discovery locates SAN fabric devices (fiber-enabled switches, bridges, and hubs) and SCSI devices connected to the SAN fabric. It also creates Business Process Views in the WorldView 2D Map of the SAN fabric and the links between the objects.

**Note:** SAN Discovery is available in CA NSM and some other CA solutions. However, discovery of SAN hosts has been disabled in CA NSM.

- SAN Discovery runs as a part of the Discovery process and writes informational as well as error messages to the Windows Event Log. SAN Discovery launches automatically during install, or you can launch it manually from the Distributed Services window as part of the CA-AutoDiscovery Service (Classic Discovery), or using the Discovery Wizard or Advanced Discovery tool in the Management Command Center.
- SAN Discovery requires prior discovery of SAN objects by Auto Discovery. It uses SNMP to gather additional information from those SAN objects. If no SAN objects exist on your network, the SAN Discovery ends quickly and logs a message stating that there were no SAN fabric elements (SAN objects) found, or that no SAN fabric exists, and there is nothing more for SAN Discovery to do.
- SAN fabric elements are those in classes designated as SAN classes. SAN classes are defined as subclasses of Switch, Bridge, or Hub classes. New SAN classes can be added easily using the Class Wizard. SAN classes require one additional class level property (CLP): SAN, Boolean datatype set to TRUE.
- You can change default run time parameters using the Classic Discovery GUI, the Management Command Center, or the Unicenter Browser Interface, or you can pass parameters to the service on the command line. (Command line arguments are not saved as defaults and only apply to the current execution of the service.)
- If any SAN components are discovered, the SAN Discovery creates a SAN Business Process View. Under this Business Process View, a set of folder objects contain SAN fabrics, SAN collections (containing one or more linked fabrics), SAN-enabled devices grouped by device type, SAN devices discovered by SAN Discovery (as opposed to discovered by IP discovery), SAN devices that have no SAN links, and an Enterprise SAN object containing all linked SAN devices. The SAN fabric object's label name is based on the Fabric ID of the principal SAN switch. All objects in the Business Process View—fabric elements and SCSI devices connected directly or indirectly to each other—share the same fabric ID.

After the initial SAN Discovery is run, you can rerun SAN Discovery manually using any of the following methods:

**Discover SAN Devices Only**

Executes an IP Discovery on the subnets you specify and only SAN devices are added to the MDB.

Once the device discovery is complete, SAN links are determined. SAN Discovery uses the newly discovered SAN objects as well as any already existing in the MDB to determine the SAN configurations within the subnets that were searched. The SAN configurations can be composed of IP and non-IP (SCSI) enabled devices.

**Typical IP Discovery**

Executes an IP Discovery on the subnets you specify. The SAN devices are discovered and identified during the Discovery process.

**No IP Discovery - Refresh SAN Links only**

Re-determines the links of previously discovered SAN components in the subnets you specify. IP Discovery is bypassed and SAN Discovery uses only those objects already present in the MDB to determine the SAN configurations in the specified subnets.

## How Discovery Uses Subnets

Discovery is used to find the devices that exist in your IT infrastructure. It organizes these devices into logical structures that can be easily understood, such as networks and subnets. It also automatically determines how the IT infrastructure is interconnected to produce a topology.

Typically, you will want to start managing your IT infrastructure by discovering the gateway (routers) on your network. Discovery maintains a list of the subnets it finds during the search process. After the routers are discovered successfully, you can select and clear the subnets you want to discover.

## How Discovery Handles a New Subnet

If you create a new subnet, an IP\_Subnet object is automatically created in the MDB. If you delete a subnet, Discovery does not automatically delete the corresponding IP\_Subnet in the MDB. You need to use the 2D Map or Object Browser to delete that subnet under IP\_Subnet class.

## How You Prepare to Run Discovery

Before running Discovery, use this checklist to ensure these prerequisites have been met:

- The computer from which you are running Discovery must be connected to the network and have a valid IP and gateway address. You can ping the gateway address to ensure TCP/IP connectivity.
- The computer that contains the MDB is running and you can connect to it.
- You have the correct SNMP community names for all of your devices in the MDB. The community name is case-sensitive. The default community name is public.

**Note:** If you are running in a non-English speaking environment, you may need to change "public" to its corresponding spelling in the language of choice. For example, you may have to specify "publico" for Spanish-speaking environments.

- You have the correct SNMP GET community names of the routers. See [Verify Correct Community Names for Routers](#).
- The Host IP Address and the Gateway Address are displayed in the Discovery Setup dialog on the Discovery tab.

If you do not see these addresses, check your Network setup in the Control Panel.

- The subnet filter, subnet mask, and subnet count are set on the Discovery Setup dialog.
- You selected the Enable the Trace Console checkbox if you want to monitor the progress of the Discovery process. See [Discovery Setup Dialog--Service Page](#).

## How You Discover a Single Network

If you have a single class A, B, or C network address, you can limit Discovery to find only the subnets and IP systems in that network.

- Replace the zeros in the subnet filter with an asterisk and provide a value for the subnet mask.
- Discovery starts by discovering the local subnet; that is, the subnet to which the local system belongs.
- Set the subnet mask to the local subnet mask. You can determine the local subnet mask by looking at the Windows Control Panel, Network, TCP/IP Protocol Setup window. Typically, the subnet mask is 255.255.255.0.
- When Discovery finds a router, it uses SNMP to get information on the subnets attached to that router. It adds only those subnets to the MDB that match the subnet filter, and then adds them to the list of subnets to be discovered in the Subnet Management dialog.
- To discover only the subnets that you have defined, specify the `-S 0` parameter on the `dscrube` command. If you specify `-S All` or a number greater than 0, the new subnets that are found in a router discovery are also discovered until all known subnets have been discovered, or the number specified is reached.

## How You Discover an Entire Intranet

If you have set up an intranet and built the appropriate gateways and firewalls to isolate your intranet from the larger Internet, you can discover your whole Intranet starting with your local subnet and router.

## How You Determine the Time Required to Ping a Class B Network

To determine how long it will take to ping an entire class B network, multiply the following values together:

- Ping timeout value (400 milliseconds)
- Number of pings per device (set during Discovery setup--3 is the default)
- Number of subnets (255)
- Number of devices per subnet (255)

For example, 400 times 3 times 255 times 255 equals 78,030,000 milliseconds.



## How Names of Discovered Devices are Determined

You can decide where Discovery gets device names in your network. The names of devices discovered are determined by the parameter used on the `dscvrbe` command, or the options you set when using the Management Command Center Advanced Discovery tool or the Classic Discovery GUI. You can use the DNS name, the IP address, or the `sysName` to name your devices.

**Note:** If you configure Discovery to remove the DNS suffix, it is only removed in the label of the object, not in the object name property. The object name property is an important key for cross-product reference in the MDB.

To use the DNS name, set one of the following options:

- On the `dscvrbe` command, use the `-F` parameter, which means to use the DNS name, if it exists, from the domain name server/host file as the object's name.
- On the Options page in the Advanced Discovery tool in the Management Command Center or the Discovery page in the Classic Discovery GUI, select the Use Domain Name Server/Host File option.

To use the IP address, set one of the following options:

- On the `dscvrbe` command, use the `-J` parameter, which means to use the Internet Protocol (IP) address as the object's name instead of the MIB-II `SYSNAME`.
- On the Options page in the Advanced Discovery tool in the Management Command Center or the Discovery page in Classic Discovery GUI, select the Use IP Address Instead of `sysName` option.

If no parameter or option is set, the `sysName` (from SNMP) is used. If the `sysName` is not available, the IP address is used as the object's name.

The names of devices discovered are obtained using the Use Domain Name Server/Host File option, which is enabled by default. This option lets Discovery call the socket function `gethostbyaddr()` to resolve the device name. On Windows, this function checks the local host file and then queries the DNS server, WINS, NetBIOS, and so on (depending on Windows network properties).

If the device's IP address has a DNS name, the DNS name becomes the object's name. If the device's IP address does not have a DNS name, if the Use Domain Name Server/Host File option is disabled, or if the `-J` flag specifies IP, then the IP address is used to name the object. Otherwise, Discovery checks to see if the device is SNMP agent-enabled. If the device is SNMP agent-enabled, the MIB-II `SYSNAME` value is used for the object's name. If the device is not SNMP agent-enabled, the IP address is used for the object's name.

**Note:** If you are combining Classic and Continuous Discovery, see [How You Can Combine Running Classic and Continuous Discovery](#) (see page 101).

## Discovery Creates Incorrect Subnets

### Symptom:

Classic Discovery runs for a long time and places devices under the wrong subnets. When I look at a subnet segment, I see incorrect devices, for example, devices with IP addresses from a different Class C network.

### Solution:

When discovering a subnet with Classic Discovery, verify that the subnet mask is set correctly. The default subnet mask that Classic Discovery uses is 255.255.255.0. If the subnet mask is different, specify the `-N subnet-mask` parameter in the `dscvrbe` command, or set the subnet mask on the Discovery Wizard or Advanced Discovery Subnets Page in Unicenter MCC. Classic Discovery also retrieves subnet masks from a router if SNMP is enabled. Be sure that subnet masks are set correctly for all router interface cards that were discovered.

## Discovering IPv6 Network Devices using Common Discovery

Continuous and Classic Discovery do not have the ability to discover or recognize devices on IPv6 networks. To discover devices on IPv6 networks in your enterprise, you must use Common Discovery, which is a discovery product that is capable of discovering and classifying all devices on IPv6 and IPv4 networks. You can install Common Discovery from the Unicenter Product Explorer.

**Note:** Common Discovery can discover only IPv6 devices for CA NSM.

Common Discovery Import is a WorldView service that populates the MDB with an IPv6 topology. IPv6 (Internet Protocol version 6) is the latest protocol for IP addresses. It supports longer addresses than IPv4, with a size increase from 32 bits to 128 bits.

**Note:** CA Common Discovery is not integrated with the CA NSM database and is therefore not notified about a new class. Also, the `dscvrbe` command, which starts Classic Discovery, does not apply. Instead, CA NSM gets data from CA Common Discovery itself. To add a class, you must manually update the rule classification file (`CmnDscvrClassification.xml`) in the `CACD\Config` directory. Its format is similar to the CA NSM rules file.

You use the Common Discovery Import service to populate the MDB with an IPv6 topology containing IPv6 network devices discovered by Common Discovery.

## Common Discovery

CA Common Discovery is a subsystem that provides the discovery and classification of all entities in your Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) network. It discovers the relationships between these entities and effectively records the network's topology.

CA Common Discovery includes the following components:

### **Discovery User Interface**

Provides the services to support an administration user interface thin client. The Apache Tomcat web service must be previously installed on the computer. The installation prompts for the Tomcat installation path, discovery server name, and port number.

You can install multiple Discovery UI components that point to a single Common Discovery server component.

### **Discovery Server**

Provides a central point for storing and querying Discovery data, options, policies, and log data. The server includes a single instance of the discovery agent. You must install at least one discovery server in a CA Common Discovery installation. Large environments can install multiple discovery servers as they are needed, but it is not necessary to have multiple discovery servers. Several CA products can share the same discovery server. The installation prompts for the discovery server port number.

Installation includes the following discovery server subcomponents:

- Request Manager
- Database (DB) Manager
- Log Manager
- Enterprise Common Services Installation

**Note:** It is not necessary to have multiple Discovery Servers. You can share a single Discovery Server across multiple CA products.

### Discovery Agent

Provides discovery data gathering. CA Common Discovery installs at least one discovery agent with a discovery server in a CA Common Discovery installation. Multiple agents can be installed at strategic locations on the network to gather data and communicate it back to the discovery server. After installation, the agent service starts automatically. It also registers itself with the discovery server. The agent registration process adds the agent to the discovery server's agent list. The server returns a set of default agent options. The installation prompts you for the discovery server name.

Installation includes the following discovery agent subcomponents:

- Request Agent
- Discovery Engines
- Enterprise Common Services Installation

### Discovery Web Client

The web client is a thin client that depends on centrally installed UI Discovery Request Client Servlets. Web clients are used to administer Discovery policies and options.

**Note:** For any Common Discovery deployment, there will be any number of web clients.

### Discovery Request Client

Discovery Request Clients leverage exposed SOAP methods that provide remote access to the Common Discovery Server. The SOAP infrastructure ensures these methods are available to any C/C++ and Java applications. Consumer applications must integrate with Common Discovery by making calls to a C/C++ or Java Discovery Request Client as appropriate in order to retrieve the discovered entities. The consuming application may also choose to integrate with the Discovery Request Client in order to dynamically invoke a discovery.

The Common Discovery installation provides UI component that consists of Java Servlets that invoke the Java Discovery Request Client.

- Large, dispersed environments may install multiple Common Discovery UI components as per customer needs.
- Multiple UI components are not a requirement. In fact, a single UI component may be shared across multiple CA products.

**Note:** For any Common Discovery deployment, there must be at least two installed Discovery Request Clients - one that is specific to the consumer application and the other that consists of the Common Discovery Java Servlets for supporting its web clients.

## CA Common Discovery GUI

CA Common Discovery performs network discovery for your CA product. You can access the CA Common Discovery GUI using the following URL:

`http://localhost:port number/cmndiscovery/DiscoveryUI.html`

## Using Common Discovery GUI

You can use the CA Common Discovery GUI to perform the following administrative tasks:

- Configure server and agent options
- Define scan and purge policies
- Run Ad Hoc scan
- Run Ad Hoc purge
- Schedule scan and purge policy requests
- Check the scan progress
- Cancel a scan
- View scan history
- View discovery log information

## Configure New Discovery Server

You can use the CA Common Discovery GUI to administer multiple discovery servers.

### To configure a new discovery server

1. Open the CA Common Discovery GUI.

The default discovery server and its default namespace get connected.

**Note:** To select a different discovery server, you can select the server from the Discovery Server drop down list. You can also select a different namespace from the Namespace drop down list.

A *namespace* is a repository of information that includes discovery policies, entities, and scan history that are specific to a CA product. During discovery server installation, CA Common Discovery creates a common namespace that can be shared by multiple CA products.

2. Click the Edit button next to the Discovery Server drop-down list.

The Discovery Servers dialog appears.

3. Click New.

The Configure new Discovery Server page appears.

4. Complete the fields as appropriate, and click OK.

A new discovery server is configured.

**Note:** You can click Find to make sure your fields are accurate. This validates that a connection to the discovery server is possible. It also populates the default Namespace drop-down list based on the available namespaces defined on the discovery server.

### More information:

[Discovery Server Options](#) (see page 144)

[Set Discovery Server Options](#) (see page 143)

## Set Discovery Server Options

The discovery server options globally control the functionality, set exclusion criteria for all scans, and control SNMP community names list on the CA Common Discovery Server. You can configure your trusted servers list and subnet hierarchy using server options.

### To set discovery server options

1. Open the CA Common Discovery GUI.

The default discovery server gets connected.

**Note:** To select a different discovery server, you can select the server from the Discovery Server drop-down list and click Connect.

2. Select the Options tab.

The Server Options page appears.

3. Set the options as appropriate, and click Save.

The discovery server options are defined.

**Note:** You can click Reset to get the last saved server options.

### More information:

[Configure New Discovery Server](#) (see page 142)

[Discovery Server Options](#) (see page 144)

## Discovery Server Options

To set discovery server options, the following information is required:

### IPv6 Display Format

Specifies the IPv6 address display format.

### Scan history deletion frequency (in days)

Specifies the retention period of scan history.

### Log Destinations

(Optional) Defines the destination where you want to route the log data. You can route the log data to the following options:

- eTrust Audit agent
- Event Management node
- System event log

**Note:** If you route the log data to a remote machine, ensure that the remote machine has the Enterprise Common Services (ECS) installed and running.

### Trusted Servers

Configures a list of CA Common Discovery servers. You must configure the discovery server's options so that both the servers are listed in each others Trusted Servers list to share information.

### Subnet Hierarchy

Configures subnet hierarchy. The IPv6 protocol gives network administrators the ability to define subnet hierarchy in their IPv6 networks. The following parameters can be used to define a subnet hierarchy:

#### Global Routing Prefix Length

Determines the number of bits in IPv6 address that precede the subnetID.

#### Bits Per Level

Associates bits in the subnetID with subnet levels in an IPv6 subnet hierarchy.

**Note:** The bits are left justified. The first subnet level comprises the leftmost n bits of the subnetID. There is a corresponding filter in scan policies that lets you use these bits to filter scan requests.

Subnet hierarchy configuration is adjusted only if scan policies do not use it. The policies that leverage the current subnet hierarchy configuration are listed on the right hand side. You can use the Scan tab to review and edit the subnet hierarchy filters individually. You can click Disable Filters to remove the subnet hierarchy filters for all listed policies.

To set global scan options, the following information is required:



### **Global SNMP Community Names**

Controls the SNMP community names list globally for your discovery server. You can avoid adding community names when you define scan policy by adding enterprise-wide SNMP community names in the list. For better performance, keep the community names list small or the discovery agent could incur overhead when attempting to connect to SNMP enabled devices.

**Note:** Global exclusion criteria and SNMP community name values are appended to the scan policy at the time it is run.

### **Global Exclude IP Addresses**

Sets global exclusion criteria for all scans.

### **More information:**

[Set Discovery Server Options](#) (see page 143)

## **Set Discovery Agent Options**

The discovery agent options control discovery functionality and resource consumption on each agent system. Agent options can be configured globally or independently so that all agents can be configured similarly, individually on an agent by agent basis, or in a combination.

### **To set discovery agent options**

1. Open the CA Common Discovery GUI.

The default discovery server gets connected.

**Note:** To select a different discovery server, you can select the server from the Discovery Server drop-down list and click Connect.

2. Select the Options tab.

The Options page appears.

3. Select the Agent tab.

The Agent Options page appears.

4. Set the options as appropriate, and click Save.

The discovery agent options are defined, and the updated data is displayed in a tabular form.

**Note:** You can click Reset to get the last saved agent options.

## Discovery Agent Options

To configure discovery agent options, the following information is required:

### Default Agent Options

Sets the following default agent options:

#### Minimum Database update frequency (in seconds)

Specifies the frequency of updates from discovery agent scan processing to discovery server.

**Note:** The value 0 indicates that updates must be sent immediately as they are recognized, but this generates network traffic.

#### Maximum Classification Engines

Controls the number of instances of a classification engine scan running on an agent system at any given time.

#### Maximum Ping Sweep Engines

Controls the number of instances of a ping sweep engine scan running on an agent system at any given time.

#### Maximum Router Probe Engines

Controls the number of instances of a router probe engine scan running on an agent system at any given time.

### Use Default Settings

Sets the default agent options defined for the agent in the row. If the agent machine has different settings than other machines, clear the check box and manually set the values in the row as required.

**Note:** A new discovery agent will adapt the default settings once it registers with the connected discovery server.

### Remove Agent

Removes the discovery agent row from the discovery server's agent list.

**Note:** Under normal conditions, when a discovery agent is uninstalled, it notifies the discovery server and its corresponding row in the table is automatically removed. In case the discovery agent is uninstalled and the network connectivity with the discovery server is interrupted, you must select the Remove Agent check box to remove the agent.

### More information:

[Set Discovery Agent Options](#) (see page 145)

## Understanding IPv6 Discovery

The Common Discovery Import service retrieves IPv6 data from one or more Discovery servers and adds the information to the MDB. You identify those servers as Discovery servers using the Unicenter MCC IPv6 Import tool, and you also use that tool to configure the service.

The first time Common Discovery Import runs, it retrieves all IPv6 data from the Discovery server or servers. Subsequent imports retrieve only data that has changed. You can, however, configure the IPv6 Import tool to retrieve everything.

The WorldView Topology view shows IPv6 devices under IPv6 Network in the left pane. Some devices may be visible under both TCP/IP Network and IPv6 Network if they are dual-stack and IPv4 discovery was previously run.

The Common Discovery Import service continuously polls the Discovery servers every 40 minutes for new, updated, or deleted entities. You can change this collection interval with the IPv6 Import tool. You can also stop the service.

## Using IPv6 Discovery

The CA Common Discovery agents gather information about IPv6 devices in your network, and the Common Discovery Import service collects that information and updates the MDB. You can identify the Discovery servers, start and stop the import, and configure discovery policy.

**Note:** For more information about Common Discovery, see the Common Discovery Help. For more information about the Common Discovery Import service and the IPv6 Import tool, see the MCC Help.



# Chapter 4: Visualizing Your Enterprise

---

This section contains the following topics:

[WorldView Components](#) (see page 149)

[Configuring Business Process Objects Using Business Process Views](#) (see page 180)

[Creating Business Process Views Using SmartBPV](#) (see page 184)

## WorldView Components

WorldView provides an extensive set of tools that let you customize any view of your enterprise according to its logical structure, from 2D maps for iconic views, to a selection of browsers for textual views.

Using WorldView, you can perform the following tasks:

- See your entire network graphically represented on the 2D Map, grouped into networks, subnets, and segments based on their logical relationships.
- Define new classes using the Class Editor (Management Command Center) or Class Wizard (WorldView Classic), allowing for characterization and modeling of practically anything.
- Create customized Business Process Views (static and dynamic) of specific processes based on different business needs, resource features, roles, geographical locations, organizational structures, and applications.
- Import and export objects to and from WorldView using the Repository Import Export utility (trix).
- Set policies for a managed object's severity using alarmsets.
- View relationships among objects using the Association Browser.
- View links between objects using the Link Browser.
- Travel back in time to view historical information about your network using the Management Command Center Historian view.
- Determine the relative importance of an object in your network by associating a weight to each object.
- View MIBs data using ObjectView.
- Import IPv6 network devices discovered by Common Discovery using the Common Discovery Import service.

- Discover, view, and manage certain Unicenter components, such as the WorldView Severity Propagation Service, the Distributed State Machine (DSM), and Enterprise Management, without installing CA NSM on the computer where these components are installed.

## Managed Objects

Managed objects represent entities stored in the MDB that can be monitored and controlled using the Enterprise Management applications. These may include hardware, software applications, databases, and various business processes.

Managed objects have two distinct characteristics:

- They represent entities that have vendor-supplied agents or a product agent.
- They are derived from an existing class in the MDB that is itself derived from the ManagedObject class.

**Note:** Only objects instantiated from the ManagedObjects class with the class property `map_visible` set to True are visible in the 2D Map.

System managed objects are different from ordinary objects. You can use a managed object to monitor and control the IT infrastructure entity that it represents. A managed object can represent anything: hardware, software, business process view, and so forth. Managed objects have the following characteristics:

### Object Properties

Relate to the state of an object. An object usually has a set of possible states determined by its properties. A property, in turn, may be either an attribute or a relationship.

The terms *property* and *method* take on special meaning when used in the context of either class or instance. Therefore, we qualify the discussion when necessary by referring to class-level and instance-level properties and methods.

### Object Methods

Determine the kind of behavior exhibited by the object. For example, an object modeling a satellite might have a method that calculates the satellite's position for display on a monitor.

### Object Attribute

Specifies a type of property containing literal values. An object representing a PC would probably have a value for the number of printer ports available.

**Object Relationships**

Specifies a type of property denoting that a class or object relates in some way to another class or object. For example, a model object may have vendor information, which forms a relationship to a vendor object.

**Topology**

Presents a set of relationships among managed objects. Using a broad definition of topology simplifies the task of modeling object topology. Topology represents the set of relationships between objects. The simplicity of this definition allows for more flexible interpretation, wider functionality, and more powerful applications.

## Viewing Your Network Topology Using the 2D Map

The 2D Map gives you direct access to the managed objects that the MDB maintains. As you launch the 2D Map, these objects appear automatically in folders and are arranged according to your network topology (after running discovery). In addition to managed objects, the 2D Map can contain both WorldView objects and map-only objects.

Two concise views simplify the deployment and maintenance of your enterprise management system--iconic and textual. By default, the maps display your objects (represented as icons on the 2D Map), one level at a time. When you open or expand an icon, the objects included in this object appear.

The 2D Map lets you show objects, delete objects, change the values of their properties, enhance their definitions with comments and user-defined fields, and add objects to the MDB. You can search and query object information in the MDB. The status of your managed objects updates in real-time.

The 2D Map uses real geographic maps to determine physical locations; for example, it has knowledge of longitudes and latitudes. Thus, you can position objects or collections of objects for realistic placement on the various maps, which feature full geographic detail such as roads and bodies of water. You can create these using the Cartografx Vector Globe product that is provided with CA NSM.

Using the 2D Map, you can create map-only objects, such as background maps, lines, shapes, and text to supplement your networks, subnets, and segments, based on your network topology.

The 2D Map also addresses additional management needs by letting you create customized Business Process Views. You can add images, such as floor layouts, to your views to see exactly where a device resides or a problem occurs.

## 2D Map

The 2D Map is a two-dimensional, geographical representation of the logical structure of your enterprise. The map simplifies managing your eBusiness resources with tools you can use to customize visual representations of your business enterprise.

To access the 2D Map, select Topology from the list box of the left pane of the Unicenter MCC. Expand the Managed Object Root tree and select an object from the list. Select 2D Map from the list box on the right pane. The 2D Map appears in the right pane.

The 2D Map has a Toolbox that lets you create new objects, copy objects, move objects, delete objects, add links to other objects, and design custom views.

To view various aspects of your network Topology, select the ManagedObjects you would like to view from the Topology pane, then select 2D Map from the right pane of the Unicenter MCC.

To view various Business Processes, select the Business Process Views you would like to display from the Business Processes Views pane, then select 2D Map from the right pane of the Unicenter MCC.

Following are some of the map features that expand your 2D Map views and display properties of managed objects:

- To open an object for a view of its children, double-click the object. The 2D Map will fly into the object, expanding as it navigates.
- To display all the instance-level properties of an object, right-click the object, select Open Viewer from the context menu to open a submenu, and click Properties. The Property Notebook opens, displaying properties for the selected object.
- To display a cross-hair containing the name, label, IP address, and severity status, hold the cursor over the object.



## Billboards

Billboard objects let you keep up-to-date information about critical (not propagated) objects in view at all times. You can take a quick look at the billboard to see if any of the children of this container are critical. They are real objects so you can enter the billboard to solve the problem. Double-click on the object to get a closer look at a critical object.

Once you create a billboard object, all of the critical children of the billboard's siblings are shown in the billboard. If the critical status of an object changes to normal, that object is removed from the billboard. A status that changes from normal to critical causes the affected object to appear in the billboard automatically. To create a billboard, click the Toolbox icon and choose Billboard from the Managed Object tree, then drag and drop it into the 2D Map.

You can see the status of any object that appears in the 2D Map at a glance, because devices with problems, along with their segments, subnets, and networks, appear in colors reflecting the severity of their problems. Alarmsets defined in the MDB and represented on the 2D Map determine the relative importance of different classes of faults by assigning each one a severity status. CA NSM provides default alarmsets that you can assign to any object, customize, and extend.

**Note:** Do not place billboard objects on the topmost parent.

## Background Maps

You can add background images to your 2D Map by choosing a background image or geomap from the Toolbox tree list. Click and drag the image by holding down the left mouse button and dropping it onto the 2D Map. The background map appears underneath the objects in the 2D Map. The Unicenter MCC supports any BMP graphic you want to use as a background. Use the context menu to remove the background image.

**Note:** The object you select in the Unicenter MCC left pane determines the contents of the Toolbox images; that is, the images available to add as background. The Toolbox is always populated with the classes and images from the current provider of the object you select. For example, if you select a host name in the left panel and then navigate, or drill down, to the agent level, the Toolbox provides images and geomaps you can set as backgrounds. However, if you select an agent from an expanded tree view in the left pane, no class or images are available in the Toolbox because the DSM provider does not expose any images for use.

You can create additional custom maps to use as background geomaps with Vector Globe, a product licensed by CA from Cartografx Corporation. You can create maps for anywhere in the world with a configurable level of detail, place these maps as backgrounds for your topology, and arrange devices by geographical latitude and longitude. Vector Globe is provided as a separate CD for you to install after CA NSM is installed.

### Save Arrangement of Objects in 2D Map

In addition to the existing object arrangements, you can also have your own object arrangements in 2D Map.

#### **To save the arrangement of objects in 2DMAP**

1. Open 2D MAP.  
The 2D Map window opens
2. Drill down to the parent layer, that is, a layer above the existing layer.  
The parent layer opens.
3. Right click on the parent layer icon and click Open Details.  
The Managed Object Notebook window opens.
4. Click on the Others tab and navigate to the `autoarrange_type` property name.
5. Set the Property Value as desired. Click OK.  
The new arrangement is saved.

### Custom Views

Custom Views implement functionality that lets you display 2D Maps with custom rendering such as colors, link tariff information, link bends, and so forth. A Custom View allows for the MDI-style layout of multiple maps or plugins across multiple Unicenter MCC frames.

You create your custom objects in the 2D Map using the Toolbox after first activating Custom Views in the Unicenter MCC view toolbar drop-down menu option.

Custom Views provide the following features:

- Text boxes
- Multiple bitmaps
- Shapes, such as Circle, Diamond (Variable), Diamond (Fixed), Ellipse, Hexagon, Pentagon (Right), Rectangle, Rhombus, Square, Trapezoid, (Up), Trapezoid (Right), Triangle (Up), Triangle (Left), and None
- Lines
- Polygons
- Bendable links
- Layering

**Note:** To convert existing .gbf files into custom views, open the .gbf file in the WorldView Classic 2D Map and resave it. The .gbf is converted and appears in the Custom View left pane of the Unicenter MCC under the private node. The custom view is a named, publishable object that contains custom rendering and layout.

## Favorite Views

Favorite View allows you to create placeholders of specific objects for quick and easy view.

## How Navigation Works in the 2D Map

You can navigate in the 2D Map with a variety of actions. The following list describes actions you can take in the 2D Map:

- The pan feature lets you move the 2D Map to see where you are in the map without getting lost. To pan the 2D Map, click the mouse on a background area (not an icon or folder title bar) and drag it in the direction you want to pan.
- The zoom feature lets you zoom in and out of the map using your keyboard and mouse. To zoom in the 2D Map, hold down the Ctrl key and the left mouse button. Then, drag the mouse upwards to zoom in or drag the mouse downwards to zoom out.
- Click the right mouse button on the background area of an object to display a context-sensitive pop-up menu. Each menu is sensitive to the type of class from which the object derives; therefore, the menu that appears may vary for each class of object. The class information decides what pop-up menu appears.
- The pop-up menu is associated with the object's class and is customizable. For example, from a WBEM object's pop-up menu, you can ping to learn if WBEM is active on that machine and you can browse the WBEM data of the object. You can customize the pop-up menu with the CA NSM SDK, the Class Editor, or the Class Wizard.
- Positioning and holding the cursor over any object displays a cross-hair cursor containing up to four instance-level property values for the object such as name, label, IP address, and severity status. Class defines the cross-hair cursor content; therefore, the information that appears may vary for each class of object.

**Note:** You can customize the cross-hair cursor data and the pop-up menu for any class at any time with the Class Editor or Class Wizard.

- Double-clicking on any object executes the default pop-up menu option. This action usually zooms into the folder that displays the children of the selected object. That is, you begin to "drill down" through the levels of your IT infrastructure.

Continue to open objects to drill down to the lowest level and check the status, identity, and so forth, of each object.

- Select multiple objects with one of the following methods:
  - Hold the Shift button on the keyboard and click the first and last objects to select a range of objects.
  - Hold the Ctrl button on the keyboard and click each object you wish to have selected.
  - Hold down the Shift key, click and hold the left mouse button on a background area, then move the cursor to expand a selection rectangle that includes all the desired objects.

## Business Process Views

A business process view can be a manual effort, created by dragging and dropping, or it can be automatic, based on a query (a Dynamic Business Process View).

A CA NSM Business Process View is a logical group of managed objects that you create based on any criteria you determine, such as geographic location, business process, security, and so on. The Business Process View acts as a filter that displays only objects relevant to the management of resources for a specific business requirement.

A Business Process View helps you monitor and manage designated segments of your enterprise. The Business Process View is the means by which you can logically group the constituent resources that perform some business critical process; for example, accounting or personnel. The Business Process View is an effective way to alert you that a key link in a chain of resources is encountering some problem that may impact the business.

The contents of a Business Process View can represent whatever you decide is important to your enterprise. You can group these views by geographical locations, organizational structures and roles within the organization, applications, resource features, or any other requirements.

You can choose to display all Business Process Views at the top level or create intermediate ones to help organize them. For example, if you have Business Process Views for each city where your company has a branch, you could organize them into regions like Northeast, South, Midwest, and so on. Instead of one long list, you would see only the regional Business Process Views when the list is collapsed.

**Note:** See the User Options dialog, Business Process Views page to organize your Business Process Views.

The Business Process Views you create are visually represented as separate folders on the 2D Map. A Business Process View is itself a managed object, is stored in the MDB, and is accessible to any user of your 2D Maps. The same adding and modifying procedures that apply to all objects apply to the Business Process Views.

To monitor the condition and status of objects, you can set triggers and thresholds, and intercept messages generated by programs participating in the process. These views can assist you in the early detection and prevention of problems, and when a problem does occur, the Business Process View provides an immediate, graphical view of the source and severity of the problem.

## Types of Business Process Views

You can create various types of Business Process Views using a variety of methods.

- **Static Business Process Views** are typically manually defined by the user. They are simply a collection of objects that are related in some way, for example, by class, or perhaps by the business process or service they support.
- **Dynamic Business Process Views** are populated automatically based on policy defined for that Business Process View. Two main Dynamic Business Process Views tools, Dynamic BPV and the Dynamic Container Service, let you define policy that results in objects being added and removed automatically. For example, you can define policy to show all critical Windows servers, all unmanaged switches, or all devices in the IP range 172.16.10.100-200, and so forth.
- **SmartBPVs** are based on communication protocols. These are collections of devices that are communicating on a given port or protocol, for example an HTML Business Process View, an SQL Business Process View, or a Microsoft Exchange BPV, and so forth. Smart BPV uses packet-sniffing technology from Sonar to detect who is talking to whom using what protocol, and to dynamically update the Business Process Views accordingly. SmartBPV is available on Windows only.
- **Business Process View Management (BPVM)** is a component that lets you create Business Process objects to monitor and control your network. You can use Business Process objects to apply new rules to your system to determine how state propagates from existing WorldView objects using methods ranging from simple counters to complex correlation rules. BPVM lets you implement policies to make automated high-level decisions about key resources and set warning and critical events when problems are detected. BPVM is available on Windows only.

The following Business Process Views are created automatically:

- The **Domain** Business Process View contains Domain objects. Each Domain object represents the Agent Technologies DSM component.
- The **WBEM** Business Process View is created by the Discovery process and contains all of the WBEM objects (devices that provide WBEM data) found in your network.
- The **Deployed Objects** Business Process View contains the state of all CA NSM components that are installed in the same DIA zone.

## Dynamic Business Process Views

A Dynamic Business Process View is a Business Process View that is populated with objects that match the specific criteria defined in a query. The process saves you the effort of manually searching for objects to populate a Business Process View.

When a device is added or removed from a Dynamic Business Process View, you can send a customized message to the Event Management Console.

### To define Dynamic Business Process Views

1. In the Unicenter MCC Business Process Views navigation pane, right-click the BPV parent object and select New, Business Process View to create a new BPV. In the Properties view for this new BPV, specify the Dynamic BPV name.
2. Reclassify the the new Dynamic BPV object to the DynamicBPV class by right-clicking on the object to open the context menu and selecting Reclassify, Business View, DynamicBPV.
3. Right-click the Dynamic BPV object to open the context menu again and select Viewers, Dynamic BPV Editor.
4. Define the rules that define membership to the Dynamic BPV folder.

**Note:** If you select any other Dynamic BPV from name drop-down list and save, the current rules that are displayed on the editor are applied to the selected Dynamic BPV. Although you may not see the new rules when you open that dynamic BPV, however, the rules have been applied.

5. Click Save to save the view and populate the folder.

## Dynamic Containment Service (DCS)

The Unicenter Dynamic Containment Service (DCS) is an extension of the Business Process View concept. The Dynamic Containment Service (DCS) maintains the contents of any designated Dynamic Container object in the MDB according to a fully configurable rule-based inclusion policy.

The Dynamic Containment Service consists of three closely related components:

- The DCS Service is a Windows Service that automatically starts and stops the DCS Engine during a reboot. The DCS Service helps to ensure that the DCS Engine is running whenever the MDB is running and that your container objects are an accurate reflection of their associated inclusion policy.

The service defaults to an automatic startup type. It can also be started and stopped manually using the Window Service Control Manager or from the command line with the following commands:

```
Net Start CA-DCS
Net Stop CA-DCS
```

- The DCS Engine implements the policy defined for each Dynamic Container object in the MDB. The Engine detects property changes, and objects are added or removed as children of a Dynamic Container object as they conform to or no longer conform to the inclusion policy.
- The DCS Policy Wizard quickly configures the engine to maintain your designated Dynamic Container objects. You can access the DCS Policy Wizard from the Start, Programs, CA, Unicenter, NSM, WorldView group. Using the Policy Wizard, you can perform the following tasks:
  - Select the repository you want the engine to run against, and configure sign-on details so that it can run unassisted as a Windows service.
  - Specify the location and granularity of the log file that the engine generates.
  - Specify the Event Management node to which events are forwarded by the engine. You must provide a hostname, or an IP address; either an IPv4 or IPv6 address is acceptable.  
**Note:** If you enter a valid compressed IPv6 address, the address gets expanded to maximal extended form. But if you enter an invalid IPv6 address, an error message appears.
  - Configure the inclusion policy for any number of Dynamic Container objects you want to have dynamically maintained by the engine.

**Note:** Configuration changes do not take effect until you stop and restart the DCS Service.



## Determining the Relative Importance of an Object in Your Network

You can define the relative importance of an object in your network infrastructure by assigning a numeric value from 1-100 that specifies the *weight*. The higher the weight, the more value an object has in your network. Each class of ManagedObject has a default weight assigned to it, and when your network is discovered, each object that is derived from that class inherits the default weight for each class. For example, the default weight of a router is higher than the default weight of a workstation because, typically, a router has more significance in a network than an individual workstation.

You can change this weight value for each object to reflect the value of the object in your network using the Status page of the Properties viewer. You may want to assign a higher weight to a particular router because the router affects more critical processes. For example, you may want to increase the weight of the router that is connected to the computers that process your company's payroll. Assigning a weight to an object helps give you a better indicator of the overall health of your network infrastructure.

## Severity Levels

The Management Command Center displays the current, real-time status of any managed object. If a managed object experiences a problem, its visual appearance in the Management Command Center changes.

CA NSM uses the severity level value to change the state of an object. The severity value is a number from 0 to 9, indicating one of the ten predefined severities provided with WorldView. The severity level value of an object at any given time is assigned using policy that you define. The severity of a managed object determines its appearance in the Management Command Center.

Using the Status Color Schemes Editor, you can override the default colors that are used to show an object's severity. This feature lets you customize the appearance of Management Command Center objects and lets you better visualize your network.

## Weighted Severity

WorldView calculates the weighted severity of an object by multiplying an object's numeric severity by the object's weight.

**Note:** The weighted severity component of CA NSM uses only the following severity to status mappings:

0=Normal

1=Unknown

2=Warning

3=Minor

4=Major

5=Critical

### Example: Calculate Weighted Severity

The default weight assigned to the router class is 60. When you discover all of the routers in your network, each router inherits a weight of 60. However, you have one router (Router A) that is more valuable to your network, so you change the weight of that router to 80.

When an object below Router A changes state, for example, a server is low on disk space and goes critical, that critical severity is propagated to Router A. The propagated severity of Router A is then critical. Because you assigned a weight of 80 to the router, the propagated weighted severity is 400.

If an object below another router (for example, Router B, which has the inherited weight of 60) also changes to a critical state, the propagated weighted severity of Router B is 300. These values are used in the algorithm to derive importance.

## Object Importance

Using a sophisticated algorithm, WorldView determines the *importance* of each managed object in your network. Letting you view the importance of each managed object in your network helps you better analyze your network by giving you a better view of the health of your IT infrastructure. For example, importance lets you quickly and easily distinguish between a printer that has a critical severity because it is out of toner and a critical server that processes your payroll.

Importance is calculated using the weight and severity levels of child and parent objects in your network. The importance of an object increases when the propagated weighted severity of one of its child objects increases. You can set the weight of each object in your network, or you can use the default weight that is preset for each managed object class.

After WorldView calculates the importance of an object, the following thresholds determine what color is used to display the object in Maximum Importance view:

- Insignificant--0-15
- Of Interest--16-40
- Minor Concern--41-60
- Major Concern--61-80
- Severe--81-100
- Ultra--101-500

You can change the default thresholds by editing the `wvapiwrap.cfg` file. On Windows, this file is located in the `install_path\CA\SharedComponents\CCS\WVEM\CONFIG` directory. On UNIX/Linux, this file is located in the `$CASHCOMP/ccs/wv/config` directory.

These thresholds map to the severity levels for the purposes of displaying the colors that represent the different levels of importance in the Management Command Center. These same colors associated with the six levels of severity are used for importance:

- Ultra appears as the same color that is defined for a severity of Critical.
- Severe appears as the same color that is defined for a severity of Major.
- Major Concern appears as the same color that is defined for a severity of Minor.
- Minor Concern appears as the same color that is defined for a severity of Warning.
- Of Interest appears as the same color that is defined for a severity of Unknown.
- Insignificant appears as the same color that is defined for a severity of Normal.

### Change Default Importance Thresholds

WorldView determines the importance of each managed object in your network. After WorldView calculates the importance of an object, certain numeric thresholds determine what color is used to display the object. You can change the default importance thresholds according to your company's requirements.

To change the default importance thresholds, edit the `wvapiwrap.cfg`. On Windows, this file is located in the `install_path\CA\SharedComponents\CCS\WVEM\CONFIG` directory. On UNIX/Linux, this file is located in the `$CASHCOMP/ccs/wv/config` directory.

## Set Policies for a Managed Object's Severity Using Alarmsets

Alarmsets associate state changes of an object with a severity level. Each alarmset has many alarmset entries that map each status text to one of 10 predefined severity levels. The severity level then determines how the object appears on the 2D Map. For each alarmset entry, you can also specify whether the severity is propagated up through the object's parent folders.

CA NSM assigns a default alarmset for each class. You can add new alarmset entries to this alarmset or you can assign a new alarmset to any object in your network.

**Note:** Alarmsets are associated at the object level, not the class level. Every object instantiated from a subclass of the ManagedObject class can be associated with an existing alarmset. This means that each managed object can be associated with a separate alarmset.

For a CA NSM object to be a managed object, it must be associated with a device MIB, an agent MIB, or some application that can report to the MDB (through SNMP traps) on the state (status) of the object. Whenever the status of a managed object changes, CA NSM compares that status setting to the alarmset entries (policies) defined by the alarmset associated with the managed object.

**Note:** If the object does not have an alarmset associated with it, the severity of the object does not change on the 2D Map. Also, if the object has an associated alarmset and if there is no entry that matches the status received, the severity of the object does not change. However, in addition to updating or setting the status of the object, the application may choose to explicitly update the severity. The severity set by the application overrides the severity defined by the alarmset entries.

## Severity Propagation Service

The Severity Propagation Service is a critical service that is installed as part of the WorldView Manager component. It keeps the WorldView status propagation up-to-date and also other statistics, such as importance, nodal sum, and variance, which show the state of WorldView objects in CA NSM user interfaces. The Severity Propagation Service must run on the MDB server.

The Severity Propagation Service is actually a DCOM server that measures the readiness of the Severity Propagation Engine, which is started by the service. You cannot see WorldView objects in the Management Command Center until the Severity Propagation Engine and the Severity Propagation Service are ready.

When CA NSM is installed, the Severity Propagation Service is registered and a SeverityPropagation user account with a strong password is automatically created for the Severity Propagation Engine. A RunAs user account with the same password is also added to the dcomcfg utility.

These user IDs are created so that the Severity Propagation Engine can stay connected when the user logs off of the computer. You may want to change the password for this user for security reasons. To do this, you must deregister the Severity Propagation Engine, which removes the user accounts, and re-register it with a new password.

**Important!** All WorldView connections to this MDB must be closed before stopping the Severity Propagation Service, which is a prerequisite for deregistering and re-registering the Severity Propagation Engine, or the severity for this WorldView repository will be incorrect.

**Note:** For information about changing the password for the Severity Propagation Engine user accounts, see [Change the Password for the Severity Propagation Engine User Accounts \(Windows\)](#)

**Important!** When attempting to start or restart CA NSM services manually on a computer that contains the WorldView Manager, you must start the Severity Propagation Service first. This action ensures that all services that use the Severity Propagation Engine initialize correctly. In particular, you must start the Severity Propagation Service before you start Agent Technology Services and DIA/DNA Services. Also, when you shut down the MDB on a computer running a WorldView Manager, you must recycle current persistent services and any instances of Unicenter MCC (both local and remote) after you restart the MDB. Current persistent services include, but are not limited to, the Severity Propagation Service and Agent Technology Services.

## How You Correctly Stop and Restart the Microsoft SQL Server Database

You may need to stop and restart the Microsoft SQL Server database. To do this safely, you must follow a specific sequence. Following this sequence ensures that you can see all objects in your network, and that the Severity Propagation Service is correctly reporting the severity of all objects in your network.

1. Run the following command to stop the RMI Server:  

```
RMI_MONITOR -k LOCALHOST
```
2. Stop the following services using the Windows Service Control Manager:
  - CA Agent Technology services
  - CA-Continuous Discovery Manager
  - CA DIA DNA (there may be a version associated with this service)
3. Stop the Microsoft SQL Server database.

4. The Microsoft SQL Server database starts by itself at the next request.
5. After Microsoft SQL Server starts, restart the following services in this sequence:
  - a. CA WorldView Severity Propagation Service
  - b. CA DIA DNA (there may be a version associated with this service)
  - c. CA Agent Technology Services
  - d. CA-Continuous Discovery Manager

**Note:** When you stop the Microsoft SQL Server database on a computer running a WorldView Manager, you must also stop and restart any instances of Unicenter MCC (both local and remote) after you restart Microsoft SQL Server.

## Viewing Object Details and Properties

Every managed object is derived from the ManagedObject class, or a subclass of ManagedObject. You can review, and sometimes modify, the values of a managed object's instance-level properties with the Properties viewer.

Each of the tabs in the notebook corresponds to a property group for the class. Each field represents an instance-level property. The number of tabs that appear depends on the number of property groups defined in the object's class definition.

**Note:** Use the Class Editor facility to modify the occurrence of class- or instance-level properties.

## Modifying Class Properties with the Class Editor

CA NSM comes with an extensive set of predefined classes to manage your IT resources. (To learn more about the predefined classes, see the Predefined System Classes section of the *Programming Guide*.)

There will, on occasion, be times when you want to modify the properties of an existing class; for example, to change the context menu associated with a class of objects or the icon that appears.

In addition, there may be instances when an application or device contains unique qualities that warrant the creation of a new class, defined as a subclass of one of these predefined classes (primarily the ManagedObject class).

Making updates of this nature is done using the WorldView APIs of the CA NSM SDK or through use of the Class Editor.

## Reviewing Class Definitions

Before modifying the CA NSM class hierarchy, you should first review the predefined classes by selecting the Class Specification category in the left pane of the MCC and the Class Editor from the right pane's drop-down list.

Class Specification and the Class Editor display the class definitions of the MDB in an organized hierarchical format. It is an effective tool that lets you see what properties exist for each class.

The left pane displays the CA NSM classes in a hierarchical structure. The right pane displays the class-level properties of the class highlighted in the tree, followed by instance-level properties when they exist for the class.

Class properties define an object's state and behavior and can be of two types, class or instance. Class-level properties may not be modified. Instance-level properties are unique for each object and may be modified.

**Note:** For a complete list and description of all existing properties and their attributes, see the class- and instance-level properties tables in Object-Oriented Programming Basics in the MCC online help.

## Viewing MIBs and WBEM Data with ObjectView

ObjectView lets you browse a device's MIB and WBEM data. The MIB consists of a collection of attributes that are retrieved using network management protocol. MIB information includes aspects of a device's performance including the number of devices connected, which interfaces are down, and the number of packets coming in compared with the number going out.

WBEM information is a standardized collection of end-to-end management and diagnostic data in enterprise networks that can include hardware, protocols, operating systems, and distributed applications.

**Note:** Use the WBEM Browser to browse WBEM data.

The left pane of the ObjectView container displays the groups of attributes belonging to a device. Expanding a category shows the particular attributes for which you can obtain values from the device's MIB or WBEM data to display in the right pane.

You can access ObjectView directly from Tools in the left pane drop-down of the Management Command Center.

Using ObjectView, you can obtain attribute information such as Object ID, Value Type, or information for an attribute group such as an attribute count and, if applicable, table information. You may also set attribute values at the device level. For example, if an interface is having problems, you can change its adminStatus from up to down.

ObjectView also provides the DashBoard Monitor to let you graph selected MIB attributes in real time.

### DashBoard Monitor

The DashBoard Monitor lets you graph selected MIB attributes in real time. It provides access to the Graph Wizard, a tool for creating dynamic graphs of device performance for further analysis. WorldView Classic's ObjectView supports Microsoft Excel and you can display the collected data in Excel spreadsheets.

Using the Dashboard Monitor and the Unicenter MCC Customize Chart, or the WorldView Classic Graph Wizard, you can create formulas for attributes you select for graphing that include usage of polling interval information.

### Customize Chart Overview

You can design ObjectView charts using Unicenter MCC. The Customize Chart dialog provides a tabbed notebook containing several pages of graph options. You can select numeric attribute values in ObjectView that are then added to the DashBoard. Click the Graph button on the DashBoard toolbar to open the Customize Chart dialog.

### Graph Wizard Overview

The Graph Wizard takes you through the steps necessary to create a dynamic graph for ObjectView. You can access the Graph Wizard from the DashBoard.

You select numeric attribute values in ObjectView, which are then added to the DashBoard. The Graph Wizard lets you customize the graph by setting options in the following windows:

- Greetings
- Style
- History
- Foreground Color
- Background Color
- Attribute Colors



In addition to these options, you can set alarms in graphs, except LED, to change the display of an attribute value when it reaches a definable threshold. You can customize alarm notification by setting colors, text, and severity.

## Viewing Relationships Among Objects Using the Association Browser

You can navigate different types of relationships among objects using the Association Browser. For the selected object, if there are associations to other objects, you see these relationships in the form of a hyperbolic tree. The hyperbolic tree helps you see what is happening in your enterprise as you traverse nodes through associations. Objects are linked by color-coded lines that represent the associations.

**Note:** Not all classes of objects displayed in the Management Command Center support the concept of associations. Also, many objects of classes that do support associations may not have any associations currently defined. When this is the case, only the selected object appears in the Association Browser.

### Context Menu

The Association Browser displays a context menu when you right-click an object. Each menu is sensitive to the type of class from which the selected object is created. The class information decides what context menu appears.

When you right-click the whitespace (background) in the Association Browser, a context menu with the following options appears:

#### **Expand**

Expands the root node object to reveal any associations. If there are associations to other objects you will see these relationships in the form of the hyperbolic tree.

#### **Collapse**

Collapses the hyperbolic tree.

#### **Show Implied Links**

Toggles the display of implied links on and off. By default, implied links are displayed in the Association Browser.

Implied links are links that appear for parent objects because somewhere deep in the hierarchy, two or more child objects of those objects are linked together. Certain SAN objects have many implied links.

### **Show Reverse Inclusions**

Toggles the display of reverse inclusions on and off. Reverse inclusions show you the parent objects associated with each node in the Association Browser. Displaying reverse inclusions is useful when administering SAN objects.

By default, reverse inclusions do not appear in the Association Browser.

### **Legend**

Displays the Association Legend.

You can also right-click an association link object to display a context menu, which varies depending upon the type of association link connecting the two objects.

## **Viewing Links Between Objects**

There are two types of links displayed in the maps: true links and implied links. True links are links between two managed objects. They may be across subnets and may actually link across whole networks. From these true links, you can calculate implied links up the topology. Implied links show that a parent's objects are linked due to its children being linked. For example, a node on a segment in California is linked to a node on a segment in New York. An implied link is generated showing that California and New York are linked. There can be many implied links between two parents. The Link Browser manages all the links.

The Link Browser lets you view all links defined among objects. The dialog shows the links in a simple list of three columns: the source object, the link object, and the destination object. These links show the true links and their source and destination objects. The first in the list are normally the true links between the selected objects, followed by the implied links of the true links.

When you invoke the Link Browser, the link you select determines the source and destination objects. All links between these two objects appear in the Link Browser.

### **Open the Link Browser**

To view the Link Browser, right-click a link in the 2D Map, Association Browser, or Instances pane and select the Browse Links context menu item.

Using the Link Browser, you can take the following actions:

- Double-click a link in the Link Browser to open the Properties viewer for the link.
- Double-click a node in the Link Browser to navigate to that node in the original pane.
- Right-click a node or link to display a context menu for that object.

## Viewing Historical Information about Your Network

You can travel back in time to view historical information about your network using the Unicenter Controller and Management Command Center Historian view. Controls in the Unicenter Controller let you set the clock to a time in the past, and the Historian displays information for the specified point in time on timelines. The Unicenter Controller opens automatically only when historical event information is available for the selected object.

The Unicenter Controller provides controls similar to those found on a VCR. There are buttons for forward play and reverse play and others for quickly moving to key points in time. A dial lets you manually turn the clock back to a desired time in the past.

For example, you need to know what the state of a Printer 6B was during the weekend. Display the Historian for Printer 6B. The Unicenter Controller automatically opens. From the Unicenter Controller, rotate the dial counter clockwise to reach that date and time. When you release the mouse button, the historical data for the printer is updated in the Historian view for Printer 6B.

The Unicenter Controller also provides the following features to assist you in managing the present time:

- You can display a list of objects that are identified as being in a "critical" state. Select these objects from the Management Command Center to be updated to present information.
- You can display specific Business Process Views that you define.
- You can display views you have previously saved as bookmarks.

## Importing and Exporting Objects to and from WorldView

**Important!** Before using the Repository Import Export utility (trix), be sure that you are thoroughly familiar and have experience with the CA SDK. Lack of experience with the CA SDK may cause damage to WorldView and objects.

You can use the Repository Import Export utility (trix) to import and export objects to and from WorldView. Trix is an effort- and time-saving utility that lets you do all of the following:

- Back up the contents of WorldView.
- Transport a copy of the contents of WorldView to another site.

For example, in a larger company, you may want to populate the main corporate WorldView with the contents of WorldViews from offices around the country or the world.

- Populate WorldView with existing objects from another WorldView.
- Populate WorldView with objects that are similar to existing objects already stored in that WorldView or a different one.

You may want add a few objects that are similar to existing objects and avoid running Discovery, which saves network resources. For example, if you have two subnets with 10 computers on each that are the same except for their names, you can use trix to export the definitions of the existing computers into a trix file, modify the name property as necessary for each applicable object and import the objects with modified name properties into WorldView.

Trix exports (copies) the object definitions (properties) from WorldView into a trix script file. You can import or export all of WorldView or any part of it.

Trix is available using the WorldView Classic or Unicenter Browser Interface on Windows platforms. You can also use the trix command from a command prompt on Windows platforms.

## Export Methods

You can export objects from WorldView in either of the following ways:

- By object
- By class

When exporting by object, you can export child objects, inclusions, and outgoing and internal links. Incoming links are not available for exporting. When exporting by class, all objects that match the class definition you specify are exported.

You can then import (copy) these object definitions in the trix script into another logical repository. You can import the trix script as is or you can modify the object definitions in the trix script and import the modified object definitions into the same logical repository or a different one. You may want to do this when you have new objects on your network that are similar to existing objects. Performing this procedure makes it unnecessary to run Discovery for a small number of new objects, thus saving network resources. You can also create your own trix scripts using the trix script syntax.

## Trix Script Files

The default trix script filename on Windows is trix.imp. The default trix script filename on UNIX and Linux is TRIX.TNG. You can assign a more meaningful prefix to the files using the Export Repository dialog.

The name is limited to four characters. After the trix script file is created a number is assigned to it starting with zero (trix0.imp on Windows, TRIX0000.TNG on UNIX and Linux).

On Windows, trix files are limited in size to 1 MB. When the trix file (trix0.imp) size approaches the limit, a second trix file (for example, trix1.imp) is created until all the object definitions have been copied into the trix files. When multiple trix files are created, an Include file lists all the created trix files. When you start trix from the command line, you can use this Include file to reference all the trix files instead of listing each trix file individually.

## How to Export and Import Objects Using WorldView Classic

Open the Repository Import Export utility using WorldView Classic by choosing Start, Programs, CA, Unicenter, NSM, WorldView, Repository Import Export. After you select the repository from which you want to export or import, the CA NSM Repository Import Export window appears.

- To export objects, choose Actions, Export Repository. After you choose the logical repository from which you want to export objects, the Export Repository dialog appears. The Export Repository dialog lets you export object definitions by class or object and determine the location of the trix script file and the prefix you will use for the trix script filename.
- To import objects, choose Actions, Import Repository. The Import Repository dialog lets you start and stop the importing of object definitions into a logical repository, add and remove script files to be imported, and monitor the progress of the importing process. Select the trix files that contain the object definitions you want to import into your logical repository. Use the Add Scripts button to add a trix file not listed in the Import Repository dialog.

## How to Export and Import Objects Using Unicenter Browser

You open the Repository Import Export utility using the Unicenter Browser Interface by clicking the Tools button on the toolbar.

- To export objects, right-click the Repository Import/Export icon in the left pane, and select Export from the pop-up menu. The Export Repository dialog lets you export object definitions by class or object and determine the location of the trix script file and the prefix you want to use for the trix script filename.
- To import objects, right-click the Repository Import/Export icon in the left pane, and select Import from the pop-up menu. The Import Repository dialog lets you start and stop the importing of object definitions into a logical repository, add and remove script files to be imported, and monitor the progress of the importing process.

## Understanding IPv6 Discovery

The Common Discovery Import service retrieves IPv6 data from one or more Discovery servers and adds the information to the MDB. You identify those servers as Discovery servers using the Unicenter MCC IPv6 Import tool, and you also use that tool to configure the service.

The first time Common Discovery Import runs, it retrieves all IPv6 data from the Discovery server or servers. Subsequent imports retrieve only data that has changed. You can, however, configure the IPv6 Import tool to retrieve everything.

The WorldView Topology view shows IPv6 devices under IPv6 Network in the left pane. Some devices may be visible under both TCP/IP Network and IPv6 Network if they are dual-stack and IPv4 discovery was previously run.

The Common Discovery Import service continuously polls the Discovery servers every 40 minutes for new, updated, or deleted entities. You can change this collection interval with the IPv6 Import tool. You can also stop the service.

## IPv6 Import Tool

CA Common Discovery discovers IPv6 devices on your network. After discovery, you import the discovered information into the MDB using the WorldView Common Discovery Import service. The IPv6 Import tool is used to configure the import service.

This tool lets you perform the following actions:

- Identify the Discovery servers that have the IPv6 device information that you want to import
- Access the web-based CA Common Discovery configuration tool to define or edit the scan policy for a Discovery server
- Start and stop the import service
- Configure the import service's collection (poll) interval
- Control whether the service imports all discovered data or only changed information
- View the number of hosts, routers, and other devices added to the MDB during an import

IPv6 Import is one of the Unicenter MCC tools that you can access by selecting Tools from the drop-down list above the left pane.

This tool has the following fields and buttons:

### Discovery Servers area

Lets you define one or more Discovery servers that provide information about IPv6 devices, include or exclude the defined servers from an import, and configure what those servers should discover.

The columns are:

#### Server

Lists defined servers.

When you double-click a server name, a connection settings dialog opens so that you can change the user name or password used to access the web-based Common Discovery configuration interface.

#### Status

Shows whether a server is accessed during an import. The status is *Included* or *Excluded*.

**Default:** Included

### **Configuration**

Contains hyperlinks that open the CA Common Discovery configuration tool. This tool lets you configure, start, and maintain Discovery scans and purges. It also provides access to administrative information like scan history and logs.

The buttons are:

#### **Add**

Opens a dialog that lets you define a discovery server by entering the server name, protocol (http or https), port (8081 by default), and user name and password for accessing the server.

#### **Delete**

Removes the server from the list in the IPv6 Import tool. This server can no longer participate in the import.

#### **Exclude**

Prevents the server from participating in the import. The status changes to *Excluded* and the icon to the left of the server name becomes red.

#### **Include**

Lets an excluded server participate again in the import. The status changes to *Included* and the icon to the left of the server name becomes green.

#### **Select All**

Selects all discovery servers in the list.

### **Service Statistics area**

Shows how many hosts, routers, and other devices are added to the MDB as managed objects during the last data collection by the Common Discovery Import service.

### **Service Configuration area**

Lets you start or stop an import, specify the polling interval, and indicate whether the service imports all discovery data or only changed information.

#### **Start Service/Stop Service**

Lets you start or stop an import. The label is Start Service when no import is taking place, and Stop Service when the Common Discovery Import service is running.



**Collection Interval**

Opens a dialog that lets you specify the number of minutes the import service should wait between each poll to discovery servers.

**Default:** 60 minutes

**Reset Import**

Controls whether the import service requests all discovered objects from Discovery servers or only objects that have been added or updated since the last collection. A dialog asks you to confirm that you want to import all objects in the discovery server database. A full import can be time- and resource-intensive.

**Default:** new information only

**Note:** For more information about Common Discovery, see the Common Discovery Help. For more information about the Common Discovery Import service and the IPv6 Import tool, see the MCC Help.

## Registering and Updating Unicenter Components Using Unicenter Registration Services

The Unicenter Registration Services shows the CA NSM deployment and its health and availability status. This service lets any component (or product) register itself, create objects representing its desired hierarchy, and update them with health status. These tasks are accomplished without installing a WorldView component or having knowledge of a WorldView component on the registering device; only an active DIA is required.

The first WorldView Manager/MDB that is installed in a particular DIA zone is designated as the Unicenter Registration Service Server. You can change the server that is registered with DIA by using the `switchwvregsrvr` utility. For more information about the `switchwvregsrvr` utility, see the online *CA Reference*.

Using the Unicenter Registration Services, all CA NSM state objects are created in one MDB that has been designated as the Unicenter Registration Services server. This configuration provides a complete holistic view of the CA NSM deployment and its status. You see the whole deployment of CA NSM in one repository.

**Note:** When this mode is enabled, you may not see your Enterprise Management objects in the local repositories if the Unicenter Registration Services server is pointing to another MDB. If you have custom processing needs that are dependent on previous product architecture, you can revert to using WorldView daemon processing and the Unispace folder. For more information, contact Technical Support.

Using this functionality, the device where the component resides is discovered as well as the component. The Unicenter Registration Service Server contains a Business Process View under ManagedObjectRoot called Deployed Objects, which contains the state of all CA NSM components that are installed in the same DIA zone. After the component is registered in WorldView, you can use all of the WorldView functionality to view and manage the component, such as defining status propagation rules and find algorithms, and accessing the component using the Management Command Center, WorldView Classic GUI, and the Unicenter Browser Interface.

When you install a CA NSM component, the component is "registered" in the MDB on the Unicenter Registration Service Server, and a proxy object is created, which represents the component. Each component sets the state of its proxy object and this state is recorded in the MDB. Components that create proxy objects are WorldView, Enterprise Management, and Agent Technology. Agent Technology groups the DSM object into a Business Process View called Domain, which also appears directly under ManagedObjectRoot.

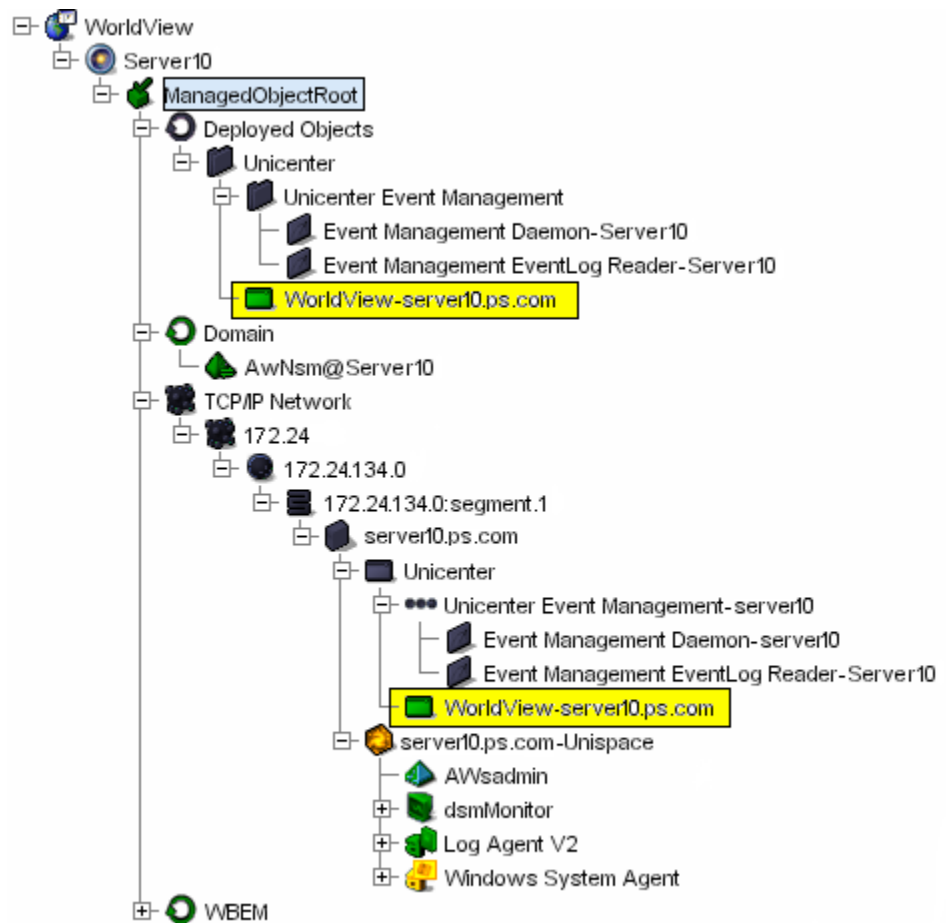
Possible states for each component include the following:

- Running
- Started
- Stopped
- Removed

You may want to use alarmsets to map the specific status text to a severity. The status texts are product-specific and help you diagnose the problem with a particular proxy object.

In addition to the state, other properties that are important to the health of the component are recorded in the MDB. When the state of an object changes, the change is reported to the MDB, and any properties that may have changed are also reported.

The following picture shows the Topology View for a computer called Server10, which is the designated Unicenter Registration Services server. Server10 also has WorldView installed on it. The Unicenter object is created under Server10 in the TCP/IP network. The proxy object, called, WorldView-server10.ps.cpm, is created as a child object of the Unicenter object. The Deployed Objects Business Process View mirrors this topology. Any status changes to the WorldView object will be reflected in both places.



## Configuring Business Process Objects Using Business Process Views

Business Process View Management (BPVM) is a component of CA NSM that lets you use Business Process objects to monitor and control your network. You can use Business Process objects to apply new rules to your system to determine how state is propagated from existing WorldView objects using methods that include simple counters and complex correlation rules. BPVM lets you implement policies to make automated high-level decisions about key resources and set warning and critical events when problems are detected.

For example, you may have a Business Process View™ that includes three web servers. Failure of one web server may be significant, but it may not have an adverse impact on your business operations. Failure of two web servers may have a major business impact, possibly bringing operations to a halt if the two servers remain offline for an extended time period. Using BPVM, you can define a rule that represents one server's failure as a warning state, and two servers' failure as a critical state requiring immediate action.

### Business Process Objects

Business Process objects represent a combination of other objects. Each Business Process object is associated with a target object, like an existing Business Process View. Business Process objects derive their state by applying rules to the associated target. Business Process objects can have any number of rules associated with them.

You can use the Business Process View Management Configuration Wizard to configure Business Process objects.

### Rules

Rules determine the Business Process object state. You can configure any number of rules for each Business Process, and you can combine rules to provide sophisticated forms of correlation for the contents (children) of the target object. Each rule is represented by an object in the MDB, under the associated Business Process, and thus independently influences the propagated process state. Rules report all threshold breaches to a designated Event Console to support automated actions.

BPVM provides several kinds of rules to influence the Business Process' state from different aspects of the target object and its children as follows:

- Child Count Rule
- State Count Rule
- Propagation Thresholds Rule
- Boolean Logic Rule
- Child Update Rule

### Child Count Rule

The Child Count Rule lets you set warning and critical thresholds on the number of children present under the target object. You can use this rule when the target object's contents are dynamic, either through manual procedures, scripting, or use of the Class Business Process View, Dynamic Business Process View, or Dynamic Containment Service utilities.

For example, if a dynamic container was configured with a policy to include business-critical servers showing a status of critical, you can use the Child Count Rule to set warning and critical states when an unacceptable number of these servers are included in the container. You could set a warning state when two servers are included in the container and a critical state when four servers are included in the container.

### State Count Rule

The State Count Rule lets you set warning and critical thresholds on the number of child objects at, or greater than, a given state under the target object. You can express thresholds as absolute values (numbers) or percentages.

For example, if a Business Process View was configured for key servers, you can use the State Count Rule to set warning and critical states when an unacceptable number or percentage of key servers is at, or greater than, a given state. You could set a warning state when 25 percent of the servers reach a critical state and a critical state when 50 percent of the servers reach a critical state.

### Propagation Thresholds Rule

The Propagation Thresholds Rule lets you control the propagation of certain states by placing thresholds on those states. You can express thresholds as absolute values (numbers), percentages (of child objects with a given severity), or elapsed time (time an object must be in a severity state before it is eligible for propagation).

For example, if a Business Process View often has one or two objects showing a down state, this condition could lead to the Business Process View permanently showing a propagated down state. This could effectively mask any key warning and critical states that would affect the management of devices in the Business Process View. You can use the Propagation Thresholds Rule to place a threshold on down state, which ensures that the Business Process object is only propagated when a significant number of devices are down.

### Boolean Logic Rule

The Boolean Logic Rule lets you correlate events for individual objects under the target object, and derive conditions from those events. You can use this rule to build lists of conditions and set warning and critical thresholds.

Conditions are defined in terms of an object, and (optionally) a state that must exist as a child of the target. With a dynamic Business Process View, you can have conditions to detect when certain key resources are included as children. With a static Business Process View, you can have conditions to detect when one or two essential resources are at a given state.

For example, if you configured a Business Process View to monitor key servers, you can use the Boolean Logic Rule to set warning and critical states when an unacceptable number of servers are at, or greater than, a given state. You could set a warning when Server A or Server B reaches a critical state and a critical state if Server A and Server B reach a critical state. You can use several instances of the rule type to effectively manage sub-elements of the overall Business Process and raise alerts when problems are anticipated or detected.

### Child Update Rule

The Child Update Rule lets you set warning and critical thresholds on the interval between updates to target object's children. You can use this rule to detect unanticipated or uncharacteristic periods of inactivity that can indicate a problem with the associated manager.

For example, if the objects in a central database are supposed to be updated every 30 minutes in batch, you can use the Child Update Rule to set warning and critical states if these objects are not updated in a timely manner. This could indicate that a network problem is preventing the objects from being updated. To detect this problem, you could set a warning state of 30 minutes and a critical state of 60 minutes.

## Integration with Event Management

Business Process and Managed object state changes (such as Start and Stop) and error events can be forwarded to the Event Console, if required. You can configure these events when you use BPVM to configure business processes. We recommend a notification level that will send key operating events to the Event Console.

### Notification Events

Notification events indicate that a threshold has been breached for a rule object and provide details about why the breach occurred. Notification events are reported to the Event Console. Examples of notification events are shown following.

When the target object changes to a critical condition, as in the Child Count Rule, the event provides details about the reasons for the state change as follows.

```
%UNI_BPVM, Target Object [ Europe Mail BusinessView ]: Rule ' ChildCount[W=3/C=4] '
for Business Process ' Business Process for Europe Mail ' is changing state to Critical
because the child count of 9 exceeds the Threshold of 4
```

When multiple conditions exist, as in the Boolean Logic Rule, the event provides details about which conditions were met as follows:

```
%UNI_BPVM, Target Object [ Europe Mail BusinessView ]: Rule ' Email Server Rule ' for
Business Process ' Business Process for Europe Mail ' is changing state to Critical
because { [ EWB-NTS-01.schaf.com WindowsNT_Server >= Warning ] [ EWB-NTS-03.schaf.com
WindowsNT_Server >= Warning ] }
```

When the target object returns to an acceptable condition, the reset event provides details about the state change as follows:

```
%UNI_BPVM, Target Object [ Europe Mail BusinessView ]: Rule ' Email Server Rule ' for
Business Process ' Business Process for Europe Mail ' is changing state to Normal
because no Thresholds have been breached
```

## Impact Events

Impact events signal a state change for a target object and provide details about the child object that caused the change, including whether the overall change is positive or negative. When event logging is set to Notification or above, impact events are reported to the Event Console. Examples of impact events are shown following.

```
%UNI_BPVM, Target Object [ Web Servers BusinessView ] for Proxy 'Proxy for Web Servers' is currently Normal
```

```
%UNI_BPVM, Target Object [ Web Servers BusinessView ] for Proxy 'Proxy for Web Servers' has WORSENEDED from Normal to Critical because child object [ TestA Windows2000_Server ] WORSENEDED from Normal to Critical
```

```
%UNI_BPVM, Target Object [ Web Servers BusinessView ] for Proxy 'Proxy for Web Servers' has IMPROVED from Critical to Warning because child object [ TestA Windows2000_Server ] IMPROVED from Critical to Warning
```

## Creating Business Process Views Using SmartBPV

Smart Business Process View Management (SmartBPV) is a component of Unicenter Network and Systems Management (CA NSM) that lets you automatically create and dynamically update Business Process Views. Through analysis of network activity, SmartBPV identifies infrastructure elements that support a specific application and automatically builds and continuously updates a focused view for management.

Sonar technology enables the business impact intelligence for SmartBPV. This technology goes beyond traditional discovery and inventory methods to dynamically map business applications to supporting IT resources.

**Note:** You should understand CA NSM Business Process Views before attempting to use SmartBPV.



## Business Process Views

A Business Process View is a simple, concise view of a logical grouping of managed objects related to a specific process. The contents of a Business Process View represent what you decide is important. You can group these views by geographical locations, organizational structures, and roles within the organization, applications, resource features, or any other requirements.

The Business Process View lets you effectively manage your infrastructure by:

- Alerting you that a key link in a chain of resources is encountering a problem that may impact the business.
- Letting you set triggers and thresholds to monitor the condition and status of objects.
- Assisting you in early detection and prevention of problems.
- Providing an immediate graphical view of the source and severity of the problem.

## Benefits of SmartBPV

SmartBPV benefits your enterprise by:

- Automatically building views that focus only on related elements within the infrastructure such as those for a specific application, business process, server type, or location.
- Showing which infrastructure components communicate directly with each other for faster problem resolution.
- Minimizing application downtime, freeing valuable IT resources to focus on new initiatives and strategic planning.
- Easily creating an automatically and dynamically updated management view for your domain of responsibility.

## How SmartBPV Works

SmartBPV provides the necessary data for Business Process Views as follows:

- SmartBPV probe collects data.
- Filters restrict the data and objects that SmartBPV adds to the MDB.
- SmartBPV dynamically updates the appropriate objects in the MDB with the collected data.
- SmartBPV analyzes network data to build and dynamically update management views (Business Process Views) that contain the infrastructure elements you identify as essential for monitoring.

## SmartBPV Examples

Examples of how you can use SmartBPV to benefit your enterprise include:

- Validating your infrastructure and determining where to place new software or apply maintenance (for example, how many Exchange monitoring agents are required and where).
- Collecting all instances of Windows servers and determining where activity is occurring between them for diagnostic and network planning purposes.
- Monitoring new or emerging protocols within your network (for example, where all Voice over IP elements reside and how they interact).

## How Optimizing SmartBPV Enhances Implementation

We recommend that you optimize SmartBPV so that it is easier to implement. Optimizing SmartBPV helps you manage the interaction of SmartBPV and Discovery.

The best practices for implementing SmartBPV include the following:

- Run a full discovery of your network before running SmartBPV.
- Do not run SmartBPV and Continuous Discovery at the same time.

If it is not practical to run a Discovery of your network before running SmartBPV, we recommend that you separate the running of SmartBPV and Discovery using one of the following methods:

- Run SmartBPV with the option to postpone discovery of unknown nodes. SmartBPV then runs without trying to discover unknown objects, and instead creates a list of objects to be discovered later. When SmartBPV starts, respond No when prompted about the deletion of SmartBPV PLOG files so that the files can be used again. Once SmartBPV completes, run Discovery to find these unknown objects and then start SmartBPV again.

To configure SmartBPV to postpone discovery of unknown nodes, modify `smartbpv.properties` to set these values:

- `DISCOVER_MISSING_OBJECTS = False`
- `DISCOVERY_SCRIPT = ./temp/SmartBPR_Discovery.script`
- `CREATE_MISSING_OBJECTS = No`

- Run SmartBPV to skip discovery of all unknown nodes and treat them as unclassified objects. Objects not already discovered will be unclassified in the repository until they are later discovered and classified.

To configure SmartBPV to skip discovery of unknown nodes, modify `smartbpv.properties` to set these values:

- `DISCOVER_MISSING_OBJECTS` = False
- `CREATE_MISSING_OBJECTS` = Yes

If it is not practical for you to run a Discovery of your network before starting SmartBPV and you require that SmartBPV be fully initialized and all objects discovered in a single step, run SmartBPV as a batch process when the load on your system and network is low. The syntax for this mode of operation is as follows:

```
smartbpv -nogui
```



# Chapter 5: Customizing Your Business Views

---

This section contains the following topics:

[Why You Need Unicenter Management Portal](#) (see page 189)

[CleverPath Portal Technology](#) (see page 190)

[Users, Workgroups, and Security Profiles](#) (see page 191)

[Scoreboards and Dashboards](#) (see page 191)

[Unicenter MP Administration](#) (see page 193)

[Workplace Templates](#) (see page 196)

[Working with Components](#) (see page 198)

## Why You Need Unicenter Management Portal

Unicenter Management Portal (Unicenter MP) delivers a consolidated view of enterprise management information provided by CA NSM components and other CA products using the familiar web portal model.

Unicenter MP complements CA NSM advanced visualization by doing the following:

- Delivering a process-oriented view of enterprise management information, focusing on intelligent delivery
- Personalizing and tailoring information to the needs of each user
- Providing an additional layer of access security to the management infrastructure

Unicenter MP helps to expose enterprise management information by means of the Web to a very broad audience, including network and systems administrators, performance groups, customers, business partners, and inside users. The portal's lightweight HTML-based interfaces provide fast and secure access to the information within and outside of the enterprise.

A new generation of Internet-based technologies introduces new challenges and opportunities. In Unicenter MP, you can tailor enterprise information to address the business needs of specific users. For example, Unicenter MP lets you provide access for the following users to information that is relevant to their business needs:

- Business customers, who routinely use the Web for other purposes, expect service providers to use the same technology to inform them of service level agreements, critical problems, status of specific resources, and simple tests they can use to verify service availability.

- Business partners need to be able to exchange enterprise management information between enterprises.
- Traditional CA NSM users (for example, enterprise management groups or network administrators) need to share information with other groups within the enterprise (for example, performance groups, application groups, IT managers, Business Executives, users, and so on).

Unicenter MP provides a clear, real-time view of information that makes it easy to understand how IT affects a particular business unit or application. It eliminates users from viewing lengthy status reports that do not apply directly to them in order to find information.

## CleverPath Portal Technology

Unicenter MP employs CleverPath Portal (CPP) technologies to deliver Unicenter content in a role-based portal environment. Unicenter MP installs an embedded version of CleverPath Portal 4.72 with the Light Search Engine only.

A Unicenter MP installation enforces the following restrictions on CleverPath Portal deployment and configuration:

- Unicenter MP does not support non-framed CleverPath templates.
- Unicenter MP supports only a subset of the databases that CleverPath Portal supports.
- Unicenter MP supports only a subset of the application servers that CleverPath Portal supports.
- Unicenter MP requires CPP to run with a private version of JRE 1.5.0.

**Note:** The same restrictions should be honored if Unicenter MP is installed on top of CleverPath Portal.

For additional information about CleverPath Portal, see the *CleverPath Portal Administrator Help* and *User Help*. Access these help systems within Unicenter MP by clicking the Help button on the top right of the interface (User Help) or clicking the link for either help system under the Help sub-tree on the Portal Administration page.

## Users, Workgroups, and Security Profiles

Unicenter MP implements role-based management of your enterprise information by controlling access to Unicenter portlets and the actions available in portlets based on the Unicenter MP workgroups to which a user belongs.

Unicenter MP lets you add users and organize them into workgroups. The workgroup a user belongs to controls access to portlets, channels, and library folders based on the view and modify permissions associated with these objects in the workgroup. Administrators can grant permissions to an individual user or to workgroups. For more information about users and workgroups, see the *Implementation Guide* and the *CleverPath Portal Administrator Help*, which you can access from the Help sub-tree of the Portal Administration page.

To control access to the actions available in a portlet, Unicenter MP uses the concept of security profiles. Each workgroup has a security profile associated with it. The same security profile can be mapped to multiple groups. Each security profile consists of a set of permissions that control individual actions in the Unicenter portlets. You, as administrator, can modify the existing profiles, or create a new one, using the Unicenter MP Administration Wizard.

You can access the wizard from the Unicenter MP Administration workplace that is available to the admin user. Also, as a member of the Admin workgroup, you can access the wizard by selecting Knowledge from the main Unicenter MP menu bar. In the left pane of the Knowledge page, select Library, Enterprise Management, Administration, \_UMP Administration Wizard.

## Scoreboards and Dashboards

Unicenter MP provides two types of views, scoreboards and dashboards. Both views are designed to display a high-level, summarized view of your systems, but they differ in the type of data they gather and the way they present it. The following descriptions describe these differences.

- A scoreboard is a summarized view that displays statistics for the objects stored in the Management Database, WorldView, or DSM. A scoreboard can include counts for a number of objects in a particular state or a number of events that have occurred during a particular interval. You can access the Portal Explorer from a scoreboard to drill down into an agent and view DSM-managed objects.
- A dashboard displays real-time information from CA NSM agents. A dashboard lets you combine (on one screen) multiple metrics from one or many agents and one or many hosts. Each metric is presented in an individual tile. Dashboards poll the data from the agents and show the metrics "as is." The dashboard is not a summarized view, and it is not just presenting statistics.

## Scoreboards and Dashboards Distributed with Unicenter MP

Unicenter MP supports several types of scoreboards and dashboards based on WorldView, DSM, Event Management, and other sources of information. A variety of predefined scoreboards and dashboards are distributed with the product, and some of these are included in the predefined workplaces. You can also view the predefined scoreboards and dashboards through the Unicenter MP Knowledge Tree.

Unicenter MP supports the following types of scoreboards and dashboards:

### **Alert Console**

Includes configuration and alert sections and lets you view and react to alerts.

### **Alert Scoreboard**

Includes configuration and alert sections and lets you view and react to a list of alerts with designated display attributes.

### **DSM-Agent Map Scoreboard**

Displays a summary status view of the data monitored by specific agents. Each row of the scoreboard presents details for a single agent.

### **DSM-Agent View Dashboard**

Displays a summary status view of the Unicenter agents you select.

### **DSM-Host Map Scoreboard**

Displays a summary status view of the agent data for a specific host. Each row of the scoreboard presents agent details for a single host.

### **DSM-Server View Dashboard**

Displays a summary status view of the Unicenter servers you select.

### **Event Console**

Shows network activity in the form of event messages. Each event message is displayed, by default, on the Event Console.

### **Events Scoreboard**

Displays a summary view of the Enterprise Management consoles you select.

### **eHealth Portlets**

Lets you publish eHealth portlets (eHealth Business Service Console, eHealth Business Service Console Ticker, and eHealth Reports) for any registered eHealth server. The portlets are published to the eHealth Performance\Report Server(\$eHealthServer\$) folder in the Unicenter MP Knowledge library, where \$eHealthServer\$ is the name of the selected server.



**Spectrum Portlets**

Lets you publish portlets for any CA Spectrum server that is connected to Unicenter MP.

**Unicenter Service Desk Portlets**

Lets you publish the Service Desk portlets from a specific Service Desk server to Unicenter MP.

**Unicenter SLA (Service Level Agreement) Scoreboard**

Displays a summary view of the SLAs you select.

**WV-Agent Status Scoreboard**

Displays a summary status of selected Unicenter agents and organizes them by agent type.

**WV-Business Process Views Scoreboard**

Displays a summary view of the Business Process Views objects you select.

**WV-System Status Scoreboard**

Displays a summary status of selected Unicenter systems and organizes them by host name.

## Unicenter MP Administration

Unicenter MP Administration is a set of tools for setting up, configuring, monitoring, and tuning the Unicenter MP server. Only members of the Admin workgroup have access to Unicenter MP Administration. The Admin workgroup is a default workgroup defined by Unicenter MP.

Unicenter MP provides an Administration Wizard that guides you through the most commonly used administration tasks.

Administering Unicenter MP includes two categories of tasks:

- General tasks required to administer the Unicenter MP server, including user and workgroup administration, starting and stopping the server, and changing the Admin user password.
- Tasks required to administer Unicenter NSM components, including setting up connections with the servers running Unicenter NSM and other products, enabling new components, tuning role-based security settings and monitoring the Unicenter MP infrastructure.

## Administration Wizard

Unicenter MP provides an Administration Wizard to guide you through the most commonly used administration tasks. The UMP Administration Wizard is the default workplace for users with admin privileges. Also, as a member of the Admin workgroup, you can access the UMP Administration Wizard from the Knowledge Library under Administration by clicking \_UMP Administration Wizard.

**Note:** Only members of the Admin workgroup have access to the UMP Administration Wizard. The Admin workgroup is a default workgroup defined by Unicenter MP.

From the wizard, you can launch the following tasks:

### **Task 1: Manage Components**

Establishes a connection to hosts running other CA components, such as WorldView, Agent Technology, and Event Management. Defining a host as a data source lets Unicenter MP obtain and display data from that host. Complete Task 1 before moving on to other tasks.

### **Task 2: Create or Modify Unicenter MP Portlets**

Lets you define scoreboards or dashboards, which are real-time, query-based summary views of your data. Scoreboards and dashboards appear in the Unicenter MP Library, making them available to your Unicenter MP users.

### **Task 3: Manage Scheduled Tasks**

Lets you select a scheduled task and change the status, suspend execution, reset the next execution, resume a suspended execution, or delete the task.

### **Task 4: Portal Administration**

Launches CleverPath Portal Administration, letting you perform administrative tasks such as creating users, creating workgroups, assigning users to workgroups, and more. Unicenter MP is based on CleverPath Portal technology.

### **Task 5: Manage Users**

Manages user profiles by letting you add, edit, or remove users and assign them to workgroups.

### **Task 6: Manage Workgroups**

Lets you define, edit, or remove user workgroups. Workgroups help organize users into logical business groups, such as systems administrators, business users, and mail administrators. As members of a workgroup, users inherit permissions assigned to the entire group.

**Task 7: Manage Unicenter Management Portal Properties**

Configures properties for Unicenter Configuration Management, Service Level Agreements (SLAs), Business Process Views (BPVs), scoreboards, reports, the knowledge base, IPv6 addresses, and security.

**Task 8: Manage Security Profiles**

Lets you define, edit, or delete security profiles. Security profiles control access to specific data and controls for all actions you can perform. Assigning a security profile to each workgroup further defines the security permissions for that workgroup.

**Task 9: Manage Global User Preferences**

Lets you define or edit user preferences. Although users can specify personal display and data handling preferences, you can override individual preference settings for all users, if needed.

**Task 10: Manage Web Reporting Servers**

Establishes connections to hosts running Web Reporting Server (WRS), making the reports running on these servers available to Unicenter MP users.

**Task 11: Manage Unicenter Management Portal Security Reports**

Lets you view the Unicenter MP reports related to Documents, Channels, Workplace Templates, Menu Actions, BPVs, SLAs, and Event Filters.

## Task 1: Manage Components

Unicenter MP relies on CA NSM and other CA products to provide the Enterprise Management data displayed in Unicenter portlets. To obtain the data from these products, Unicenter MP must be configured with such information as location of servers, user ID and password for databases, and so forth. The type of information provided depends on the Unicenter product to which that portal connects.

You must complete the task of defining connections to Unicenter management servers before you can view any Unicenter information in Unicenter MP. The Unicenter MP Administration Wizard guides you through this task. You can access the wizard from the Unicenter MP Administration workplace that is available to the admin user. Also, as a member of the Admin workgroup, you can access the wizard by selecting Knowledge from the main Unicenter MP menu bar. In the left pane of the Knowledge page, select Library, Enterprise Management, Administration, \_UMP Administration Wizard.

Clicking Task 1.Manage Components opens the Manage Components page, where you can view and define component connections, discover data sources, and perform other tasks related to component connections.

**Note:** For detailed information about the rest of the tasks you can perform from the Unicenter MP Administration Wizard, see the Unicenter Management Portal Help.

## Workplace Templates

Workplaces are custom pages that display only the Unicenter MP data you want. You can fully customize the content and layout to suit your needs. Unicenter MP provides the following templates for creating workplaces:

### **Empty Workplace**

Contains no preconfigured attributes. If another template does not fit your needs, use this one.

### **Application Servers Status**

Includes only application server data, such as the Application Agents Status or Application Events Summary scoreboards.

### **Database Servers Status**

Includes only database server data, such as the Database Agents Status or Database Events Summary scoreboards.

### **Mail Servers Status**

Includes only mail server data, such as the Mail Agents Status Breakdown or Mail Events Summary scoreboards.

### **Messaging and Transaction Servers Status**

Includes only messaging and transaction server data, such as the WorldView Transaction/Messaging scoreboard.

### **My Unicenter Workplace**

Includes the basic data to get you started with Unicenter MP.

### **Network Status**

Includes only network data, such as the Network Agents Status Breakdown or Network Events Summary scoreboards.

### **Systems Status**

Includes only systems data, such as the System Agents Status Breakdown or System Events Summary scoreboards.

### **UNIX Systems Status**

Includes only UNIX systems data, such as the UNIX System Agents Status Breakdown or UNIX System Events Summary scoreboards.

**Web Servers Status**

Includes only web server data, such as the Web Agents Status Breakdown or Web Server Events Summary scoreboards.

**Windows Systems Status**

Includes only Windows systems data, such as the Windows System Agents Status Breakdown or Windows Agent Events Summary scoreboards.

## Create Workplaces from Templates

You can create workplaces to help customize the content you want to see when you log in to Unicenter MP. When you use templates to create workplaces, your workplace is populated by data by default that relates to the template you selected. You can edit the content that is included by default so that it fits your needs.

**To create workplaces from a template**

1. Click Knowledge, expand Library, Enterprise Management, My Unicenter, and click Create Workplace.

The Manage Workplaces page appears.

2. Select a workplace template and click Next.

The workplace page for the template you selected appears.

3. Enter a name for your workplace in the Title field and click Finish.

The template is created and saved, and a confirmation screen appears. The workplace is populated with default content related to the template when you access for the first time. Edit the content by clicking Add Content on the workplace screen.

**Note:** If you select the Empty Workplace option, you must go through additional steps to add content and a framework for your workplace. Therefore, use one of the existing workplaces if you want the workplace to be populated by content related to the template by default.

For more information about creating an empty workplace, adding content to workplaces, and assigning workplaces to users and workgroups, see the chapter "Customizing Unicenter MP" in the *Implementation Guide*, or see the CleverPath Portal Administrator Help, which you can access from the Help sub-tree of the Portal Administration page.

**More Information**

[Workplace Templates](#) (see page 196)

## Working with Components

You can establish connections in Unicenter MP to various CA NSM components to gain access to the data and functionality provided by those components. You can also establish connections to other CA products for a view of interfaces and data provided by those products. You can connect to the following CA NSM components and other CA products to view and take action on data from these components and products in the Knowledge Library:

- CA eHealth
- CA Spectrum
- Agent Technology
- Alert Management
- Event Management
- Service Desk
- Service Metric Analysis
- WorldView

You can establish connections to the following Unicenter NSM components and other CA products to gain access to data provided by these components and launch component and product interfaces from the Portal Explorer:

- Active Directory Explorer
- Wily Introscope
- Unicenter Configuration Management
- Unicenter NSM Knowledge Base
- Unicenter Systems Performance

## Working with Unicenter WorldView

Unicenter MP lets you effectively monitor the enterprise infrastructure resources that are discovered and stored in the Unicenter WorldView repository. You can monitor resources in two dimensions:

- Vertically by business process view
- Horizontally by a managed resource group, such as a network, system, database, application, and server

The monitoring always starts with a scoreboard, which is a high level status summary of the IT resources. During the installation, Unicenter MP creates many scoreboards for both business process views and resource groups, so you can immediately start to use them once installed. Unicenter MP also provides the tools and facilities to let you customize existing scoreboards or create new scoreboards.

**Note:** Unicenter MP, by default, counts only managed WorldView objects in business process views, system status, and resource status scoreboards. However, you can change the settings to count unmanaged objects for business process views scoreboards.

The Portal Explorer, the Severity Browser, and the Severity Tree let you view detailed information about your IT resources. You can launch these detailed interfaces from the status scoreboards. You can also access the Business Process View Explorer, which lets you view Business Process Views in the context of the Portal Explorer.

You can create reports for your IT resources in the WorldView repository to monitor the resource status or to find the topology or containment information. You can create a sophisticated schedule in Unicenter MP to run these reports at certain times and intervals and notify the appropriate people when they are published.

## Business Process View Scoreboards

Business Process View scoreboards provide a summarized status view of the various resources in your enterprise, based on Business Process Views. The information can be obtained from one or multiple WorldView servers. With this type of scoreboard, you can monitor the status of the enterprise infrastructure that is relevant to your business function - quickly identifying problem areas that might impact your business processes.

Business Process View scoreboard creation and publication usually occurs as follows:

1. The Unicenter MP administrator analyzes and determines which users need information about which Business Process Views.

For example, users in London need to know the status of the London Data Center Business Process View; users in New York need to know the status of the New York Data Center Business Process View; Lotus Notes administrators need to know the status of the Lotus Notes Servers Business Process View; and Unicenter MP administrators need to know the status of all the Business Process Views in the enterprise.

2. The Unicenter MP administrator uses the Business Process Views Scoreboard Publishing Wizard to create and publish different and appropriate scoreboards for each group of users. For example, if the workgroup LondonUsers exists for London users, NYUsers for New York users, LNAdmins for Lotus Notes administrator, and Unicenter MPAdmins for Unicenter MP administrators, then the Unicenter MP administrator can create and publish four different Business Process Views scoreboards, each containing the appropriate Business Process Views for each workgroup. The administrator can also configure the scoreboards so that when they publish, automatic email notification is sent to the appropriate users.
3. Users connect to Unicenter MP and see the Business Process View scoreboard that is published for the workgroups to which they belong.

**Note:** A Business Process View scoreboard can show Business Process Views information from multiple WorldView repositories.

### Resource Scoreboards

Other than using Business Process View scoreboards to track the status of the enterprise resources important to your business processes, you can use resource scoreboards to horizontally track the status of your enterprise by the resource groups (classes) such as network, system, application, and so on.

Unicenter MP creates a set of resource scoreboards during installation. You can customize these scoreboards or create your own scoreboards to accurately monitor the resources important to your environment.

**Note:** Unicenter MP counts only managed WV objects in resource status scoreboards.



## Portal Explorer

The Unicenter MP Portal Explorer lets you view the relationships between BPVs or managed objects in tree form and view the details of a selected object in the object tree. You can launch the Portal Explorer from either Business Process View scoreboards or Resource scoreboards.

The Portal Explorer is broken into an interactive tree in the left pane and the corresponding views and tasks in the right pane. The left pane includes a tree structure list of objects for navigation purposes. For example, when the Portal Explorer launches from a Business Process View scoreboard, the left pane displays Business Process View objects. When the explorer launches from a managed resource scoreboard, the left pane tree shows the resource group (class) as the root, and underneath the root, lists the objects that belong to the class.

Upon selecting an object in the left pane, the right pane displays various tabbed views about that object. The displayed views depend on the type of selected object. The views include the following:

### **Object Severity Browser**

Appears for WorldView objects and DSM objects.

### **Notebook**

Appears for WorldView objects and DSM objects. The view shows the property and value pair of the selected object. For a WorldView object, the properties are grouped and showed in different sub-tabbed pages.

### **Explorer**

Appears if you select a virtual root object, such as a resource group. The view displays the name of the objects that belong to the selected resource group (class).

### **Event**

Appears for DSM objects.

### **Event Console**

Appears for Node (Hosts and Workstations) WorldView objects. The view is actually the event console that shows all events that are sent from the selected node.

### **Alert View**

Appears for alert-aware WorldView objects, including BPV, Node, Network Segment, Router, and so on. The view shows in the alert console all the alerts that are associated with the selected object.

By right-clicking to select a Worldview object in the tree, you may see menu pop ups. You can act on the object based on the available menu items. Also, by selecting an object in the right-hand pane, menu actions may appear depending on the type of object. For example, you can launch Performance Reports, agent dashboards, and other interfaces such as Unicenter Configuration Manager and Active Directory Explorer.

As a Unicenter MP administrator, you can set up a security policy to secure some views from certain roles of users in your organization.

## Severity Browser

The Severity Browser shows status details of the objects that are contained in a selected Business Process View or managed object. The two types of severity browsers are:

- BPV severity browser
- Object severity browser

The status details in these views include the object name, label, severity, propagated severity, maximum severity, class name, and IP address.

In the browser, you can configure the starting object status and object level.

You can launch the severity tree in the severity browser to quickly identify the root cause of an abnormal Business Process View object or managed object.

## Working with Agent Management

In Unicenter NSM, agent management deals with your ability to monitor various enterprise resources like systems, databases, and applications. The agents collect the specified metrics from the resources and report back to the designated manager. Data returned from agents helps you to quickly take corrective actions and keep your systems running properly. The status of the agents is managed by the DSM (Distributed State Machine). You access the agents through a designated DSM manager host in the enterprise.

Most agents let you customize what you want to monitor through the use of watchers. You can define and modify the watchers with the classic GUIs or with Unicenter MP.

The agent scoreboards in Unicenter MP provide summary views of the various agents and their status for different resource groups. This gives you a high-level view of the status of their resources.

The agent dashboards in Unicenter MP provide a more detailed state of the agents. You can drill-down into a single agent and also add, modify, or delete the watchers for that agent along with other configurations of the agent.

You can also schedule and run agent reports to show the status of DSM objects that meet certain criteria. Unicenter MP provides a set of predefined reports, and you can also create reports by customizing the predefined reports or using provided report templates.

Before accessing the agent scoreboards and dashboards or running reports, you need to define a connection to a DSM host. Specify this DSM connection when you create the scoreboards or dashboards.

### Agent Map Scoreboards

Unicenter MP lets you create and publish high-level, summarized views of your systems, called scoreboards. Agent Map scoreboards are one type that provide a simple visual representation of the current state of an agent and its sub-components. With this scoreboard, users can monitor the status of agents that are relevant to their job function, quickly identifying problem areas that might impact job processes.

The color-coded map shows the state of each agent. Other attributes may appear for the agents in addition to status. You can view the full attribute name by placing your cursor on the attribute.

Each row in the scoreboard represents a single agent, and you can select each row to make available any actions that you can perform on the agent in the Select Link drop-down list, such as launching the Portal Explorer and launching agent and server dashboards. You can also click the icon next to the agent name to quickly perform the default menu action for the agent (which is often launching the Portal Explorer).

This type of scoreboard is particularly useful for network administrators, systems administrators, and database administrators who need to monitor specific components in the systems for which they are responsible.

**Note:** When creating a scoreboard, only select the exception level that is meaningful. That is, do not select Normal status when all you really need to see are agents in a Warning state and worse. Also, only select agents of the same type (for example, mibmixed or non-mibmixed).

### Host Map Scoreboards

Host Map scoreboards are one type of scoreboard that provide a simple visual representation of the current status of all agents, grouped by host. With this scoreboard, users can monitor the status of agents that are relevant to their job function, quickly identifying problem areas that might impact job processes.

Each row in the scoreboard represents a single host and its agents. The state column shows the status of the host, while additional columns show the status of each agent. You can view the full agent name by placing your cursor on the agent column heading.

You can select each host row to make available any actions that you can perform on the host in the Select Link drop-down list, such as launching the Portal Explorer and Server Dashboards. You can also click the icon next to the host name to quickly perform the default menu action for the host (which is often launching the Portal Explorer).

This type of scoreboard is particularly useful for Network Administrators, Systems Administrators, and Database Administrators who need to monitor specific components in the systems for which they are responsible.

**Note:** When creating a scoreboard, only select the exception level that is meaningful. That is, do not select Normal status when all you really need to see are agents in a Warning state and worse. Also, only select agents of the same type (for example, mibmixed or non-mibmixed).

## Agent View and Server View Dashboards

Unicenter MP lets you create and publish high-level, real-time views of your agent data, called dashboards. Agent view and server view dashboards are two types that provide a personalized, consolidated view of the agents through secure, web-based interfaces.

Agent view dashboards display summary information about only the agents you select. However, server view dashboards display summary information about all agents on a server you select.

Both dashboard types provide the following information:

- A high-level, normal or general-health view of the enterprise
- A diagnostic view of exceptions that occur in the enterprise

Based on these views, Unicenter MP provides the following two modes of viewing the dashboards:

### **Normal mode dashboards**

Provide a high-level view of the agents by displaying all of the agents' tiles. For a given agent, the agent dashboards display one tile for each monitored group. By clicking on the link within the tile, you can drill down into a specific monitored group to get more information.

### **Exception mode dashboards**

Provide a diagnostic view of the agent by displaying only the tiles that have a specific abnormal status. For a given agent, you can specify the exception level, and only the tiles with that status or worse display.

Based upon the content or tile-definition, the agent view dashboards can be one of the following:

### **Default dashboards**

Contain all available tiles for a given agent. They are defined in the tile-definition called "General." The default dashboard is available out-of-the-box.

### **User-defined dashboards**

Contain the tiles that you define in a specified tile definition. In certain cases, you may only be interested in monitoring a small set of resources. In such cases, you can define a tile-definition containing only those resource groups. The dashboard only displays resources specified in that tile-set.

**Note:** The dashboard can still be shown in either Normal or Exception mode.

### **Server Exception dashboards**

Combine information from all of the agents running on that server and present in one dashboard. It includes all of the tiles that are in exception mode for all of the agents.

You can create agent view or server view dashboards using the Administration Wizard, or you can launch the dashboards in context from scoreboards, reports, and the Portal Explorer. You can also publish dashboards to the Unicenter MP Knowledge Tree.

**Note:** Status information for the agents is obtained from the DSM. This means that Unicenter MP should be configured to obtain the DSM information through a Component Manager connection.

## **Working with Unicenter Event Management**

Unicenter Event Management is a CA NSM component for collecting and responding to the events in your enterprise or a logical segment of your enterprise.

The Unicenter MP Event Management component presents the event messages in the event logs that the Unicenter Event Manager manages. By leveraging Unicenter MP filter management, you can create consoles and scoreboards in Unicenter MP to get the events and statistics of events from different areas of your enterprise. Also, you can set up event managers to send certain types or groups of messages to Unicenter MP to present as notifications.

From Unicenter MP, you can act on event messages whenever necessary in the following ways:

- Change display attributes
- Export up to the last 10000 event messages or selected messages into HTML, CSV, and PDF reports
- Search for specific messages
- Acknowledge selected held messages simultaneously

Events scoreboards in Unicenter MP provide the following four types of statistics for event messages:

#### **Summary scoreboards**

Provide the total number of messages.

#### **Breakdown scoreboards**

Provide a status bar for all messages of a group or filter and separate counts for each status type found.

**Last N**

Provides the last N number of messages that match the scoreboard filter criteria.

**Dynamic chart**

Provides a dynamic chart of the message severity breakdown. The chart is updated periodically.

The Event Console shows the detail of the messages. You can launch the Event Console from event scoreboards or the Knowledge Tree. In the Event Console, you can act on events whenever necessary, and you can search for events, filter events, and publish events to an HTML, CSV, or PDF report. You can set up the security policies to allow or disallow a user to take certain actions. The Event Console is presented in the Portal Explorer for a Node (Host and Workstation) object and is accessible from the Knowledge Library.

In Unicenter MP, you can specify users from certain workgroups as event administrators. Event administrators can change the data configuration of event consoles and scoreboards. Non-event administrators can change the presentation configuration of event consoles, but need to get permission to change the scoreboard configurations.

**Event Scoreboard**

Event scoreboards provide summarized views of messages regarding events that may occur within your IT infrastructure. Predefined event scoreboards reflect various aspects of your infrastructure including applications, databases, mail servers, networks, systems, transactions, and messaging. You may also create customized scoreboards to track events from specific sources that are significant to you.

Individual scoreboards for each resource type let you see only event messages for resources that are of interest to you. For example, if you want to know how many critical messages exist for your database server only, you can select the appropriate scoreboard and drill down into the Event Scoreboard Detail to view specific events.

You can also assign specific scoreboards to personnel who are responsible for that aspect of your infrastructure. For example, you can create a scoreboard that captures events coming from only your SNA Manager Agents, Switch Status, and Chassis Monitoring Agents and assign it to your network administrator.

Consolidation of these scoreboards centralizes management of different aspects of your infrastructure, ensuring that important events are acted upon in a timely manner.

## Event Console

Using filters, you can create an Event Console to scope the events from a specific area of your enterprise infrastructure. The Event Console lets you monitor the status of the events, respond to abnormal events as they occur, and rectify critical situations immediately.

The Event Console organizes the events in pages and shows several specified display attributes.

Using the filtering and searching facility provided in a console, you can narrow your search further and view the events that you want.

Due to the sequential nature of the event log, event messages in a console are sorted by the creation time of the messages in descending order. The most recent event appears at the top of the first page.

By default, the console is automatically refreshed within one minute. You can change your preference to disable the auto-refresh or increase the refresh interval.

You can take numerous actions on events from the Event Console, including acknowledging events, viewing event annotations and details, and publishing held or log messages to a report.

A predefined Event Console, is published at Library, Enterprise Management, Unicenter Event Management. You can use publishing tools to publish a customized Event Console at the same location in the library.

## Event Actions

Once you have selected an event message in the Event Console, you may take several actions, as follows:

- Acknowledge
- Reply
- View Detail
- View Annotations
- Export to an HTML, CSV, or PDF report

**Note:** Acknowledge is available only on messages that are waiting to be acknowledged, while Reply is available only for WTOR messages. The type of message is represented by different icons in the Attribute column in the Event Console.

If you are a Unicenter MP administrator, you can set up the security policies on the action a user can take on events.



## Manage Event Filters

You can create and manage event filters to group the events from different areas of your enterprise infrastructure. Thereafter, you can use the filters to create event scoreboards and consoles and assign the scoreboards and consoles to the different roles of administrators in your organization so that they are able to monitor the events that are relevant to their jobs.

Using Unicenter MP filter management tools, you can create, modify, and edit event filters.

To start the filter tool, click Knowledge and select Library, Enterprise Management, Unicenter Event Management, Manage Event Filters.

The Filter Group page appears. You can manage your event filters from this page.

## Working with Unicenter Alert Management

The Unicenter Alert Management System (AMS) is a CA NSM component for organizing and tracking the most important events in an enterprise or a logical segment of an enterprise. It significantly improves your ability to focus on and manage the highest severity IT events.

Alert priority is calculated from alert impact and urgency. The lower the priority value is, the higher the priority. Alerts are classified in classes and are organized into queues. The alert class defines actions to take upon the status change of the alert. The alert can be defined to associate with a knowledge base developed by the IT staff to assist in the evaluation and handling of the alerts. With the alert queue context menu facility, AMS also provides external interfaces to allow linkage of alerts to user data and handling procedures and third party applications and knowledge bases. For example, database alerts could be defined to allow the launch of database management tools and interface to the database management online documentation.

AMS lets you link to Unicenter Service Desk to let you create requests on an alert, view alert requests, and search for information that is related to an alert.

While both the classic interface and MCC provides various tools to manage Unicenter AMS, Unicenter MP Alert Management only focuses on monitoring alerts. The high-level queue-based alert scoreboard shows the number of alerts that fall into priority ranges in the different queues. Leveraging Unicenter MP filter management, you can create various alert consoles that show the alerts from different areas of your enterprise infrastructure. You can also set up the security policies to allow or disallow a user to take certain actions.

From Unicenter MP, you can act on alerts whenever necessary in the following ways:

- Create requests, incidents, and problems from an alert
- View details
- Acknowledge selected alerts simultaneously
- Transfer selected alerts simultaneously
- Raise an alarm
- Consolidate or unconsolidate alerts
- Add, view, modify, and delete annotations
- Close alerts
- Change urgency
- Change display attributes
- Export up to the last 5000 alerts or selected alerts into HTML, CSV, and PDF reports
- Search for specific alerts
- Access URLs associated with alerts
- View alert audit trails

The alert view is presented in the Portal Explorer for Business Process Views and some managed objects, such as nodes (Host and Workstation), network segments, and routers.

Unicenter MP provides tools for you to manage alert scoreboards and the Alert Console. While all users can change the scoreboard presentation configuration, only alert administrators can change the data configuration.

## Alert Scoreboard

Alert scoreboards display a graphical representation of your alert breakdown in different priority ranges. Alert scoreboard data is obtained from the AMS alert queue.

Alert scoreboards let you perform the following actions:

- Change the graphical presentation
- View alerts from specific queues
- Configure the scoreboard if you are an alert administrator

By default, the scoreboard is set to refresh automatically within one minute. Through My Preferences, you can disable the auto-refresh function or change the refresh interval.

Two predefined alert scoreboards are published at Library, Enterprise Management, Unicenter Alert Management, Configured Scoreboards. You can use the publishing tool to create customized scoreboards at the same location in the knowledge library.

## Alert Console

The Alert Console organizes the alerts in pages and shows the alerts with specified display attributes. You can create an Alert Console to view the alerts from a specific area of your enterprise infrastructure. The console also lets you do the following activities:

- Navigate the alerts page by page
- Sort alerts by all displayable properties
- Change the display attribute of an alert
- Take a number of actions on alerts
- Use the filtering and searching facility tool to further specify and get the alerts in which you are really interested
- Export alerts to HTML, CSV, and PDF reports
- View URLs associated with alerts

By default, the Alert Console is automatically refreshed within one minute. You can change your preferences to disable the auto-refresh or increase the refresh interval.

A predefined Alert Console, which shows all alerts, is published at Library, Enterprise Management, Unicenter Alert Management. You can use publishing tools to publish a customized Alert Console at the same location in the library.

## Alert Actions

Once you have selected an alert, you can select one of the following actions from the 'Select and' drop-down list and click Go:

- Create Request
- View Requests
- Search Knowledge Tool
- eHealth Report Server
- Business Service Console
- At-a-Glance Report
- Alarm Detail Report
- View Detail
- View Annotations
- Acknowledge
- Alarm
- Transfer
- Unconsolidate
- Consolidation Detail
- Close
- View Audit Trail

Unicenter Service Desk-related actions are available only when a connection to a Service Desk server is established in Unicenter MP. The following actions are available only when the Service Desk server connection is running in ITIL mode:

- Create Incident
- Create Problem
- View Requests only
- View Incidents only
- View Problems only

eHealth-related actions (eHealth Report Server, Business Server Console, and At-a-Glance Report) apply to eHealth alarms, alerts, and exception alerts. The Alarm Detail Report action applies to eHealth alarms only. eHealth actions do not appear for non-eHealth alerts.

eHealth alerts are created with corresponding eHealth server information. Even if there is no eHealth server registered with Unicenter MP, the eHealth actions still appear for eHealth alerts.

The first time you access an eHealth server in a Unicenter MP session, a login window appears (unless you have enabled EEM security in Unicenter MP). After successful authentication, access to the server remains valid for the rest of the session.

**Note:** If you are a Unicenter MP administrator, you can create security policies on the actions a user can take on alerts.

## Working with Unicenter MP Notification

The Unicenter MP Notification component lets users view notifications sent to them or the workgroups to which they belong.

**Note:** Notifications sent to a specific workgroup are received only by the members of that group. Even if you are a member of a workgroup that is superior in the user hierarchy to the specified group, you will not receive the notification unless the notification is targeted to you. For example, a notification sent to members of the TNDUsers workgroup is received by the members of the TNDUsers group, but not by members of the admin workgroup, even though the admin workgroup is superior in the user hierarchy.

A notification contains the following properties:

**Severity**

Specifies the severity level of the notification (Normal, Warning or Critical).

**Time**

Specifies the time when the notification was sent.

**Message**

Specifies the notification text.

**Category**

Specifies a category to group the notifications.

**URL**

Links to more information about the notification.

**Group/User**

Specifies the workgroup or user name to which the notification is sent.

To send a notification to Unicenter MP, an integrated application must invoke the Unicenter MP notification utility java class. With the automatic responding facility, you can set up the Unicenter Event Manager to send the notification to Unicenter MP when it receives certain types of event messages.

Notifications are saved in the MDB. By default, the notifications never expire. However, by using Task 7: Manage Unicenter Management Portal Properties in the Administration Wizard, you can set the expiration time in hours. Unicenter MP automatically removes a notification once it expires.

A predefined notification (My Notifications) is published at Library, Enterprise Management, Notifications. The My Notifications link is also in the My Unicenter workplace in the My Unicenter menu. This link shows all notifications sent to the workgroup that the current user belongs to.

Another predefined notification (UM Portal Notifications) is published at Library, Enterprise Management, Notifications. The link is accessible for admin users only. This notification shows those notifications sent from various Unicenter MP components that report their status.

Through the predefined UM Portal Notifications notification, you can create and publish a notification with a filter on the notification properties.

## Working with Unicenter MP Reports

Web Reports let you view all of your historical performance data through an Internet browser in a way that is meaningful to you. In Unicenter MP, a report can be one of the following:

- Configured reports are out-of-the-box reports you can immediately use to access meaningful information from supported products. For example, you can use the All Windows System Status report to view details about all of the Windows servers monitored by agents from the DSM connected to Unicenter MP.
- Report templates provide a structure and let you select the specific criteria you want to display in your report. You can use report templates to specifically tune the information and create your own configured reports or published reports.

For example, if you are concerned about x factor within a supported product, you can execute the x factor configured report, which provides summarized information on x factor activity for that product. But if you want to view more specific information on y within the x factor, you can fill in the y within the x factor template provided by that product to define your own configured report, or publish the report into the product tree so you can retrieve it. Web Reports provide several report templates and configured reports across supported products that let you see your data the way you want to see it.

- Published reports are the resulting reports that display the actual report data.

Report data represents only a snapshot in time, and the data is quickly dated after the report is generated. However, you can schedule your reports to run at specific intervals, ensuring that your report data is continually updated.

When you schedule a report, you can specify the date and time on which the report will run. You can specify whether the report will run on certain days of the week, a certain day of the month, or certain months of the year. You can specify what date the report starts running, whether the report runs multiple times in one day, and on what date the report stops running. You can also schedule to send automatic email notifications to any number of users every time a report is published.

You can also launch other interfaces or perform actions on selected objects in reports. When you select an object, the 'Select and' drop down list is populated with the interfaces you can launch for the object and the actions you can perform on it. For example, you can launch the Portal Explorer for an object, or you can open Unicenter Configuration Manager for an agent with the Policy Configuration option.

## Working with Unicenter Service Metric Analysis

Unicenter Service Metric Analysis lets network and systems management personnel responsible for performance and reporting produce reports and statistics for resources available inside and across the IT infrastructure. Part of Service Metric Analysis is the Service Level Agreement (SLA) component, which lets administrators and business managers understand and manage predefined service goals across all IT resources.

The SLA component of Unicenter MP provides a high-level summarized view of SLAs with the ability to drill down into detailed Unicenter Service Metric Analysis reports. The SLA component of Unicenter MP includes the following features:

- SLA scoreboards
- SLA Scoreboard Wizard
- Publishing Unicenter Service Metric Analysis reports under Unicenter MP

For more information about how to publish Unicenter Service Metric Analysis reports under Unicenter MP, see the *Unicenter Service Metric Analysis User Guide*.

**Important!** The SLA component of Unicenter MP requires Unicenter Service Metric Analysis Version 2.0 or newer to be installed and running in your enterprise.

## Working with Unicenter Service Desk

Unicenter MP provides integration with Unicenter Service Desk. After adding a connection to a Service Desk server, you can publish Service Desk portlets into the Unicenter MP Library to do the following tasks:

- Open, update, or close Service Desk incidents
- Ask the Service Desk a question or search the knowledge base
- Review questions and answers encountered by others

When establishing a connection to Service Desk, you can specify whether Service Desk is configured to run in ITIL mode. The ITIL Service Desk interface supports additional data objects, and additional alert actions are available in Unicenter MP when you connect to an ITIL Service Desk server.

You can specify how to log in to Service Desk from Unicenter MP. You can log in using the Service Desk login screen, Unicenter MP login credentials, or a specific user name and password that you specify. You can also create a Service Desk portlet, which lets you open the Service Desk screen within Unicenter MP.

If you also added a connection to the Unicenter Alert Manager, you can take the following actions on alerts in the Alert Console related to Service Desk:

- Create Alert requests
- Create Alert incidents
- Create Alert problems
- View Alert incidents
- View Alert node requests
- View Alert node incidents
- View Alert node problems
- Search the knowledge tool

**Note:** For more information about how to define Service Desk connections and publish Service Desk portlets in Unicenter MP, see the Unicenter Management Portal Help.



## eHealth Integration with Unicenter MP

Unicenter MP provides access to the eHealth suite of products so that you can monitor the performance of eHealth objects and generate reports. The eHealth Suite delivers comprehensive fault, availability, and performance management across complex, heterogeneous systems and application environments. eHealth collects a wide variety of data from your network infrastructure to generate alarms and reports.

From Unicenter MP you can access the following features of eHealth:

### **Business Service Console**

Provides a high-level view of the availability and performance of business services across an organization. The eHealth Business Service Console (BSC) is a Web-based tool that offers customized business views, immediate notification of performance problems, and drill-down capability for fault resolution. You can access the BSC in Unicenter MP from the knowledge tree or the Alert Console.

### **Business Service Console Ticker**

Shows a minimized view of the Business Service Console in the corner of your desktop. The BSC Ticker provides a subset of the hierarchy represented in the BSC and easily expands to show the full console. If a performance problem occurs, an acknowledgement indicator shows whether someone is working to resolve it, and a duration indicator shows how long the object has been in the current state of red or yellow. You can access the BSC Ticker in Unicenter MP from the knowledge tree.

### **eHealth Report List and Web Interface**

Shows all reports that have been generated on the eHealth system. Access the eHealth Report List in Unicenter MP from the knowledge tree or from the Alert Console (for a specific alarm). The Report List is part of the eHealth Web Interface, which provides access to reports and applications through Web browsers on local systems. When you open the Report List, you can also view other areas of the eHealth Web Interface.

### **At-a-Glance Reports**

Show the overall performance of an eHealth object for a specified time period. The reports consist of several charts that display different performance statistics on one page. You can access At-a-Glance reports in Unicenter MP for WorldView Topology and DSM from the Portal Explorer and for alerts that represent eHealth alarms from the Alert Console.

### **Alarm Detail Reports**

Show the availability and performance history over time of an eHealth object that caused an alarm to be generated. You can access Alarm Detail reports in Unicenter MP from Portal Explorer or the Alert Console.

### Trend Reports

Plot a variable for an object over a period of time in chart form. Trend reports can also show variables for groups of objects. The reports can reveal patterns over time and relationships between objects and between variables. The available Trend reports are Availability, Bandwidth, and Error, depending on the type of managed object. You can access Trend Reports in Unicenter MP for WorldView Topology and DSM objects from the Portal Explorer.

### eHealth Alarms and netHealth Exceptions

Create Alert Management System alerts in Unicenter MP automatically, based on policy that you deploy. When alarms are closed in eHealth, the associated alerts are closed. Likewise, if an alert associated with an eHealth alarm is closed through AMS, the alarm is also closed. Access alerts representing eHealth alarms and netHealth exceptions in Unicenter MP from the Alert Console.

Unicenter MP also offers single sign on capability with Embedded Entitlements Manager (EEM) security. When you enable EEM security in Unicenter MP, you do not have to enter credentials to access eHealth features.

## How the eHealth Integration Works

Unicenter MP provides access to eHealth reports and applications so that you can monitor the availability and performance of eHealth objects. The eHealth integration works as follows:

- An administrator specifies connection parameters to eHealth servers. Multiple servers are supported. If you have multiple connections, the default connection is used for launching reports for a managed object from the Portal Explorer.
- The Knowledge Tree provides links to the eHealth Business Service Console, Business Service Console Ticker, and eHealth Reports.

**Note:** Accessing eHealth Reports takes you to the Report List, which is a part of the eHealth Web Interface. When you access the Report List, you can also access other areas of the eHealth Web Interface.

- A workplace can display the Business Service Console Ticker, but not the Business Service Console or eHealth Reports. They take over the whole browser window.
- The Portal Explorer provides links to the eHealth reports. You are prompted for a user name and password the first time you access a report during a Unicenter MP session.
- Trend and At-a-Glance reports are displayed for WorldView Topology and DSM objects. At-a-Glance and Alarm Detail reports are displayed for AMS alerts that represent eHealth alarms.

- The Alert Console provides access to the Business Service Console and the Report List, Alarm Detail reports, and At-a-Glance Reports for alarms appearing as alerts.
- The following new scoreboards are available for eHealth objects in the WorldView Topology and alerts associated with eHealth alarms:
  - eHealth Trap Status Breakdown Scoreboard (EM)
  - eHealth Interface Status Breakdown Scoreboard (WV)
  - eHealth Status Breakdown Scoreboard (WV)
  - All eHealth Alerts Breakdown (Alert)
- When you try to access any of these features for the first time in a Unicenter MP session, you are prompted for login credentials. However, Unicenter MP offers single sign on capability with EEM security. When you enable EEM security in Unicenter MP, do you not have to enter eHealth login credentials to access eHealth features.

## Working with SPECTRUM

SPECTRUM is a network fault management system that provides proactive management of your network infrastructure through root cause analysis, impact analysis, event correlation, and service level management. It proactively enables you to manage complex, heterogeneous multivendor network environments by automatically identifying the cause of network problems, suppressing symptomatic alarms, and highlighting business impact, enabling you to meet and exceed service level agreements.

Unicenter MP provides an integration point with SPECTRUM. Unicenter MP uses the Unicenter NSM integration with SPECTRUM to provide access to SPECTRUM objects and interfaces. From Unicenter MP, you can do the following:

- Create a connection to a SPECTRUM server
- Access the SPECTRUM OneClick interface
- Access the SPECTRUM Business Services page
- Publish SPECTRUM portlets
- View SPECTRUM objects in scoreboards and the Portal Explorer
- View SPECTRUM events in the Event Console and alerts in the Alert Console
- Launch the SPECTRUM OneClick interface from the Portal Explorer

To work with SPECTRUM, you must define a connection to the WorldView server that has been integrated with SPECTRUM and contains SPECTRUM objects.

## Additional Component Integrations

Unicenter MP integrates with the following additional components or products without creating a new folder or new content in the knowledge tree. These integrations provide additional data and functionality or access to additional interfaces or features, mostly through the Portal Explorer or report and scoreboard menus:

### **Unicenter Configuration Manager**

Lets you launch the Unicenter Configuration Manager interface from the Portal Explorer to manage agent configurations.

### **Unicenter NSM Knowledge Base**

Provides access to articles in the CA NSM Knowledge Base that provide solutions to common issues and problems.

### **Systems Performance**

Lets you launch Systems Performance Reports from the Portal Explorer for agents or servers that contain vital performance information.

### **Active Directory Explorer**

Lets you launch the Active Directory Explorer from the Portal Explorer, which is an interface for managing objects in an Active Directory.

### **Wily Introscope**

Lets you launch the Introscope Agent Dashboard from the Portal Explorer when an Introscope Agent object is selected. This dashboard provides useful application management information for servers managed by Wily Introscope.

# Chapter 6: Monitoring Your Enterprise

---

This section contains the following topics:

[Using Agent Technology to Monitor Resources](#) (see page 221)

[Understanding Unicenter Remote Monitoring](#) (see page 222)

[Understanding Resource Monitoring](#) (see page 227)

[Understanding Systems Management](#) (see page 243)

[Understanding Configuration Manager](#) (see page 264)

## Using Agent Technology to Monitor Resources

To facilitate comprehensive and integrated network polling and administration, CA NSM uses Agent Technology to automate manager tasks and responses to events. Agent Technology monitors and reports the status of your resources and applications and lets you manage those resources (also called managed objects). The status of a resource is displayed within the Management Command Center and on the WorldView 2D Map.

This chapter explains the basic functions of four aspects of monitoring your enterprise:

- **Unicenter Remote Monitoring**

A monitoring option that can be deployed instead of Agent Technology. Its manager (called the agent) runs on the Windows platform, and it can quickly discover and begin to monitor your resources. No installation is required on the remotely monitored machines.

- **Resource Monitoring**

The agent side of Agent Technology, which gets installed on remotely monitored devices. Specific agents monitor specific system resources, such as CICS resources, Active Directory Services resources, log files, UNIX/Linux system resources, Windows Management Instrumentation Resources, Windows System resources, and z/OS resources.

- **Systems Management**

The manager side of Agent Technology, which gets installed only on the management server, known as the Distributed State Machine (DSM). It controls the discovery of remote agents and monitored resources and maintains the status of these resources based on information received from the agents.

- **Configuration Manager**

Unicenter Configuration Manager provides an interface for the reporting and management of configuration information for remote and distributed Agent Technology and Event Management components.

## Understanding Unicenter Remote Monitoring

To determine when and how to use Remote Monitoring effectively, you need to understand both the advantages and disadvantages of deploying this non-intrusive monitoring technology.

### **Advantages of Remote Monitoring include the following:**

- **Faster deployment time**  
Because the agent does not have to be installed on each monitored resource, you can quickly get it up and running. For example, instead of installing an Agent Technology agent on 200 or more managed nodes, you install a Remote Monitoring Agent on only two machines that monitor those 200 machines.
- **Quick return on investment**  
The faster deployment means you can quickly see results of your investment.
- **Reduced support and management costs**  
Because a single agent machine can monitor hundreds of network resources, maintaining and supporting your monitoring environment requires less time and uses fewer human or hardware resources.
- **No performance interference**  
The agent, which also acts as the manager, runs on a separate machine and, therefore, will not disturb your production environment.

### **Disadvantages of Remote Monitoring include the following:**

- **Increased network traffic**  
Because remote agents must pull all data back to the agent machine for processing, it causes more network traffic than traditional Agent Technology.
- **Slightly less diverse data**  
The remote agent is unable to gather some of the data gathered by a traditional agent, because it does not reside on the monitored resource. Depending on the type of information that you need to monitor and, because, Remote Monitoring Agent may not be able to discover and monitor a specific resource that is important to your environment, you might decide to deploy the traditional Agent Technology.

**Note:** For more information about Remote Monitoring, see CA NSM - Remote Monitoring online help.

## Remote Monitoring Architecture

This topic explains the Remote Monitoring architecture, which will help you determine if Remote Monitoring Agent is best suited to monitor your environment.

Remote Monitoring consists of the following three major components:

- Administrative Interface

The Administrative Interface is the client application used to discover resources, configure resources, view status updates, and manage the metrics used to monitor each resource. This component runs on Windows computers only.

**Note 1:** Although the Administrative Interface provides the graphical user interface (GUI) used to discover resources, the agent actually does the work to probe the network and discover resources. Therefore, the Administrative Interface is not required to have administrator privileges to access the monitored resources.

**Note 2:** You must have the appropriate privileges to discover resources. Before you start the discovery process, ensure that the agent computer has the appropriate administrator privileges to gain access to your network resources or that an administrative account is associated with each node to be monitored.

- Agent

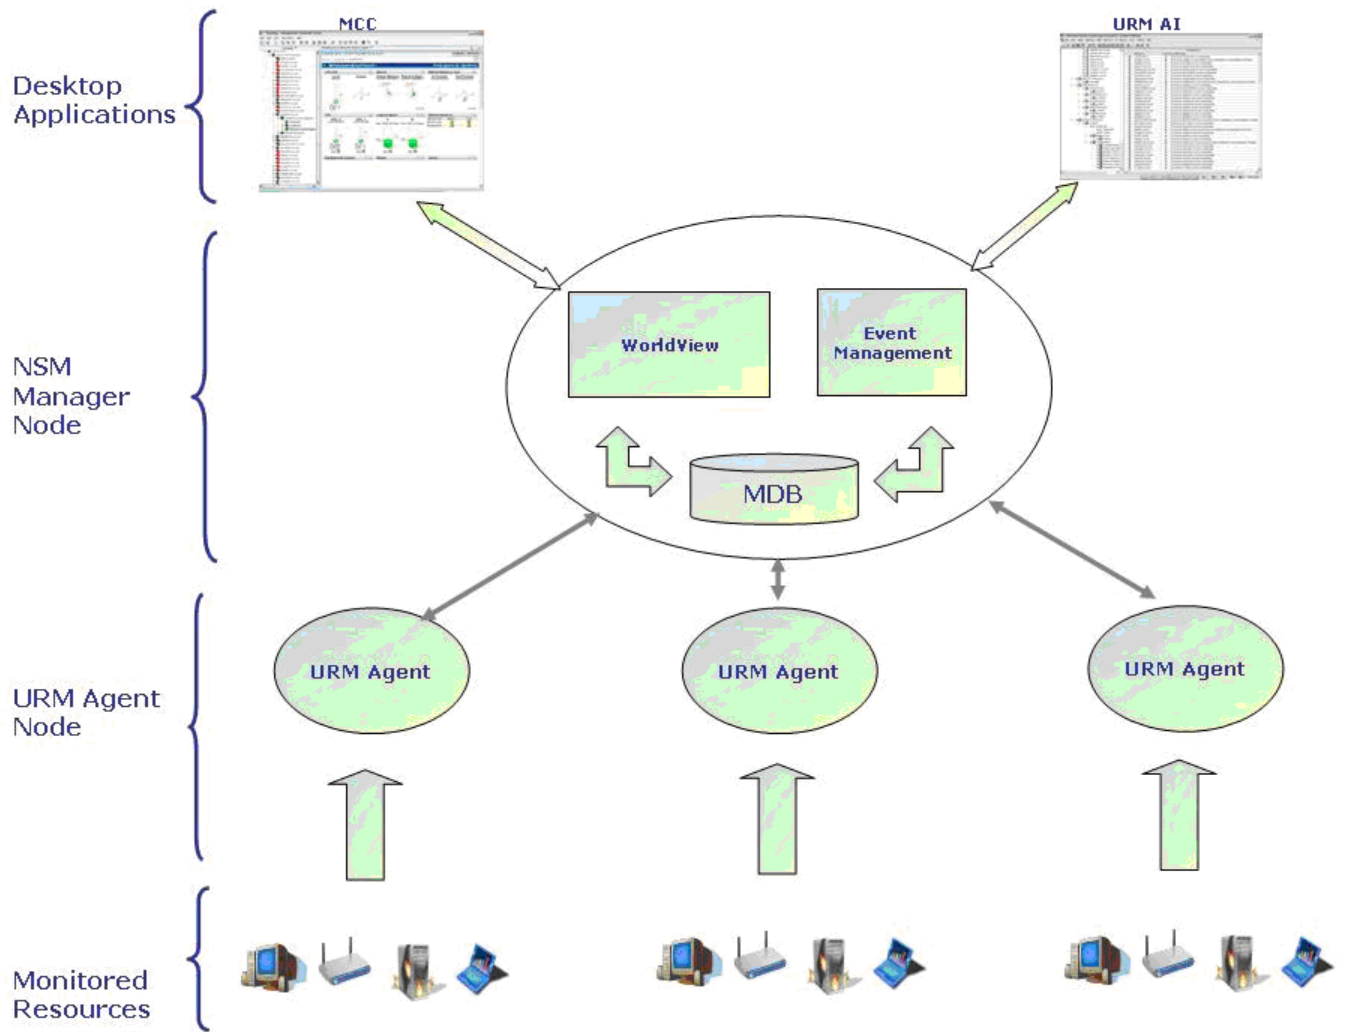
The Remote Monitoring Agent is responsible for polling all monitored resources and determining if an error has occurred. The agent can broadcast monitored resource status to a Unicenter Event Management Console, WorldView repository, and to any number of Remote Monitoring Administrative Interfaces.

The Agent runs on Windows computers only. Although it can reside on the same computer as the Administrative Interface, you can also install it on a separate Windows computer and access it from a remote Administrative Interface.

- Data Store

The data store contains all configuration information used to determine the current status of a resource. Data gathered from the most current poll is compared to the values stored in the data store, and any changes to a resource's state are communicated back to the agent. The data store is installed on the same computer as the Agent.

The following diagram illustrates how these components work together:





## Resource Types You Can Monitor

Remote Monitoring lets you monitor multiple platforms and resource types throughout your network. The following table lists all operating systems, the versions currently supported, and the type of information you can monitor for each:

<b>Operating System</b>	<b>Versions</b>	<b>Information Types Monitored</b>
Windows	<ul style="list-style-type: none"> <li>■ 2000 Professional, Server, Advanced Server, Datacenter (Intel x86)</li> </ul>	Event logs Services
	<ul style="list-style-type: none"> <li>■ 2003 Standard Server, Datacenter, Enterprise Server, Small Business Server (Intel x86, AMD-64, EM64-T, IA-64)</li> </ul>	System Metrics Detailed Metrics Registry Keys
	<ul style="list-style-type: none"> <li>■ 2003 R2 Standard, Enterprise, Datacenter (Intel x86, AMD-64, EM64-T, IA-64)</li> </ul>	
	<ul style="list-style-type: none"> <li>■ XP Professional (Intel x86, AMD-64, EM64-T)</li> </ul>	
	<ul style="list-style-type: none"> <li>■ Windows Vista Business, Enterprise, Ultimate (Intel x86, AMD-64, EM64-T, IA-64)</li> </ul>	
	<ul style="list-style-type: none"> <li>■ Windows Server 2008 (Intel x86, AMD-64, EM64-T, IA-64)</li> </ul>	
AIX	<ul style="list-style-type: none"> <li>■ 5.2 (POWER)</li> </ul>	System Metrics
	<ul style="list-style-type: none"> <li>■ 5.3 (POWER)</li> </ul>	Detailed Metrics
FreeBSD	<ul style="list-style-type: none"> <li>■ 6.2 (Intel x86)</li> </ul>	System Metrics Detailed Metrics
HP-UX	<ul style="list-style-type: none"> <li>■ 11iv1 (PA-Risc-64)</li> </ul>	System Metrics
	<ul style="list-style-type: none"> <li>■ 11.23 (PA-Risc-64, IA-64)</li> </ul>	Detailed Metrics
	<ul style="list-style-type: none"> <li>■ 11.31 (PA-Risc-64, IA-64)</li> </ul>	
Linux	<ul style="list-style-type: none"> <li>■ Red Hat 4.0 (Intel x86, AMD-64, EM64-T, IA-64, S/390)</li> </ul>	System Metrics Detailed Metrics
	<ul style="list-style-type: none"> <li>■ Red Hat 5.0 (Intel x86, AMD-64, EM64-T, IA-64, S/390)</li> </ul>	
	<ul style="list-style-type: none"> <li>■ SLES 9 (Intel x86, AMD-64, EM64-T, IA-64, S/390)</li> </ul>	

Operating System	Versions	Information Types Monitored
	<ul style="list-style-type: none"> <li>■ SLES 10 (Intel x86, AMD-64, EM64-T, IA-64, S/390)</li> </ul>	
Mac OS X	<ul style="list-style-type: none"> <li>■ 10.2 (PPC)</li> <li>■ 10.3 (PPC)</li> <li>■ 10.4 (Intel, PPC)</li> <li>■ 10.5 (Intel, PPC)</li> </ul>	<p>System Metrics</p> <p>Detailed Metrics</p>
Solaris	<ul style="list-style-type: none"> <li>■ 8 (UltraSPARC)</li> <li>■ 9 (UltraSPARC)</li> <li>■ 10(UltraSPARC, Intel x86, AMD-64, EM64-T)</li> </ul>	<p>System Metrics</p> <p>Detailed Metrics</p>
Tru64	<ul style="list-style-type: none"> <li>■ 5.1b (Alpha)</li> </ul>	<p>System Metrics</p> <p>Detailed Metrics</p>

In addition to monitoring these platforms, Remote Monitoring provides IP resource monitoring. This type of monitoring lets you gather the following information:

**State**

Indicates whether the system is responding.

**Response time**

Determines whether the response time is reasonable.

**State of selected ports**

Issues an alarm based on a state change, such as a port that is responding when it should be turned. off (not responding).

## Securing Access to Remote Monitoring

By default, all users are given full access to the Remote Monitoring features upon opening the application, and no login is required. However, you may have users who do not need to make configuration changes, but only need to monitor the status of your resources using the Remote Monitoring Administrative Interface. In this case, you can implement a role-based security scheme so that only administrators can access and change your monitoring configurations.

This role-based security access is an optional feature that provides the following two levels of security:

### **Administrator**

Provides full access to the application.

### **User**

Limits access to viewing the resource status information.

To implement this type of security, define one or more administrator accounts. Defining an administrator account puts the role-based security scheme into effect, and this security stays in effect as long as at least one administrator account is defined.

When this security is in effect, the default role is User. This means that upon opening the application, all configuration editing features are disabled. To gain administrative rights to the application, administrators must explicitly log in to the application, using the account you have defined. Upon successful login, the administrator is given full access to the application.

## Understanding Resource Monitoring

To determine when and how to use Resource Monitoring effectively, you need to understand what this monitoring technology can do for you.

### Basic Concepts

After startup, the system agents immediately start monitoring the system resources based on a predefined configuration. Lists of available (auto-discovered) system resources let you easily customize your agents during runtime to meet the specific requirements of your system.

An agent monitors system resources on the base of watchers. A watcher is the term used for any instance of a monitored resource that has been added into the agent's configuration. The agent evaluates the status of a specific resource according to the assigned watcher configuration.

To prevent losing a change in its configuration, for example, as a result of a power failure, the agent writes back its configuration data periodically. The duration of this period can be specified with the start command of the agent.

Some of the system agents support Auto Discovery. For some specific resource groups the corresponding agent adds watchers into its configuration automatically by applying filter conditions to the available lists. The agent uses the default values from the MIB to specify the properties of these watchers.

## General Functions

Most of the system agents support the general functions listed in the following sections. The descriptions in this section provide a brief overview. For further details, procedures, and examples, see the corresponding references.

### Auto Watchers and Available Lists

At startup the agent automatically discovers the system for monitored resources, but it depends on the type of the resource, whether the agent automatically creates a watcher for it or not. If a resource type appears in the form of only a few instances that shall always be monitored, it may be suitable for customers that corresponding watchers are automatically created (for example: CPU, Network Interfaces).

However, if a resource type appears in the form of many instances, for example file systems on UNIX servers; you may want to specify a particular subset of these instances that shall be monitored by the agent. For this case the agent does not create watchers automatically, but creates a list of the available objects (instances) of a resource type that can potentially be monitored.

Based on filter conditions of the available list you can specify a set of instances that you want to monitor and define an auto watcher for this set. Then, the auto watcher automatically creates individual watchers for those instances that match the filter condition. For example, you can specify a filter condition for the mount devices of the file systems and create an auto watcher for swap file systems only. Such an auto watcher creates individual watchers for each available swap file system on that server.

For monitoring files and processes the agent provides one-to-many watchers instead of auto watchers to monitor a specific set of instances by a single watcher. If the status of this set changes to warning or critical, the agent creates a culprit list that contains all monitored instances that caused the status change.

For example, you can specify a filter condition for the process path to monitor all processes that belong to `c:\Windows\system32` by a single watcher. In the case of a Down status the agent creates a list of items (process-ID:utilization value), which identifies the processes that caused this status. The sort order and length of this list depends on the severity of the violation, for example:  
408:222|409:333|475:444

## Call-Back Mechanism

The call-back mechanism of system agents enables you to assign an automated task or action to a particular event within the agent layer of the CA NSM architecture. This assignment is accomplished by means of a call-back reference which can be set up for each functional area of the agent, such as one call-back reference for CPU, one call-back reference for logical volumes, one call-back reference for files, and so on.

These call-back references can only be defined in an agent's call-back configuration file (for example: `caiUxsA2.cbc`) that can be secured by access rights. This configuration file is stored in the `Install_Path/SharedComponents/ccs/atech/agents/config/cbc` directory. It contains an entry for each call-back reference, and associates with this reference the full path and name of the script or application to run. Additionally, parameter information can be passed to the script or application, as well as a user ID that should be used to execute the script or application.

The advantage of using this additional level of indirection or call-back reference is that the name of this reference can be safely shown in the MIB without causing any security exposure, because the actual path and name of the call-back script or application is hidden within a secured file. This reference also enables you to remotely check in a secure way if a call-back reference has been configured for the respective monitored area.

**Note:** In the MIB the call-back reference name is defined as read-only. Therefore it cannot be set or modified by Agent View or the MIB Browser. The reference name can only be configured through a definition in a configuration set.

To provide improved functionality, you can specify that the agent will pass a set of predefined or user-defined parameters to the call-back script or application upon instigation. These predefined parameters will contain the following information:

- New watcher state (for example: OK, Warning, Critical)
- Type of element being watched (for example: FSys)
- Instance name of element being watched (for example: /var)
- Name of the monitored resource property that caused this status change (for example: Space, Inodes, Mount)
- Other miscellaneous var-bind information sent with the trap (for example: file system space and warning/critical thresholds)

By passing these parameters to the call-back script or application, it will enable you to build powerful scripts. These scripts can perform different actions depending on the state of the monitored resource.

## Cluster Awareness

Basically, support of monitoring clusters with CA NSM system agents is based on the CA High Availability Service (HAS). HAS is a set of extensions to Unicenter which enables Unicenter components to operate within a cluster environment, to function as highly available components, and to failover between cluster nodes gracefully. The system agents (caiUxsA2, caiWinA3, caiLogA2) use CA HAS and are cluster aware. This means even though those agents are running multiple times within the cluster (on each physical cluster node) only one agent monitors a shared cluster resource such as a shared disk.

No specific configuration is required for using these agents in a cluster, except for monitoring processes. The appropriate name of the cluster resource group (cluster service) must be specified when creating a process watcher.

**Note:** For more information, see the section Cluster Awareness and the appendix "High Availability Service" in the *Inside Systems Monitoring* guide, and the appendix "Making Components Cluster Aware and Highly Available" in the *Implementation Guide*.

## Configuring Resource Auto Discovery

Configurable resource auto discovery eases implementation phases, reduces the need for manual configuration, and discovers new resources dynamically, as they become available. An additional configuration group filter attribute serves as the criteria for an automatic resource detection and watcher creation mechanism.

## Editing Watchers

All the watchers of the system agents are editable. No watchers have to be removed and then re-added. If attributes of a watcher (for example, thresholds) are modified, the status of the watcher will be re-evaluated based on the current poll values. Therefore, modifying a watcher does not invoke polling.

## Evaluation Policy

For analog metrics of one-to-many watchers there are several possibilities to calculate the metric value. An evaluation policy makes this evaluation watcher-specific. If the result violates the monitoring conditions, a culprit list is determined. The form of the culprit list depends on the evaluation policy setting and different kinds of thresholds (rising/declining) or minimum/maximum ranges.

The supported evaluation policies are: sum, max, min, average, and individual.

## Generic Resources Monitoring

The UNIX System Agent and the Windows System Agent provide the generic resource monitoring concept that lets you extend the monitoring capabilities of Hardware monitoring and Programmable Resources monitoring by using external scripts or programs. These scripts must be “registered” in the Generic.ini file and have to provide a special output format for the evaluated data.

## History Group

The History Table lists the last  $n$  enterprise-specific status traps the agent raised. The value of  $n$  is a configurable attribute in the history group (<xyz>HistoryMaxEntries). Setting this value to 0 causes the agent not to store any trap history.

The trap history collection can be switched on and off on a per resource group basis. This feature is especially useful, if toggling watchers cause the trap history table to be filled again and again.

## Independent Warning and Critical Thresholds

The system agents allow warning and critical thresholds to be set independently for all relevant functional areas.

## Loss and Existence

For the most resource groups the system agents offer a status, which reports the loss or the existence of the resource from the watcher's point of view. The watcher reports a resource as lost or nonexistent, if it is unable to access the resource.

Beside the physical loss of monitored system resources, a logical loss has to be considered. For example: print queues can be unavailable for various reasons. The UNIX System Agent implements configurable logical and physical loss status monitoring. The propagation and evaluation of detected resource outages can be fine-tuned on a per instance basis.

## Message and Action Records

For many system agents the CA NSM r11.2 DVD ships files that contain definitions of all possible Event Message records as well as Action records. This considerably simplifies the creation of customer specific evaluations for the NSM event console.

Furthermore the CA NSM AEC component provides predefined correlation rules for the CA NSM r11.2 system agents.

## Minimum and Maximum Metrics

Minimum/Maximum metrics are binary metrics. They are used to monitor resources which have quantity characteristics that should stay within a specific interval. The agent provides two forms of minimum and a maximum metrics:

### Standard

This type provides a minimum and a maximum threshold (monitoring condition) and a monitoring level to determine the status of the resource. Detected resource values, which are greater than the minimum threshold and less than the maximum threshold, or which are equal to the minimum or maximum threshold, define the Up status for this metric. All other values define the down status.

### Extended

This type provides a minimum and a maximum range which are monitored through critical and warning thresholds leading to effectively four threshold borders:

`CritMin <= WarnMin <= WarnMax <= CritMax`

The logic of the metric can be changed by using additional policies, for example, the evaluation policy.



## Modification Policy

Files and directories can be monitored for being modified or unmodified. In both cases the dates of the corresponding files are used, that is, the file or files addressed by a file watcher or the entries in a directory including the directory itself (.) and all subentries if the recursive option is set.

## Overall Status of Each Functional Area

The system agents enable the Agent View (Abrowser) to propagate the most severe state of resources reported on the resource type specific windows to the Status Summary window. The Status Summary window summarizes the status of all monitored resources. It also displays the total number of monitored resources for each object type and the overall status according to the agent.

## Overloading Thresholds

In most cases, you define thresholds as percentages, but sometimes it is useful to define absolute values instead. Percentages are suitable where a high degree of resolution is not required. Additionally, they can provide generic values across many machines. Absolute values enable a far higher resolution. The overloaded thresholds concept lets you configure thresholds with the following scales:

- **Absolute used values**  
An example of this is defining the absolute number of MB that can be used on a logical volume before a state change occurs.
- **Percentage used values**  
This type of overload is indicated by appending a percent sign (%) to the threshold value. An example of this is the percentage of total logical volume space that can be used before the state change occurs.
- **Absolute free values**  
This type of overload is indicated by appending an F symbol to the threshold value. An example of this is defining the absolute number of bytes that should be left unused on a logical volume.

The agent will always convert the overloaded value entered by the client into an absolute used value and store this value in the MIB. This value is used for validation and status checks. The overloading must be the same for warning and critical thresholds. Not all kind of overloading is possible by all thresholds. For details see the MIB description.

Through MIB Browser, the manner in which the client distinguishes the type of overload is by appending the percent (%) sign or F symbol to the value. In Agent View, this translation is performed dynamically, using slider widgets and graphical controls.

### Periodic Configuration Write-Back

The system agents perform periodic configuration stores. To minimize overhead, an appropriate concept ensures that only configuration information that has changed since the last store operation is written back. If the system is being closed down, only recent configuration changes need to be stored, rather than the entire configuration.

### Poll Method

For each resource group the agent provides a method, which lets you disable the polling of any metric for that group completely. You can allow polling only triggered by the poll interval or allow polling also by a query. This property can be used to save performance in the agent.

### Resource Monitoring at an Instance Level

The system agents allow individual object instances to be monitored for all relevant functional areas.

### Resource Selection Capabilities

The system agents simplify the definition of new watchers by implementing a selection or available list from which the administrator can choose the specific resource they wish to monitor. The list will be generated, on demand, as per user-defined filter criteria.

### Status Deltas

For resources whose growth can consume finite resources on the machine (such as data files, and so forth), the concept of delta monitoring has been employed where feasible. This allows the agent to record the difference between the size of the resource during the last polling interval, and the size of the resource returned by the current poll. If this difference exceeds a client-defined threshold, an alert is issued. As a monitored object such as a file can contract as well as expand, it is also possible to calculate a negative value for a delta. The delta reported by the agent is always a positive or negative integer that simply reflects either the factor of growth or contraction of the resource. In the case of overloading the delta value may appear as a decimal value, for example: 99.86%.

To allow you greater flexibility when configuring the delta watchers, a type of overloading is implemented. This allows you to specify a threshold for growth, shrinkage or change in both directions. In addition to this it is possible to use the percentage type of overloading as well. You can define thresholds in the following formats:

- n- absolute shrinkage
- n+ absolute growth
- n absolute change in both directions
- n%- percentage shrinkage
- n%+ percentage growth
- n% percentage change in both directions

The threshold will always be entered as a positive value even if it is used to threshold against shrinkage. The actual delta value stored in the MIB is a positive or negative value to indicate the change as growth or shrinkage.

## Status Lags

To provide meaningful monitoring for resources that can peak for a very short period without a problem occurring, the agent can be configured to check for several threshold breaches before the state changes. This is configured by lag attributes. The lag specifies the number of consecutive threshold (b)reaches on which state changes. If the lag is set to one then the status behaves as if there is no lag. If the lag is set to two then the threshold needs to be (b)reached twice in a row to change the state.

The agent offers an aggregate lag attribute for all resources having an aggregate status. This lag defines the number of consecutive poll intervals on which any status of the monitored resource is not in the OK or Up state, before the aggregate status changes.

## SNMPv3 Support

SNMPv3 support is encapsulated in `aws_admin`. CA NSM r11, r11.1, and r11.2 system agents support SNMPv1 or SNMPv3, depending on an `aws_admin` configuration option.

## Traps with Total Values

The warning and critical values in the traps are absolute values even if you have percentage thresholds defined. Without a total value you are unable to judge the scale. For this reason the total value is added to the status and info traps.

## Watcher

An agent monitors IT resources on the base of watchers. A watcher is the term used for any instance of a monitored resource that has been added into the agent's configuration. The agent evaluates the status of a specific resource according to the assigned watcher configuration.

Usually a watcher consists of a set of metrics which are used to compare the detected values of monitored resources with monitoring conditions by considering settable monitoring levels. The result of this comparison is the status of the monitored resource according to the metric settings. The status of the watcher is the worst case aggregate of all associated resource statuses. If the aggregate status of a watcher changes, an info-trap can be sent to the manager. The info-trap contains information about the monitored resource that caused the status change.

Two basic watcher types can be distinguished:

### **One-to-one watcher**

A watcher is mapped to a single resource that shall be monitored. Characteristics of the monitored resource are evaluated by appropriate metrics. For example, a file system is monitored by a single watcher and different metrics are used to detect the status of file system characteristics such as size.

### **One-to-many watcher**

A watcher is mapped to a set of resources (instances) that shall be monitored. Common characteristics of these instances are evaluated by appropriate metrics. Unlike the one-to-one watcher a culprit list is provided to identify those instances that cause a status change of the watcher. Additionally, an evaluation policy defines for one-to-many watchers, how metric values, statuses total values and culprit lists of monitored instances are calculated. For example, processes or files can be monitored by one-to-many watchers.

## Monitoring System Resources

This section describes the resources that can be monitored by system agents.

## Active Directory Resources

Active Directory Management provides an enterprise-wide view of your Active Directory environment and supports the Active Directory Knowledge Base.

The Active Directory Explorer (ADE) is part of Active Directory Management. It is the main user interface for monitoring the Active Directory environment. ADE provides an instant view of the aggregated states of your forests, domains, sites, domain controllers, site links, and subnets. It lets you drill down into any of these components, providing a highly detailed enterprise and component-level view of your Active Directory environment's behavior.

Active Directory Management consists of the following components:

- Active Directory Enterprise Management Service (ADEM) installed on the NSM manager system
- Active Directory agent installed on each monitored domain controller

## Active Directory Enterprise Manager

The Active Directory Enterprise Manager creates and maintains all Active Directory objects according to the following enterprise-wide Active Directory resources it monitors:

- Forests
- Domains
- Sites
- Site-links
- Subnets
- Domain Controllers

The Active Directory Enterprise Manager queries the Active Directory for information about these resources. Additionally, it polls the Active Directory Agents on all monitored domain controllers in all forests for domain controller-specific metrics and statuses.

The Active Directory Enterprise Manager analyzes the information it gathers from enterprise-wide Active Directory resources and displays it through Active Directory Explorer. Based on this information it provides an enterprise-wide view of your Active Directory resources.

## Active Directory Agent

The Active Directory Services Agent can run on any Windows 2000 server platform or higher if the system is a member of an Active Directory domain. However, the complete monitoring capabilities offered by the agent are only available on a system that is defined as an Active Directory domain controller and as a DNS server. On other systems within an Active Directory Services domain, information on disk space resources, on extended resources, and on one or more performance resources is not available.

The Active Directory Agent monitors the following critical areas:

- Domain Controllers (pertinent to all servers)
- Disk Space (pertinent to domain controller)
- Active Directory Services Events (pertinent to all servers)
- Active Directory Services Performance (pertinent to domain controllers)
- File Replication Service Events (pertinent to domain controllers and FRS servers) and Distributed File System Replication Events (pertinent to Windows 2008 domain controllers and Windows 2008 Distributed File System Replication servers)
- File Replication Service Performance (pertinent to domain controllers and FRS servers)
- Domain Name Service Events (pertinent to all servers)
- Domain Name Service Performance (pertinent to DNS servers)
- Extended Resource Monitoring (pertinent on domain controllers)

**Note:** When you install the agent on a member server, only the subset of the previously listed resources pertinent to all servers is available for monitoring.

## CICS Resources

The CICS Agent provides status, event, and configuration information about a CICS region and the transactions that are executed within it. The agent enables you to monitor the key resources, such as DSA and memory, of your CICS regions. The agent can monitor individual resources as well as the "health" of an entire region, allowing you to quickly determine the cause of a problem.

The CICS Agent puts you in control by allowing you to determine the warning and critical thresholds for each monitored resource. The agent monitors these resources and, whenever a user-defined threshold is exceeded, sends an SNMP trap.

The CICS Agent runs in IPv6 environments.

# Chapter 7: Host Resources MIB

---

The Host Resources MIB (RFC 2790), defined by the Internet Engineering Task Force (IETF), provides agent-less monitoring for generic host systems. Host computers are independent of the operating system, network services, or any software application. The Host Resources MIB provides the standard for monitoring a range of parameters like CPU, memory, disk, software installed, and processes running.

The Host Resources MIB is available out of the box, and thus provides basic monitoring capability immediately with no need for any agents. It can also run along with the CA NSM agents. The Host Resources MIB is useful for providing basic information on less critical systems. For critical systems, you may want to use the other CA NSM agents because they provide more information and send traps to the Event console. The Host Resources MIB does not support traps but sends status information through poll response. Therefore, any system or resource changes may take longer to be detected than with traditional CA NSM agents.

The MIB monitors the following resources:

## **System**

- system reboot
- number of users logged in
- number of processes running

## **CPU**

- device status
- CPU load

## **Memory**

- physical memory status
- virtual memory status

## **Disk file system**

- usages

## **Process**

- services running or not running (single process instance)

You can view the Host Resources MIB on the Management Command Center, Agent View (abrowser), Node View, and MIB Browser.

## Log Agent Resources

The Log Agent monitors log data and files by using pattern matching based on regular expressions.

The Log Agent can be configured to monitor ASCII log files and to facilitate the detection of faults in applications running under the monitored system. Besides the monitoring of single log files the agent offers the monitoring of all files in a subdirectory, and the usage of wildcards in file names. Furthermore, the Log Agent can monitor the Windows Event Log, the UNIX Console, files with single lines and files with ASCII control characters.

The Log Agent has only limited support for binary files. It replaces each null character (0x00) by blank before matching against the pattern. These blanks are shown in any trap resulting from a match.

**Note:** The Log Agent handles Unicode encoded log files as binary files.

## Script Agent Resources

The Script Agent provides a rapid way to extend the monitoring capabilities of CA NSM. It leverages or extends existing business logic to manage areas of production systems that are not covered by other CA NSM monitoring components.

The Script Agent's purpose is to run any number of scripts based on a user-defined schedule and to derive a status from the output or return code of each script. The agent monitors these resources and sends an SNMP trap when a user-defined threshold is exceeded.

The Script Agent uses a configuration set that is loaded at agent startup to specify script instances. Script instances define the path to the script, the execution interval, and the return code or the regular expressions that derive a status of Normal, Warning, or Critical. You can specify any number of script instances in a configuration set. The agent reports the status derived for each script instance to the DSM and ultimately to WorldView.

The Script Agent consists of an Agent Technology agent, DSM policy for discovering and managing the agent, .wvc policy for defining the agent classes and states in WorldView, and an Agent View for you to visualize the status of active script instances. The agent also provides an extension to the Unicenter Configuration Manager so you can centrally deploy and manage the agent policy and scripts.

The agent runs under the control of the local Agent Technology Services and implements the policy defined in the configuration set it loads at startup. It maintains the status of the watcher by creating a watcher for each defined script instance and then parsing the script return code or output and comparing that return code or output against the status policy.



The DSM policy discovers the agent and script instances and uses traps or polls to determine the current state of each instance. The .wvc file populates the scriptAgt class to WorldView. Because the scripts often represent elements of key business logic that are being monitored for health and availability, you can include the class in Business Process Views.

Windows, Linux, and most current UNIX platforms support the Script Agent.

## SystemEDGE Agent

The CA SystemEDGE is a lightweight SNMP (Simple Network Management Protocol) agent capable of retrieving, monitoring, and publishing operating system metrics on a wide variety of platforms. It lets remote management systems access important information about the system's configuration, status, performance, users, processes, file systems and much more. The agent includes intelligent self-monitoring capabilities that enable reporting and managing of exceptions and that eliminate the need for excessive polling.

CA NSM r11.2 supports CA SystemEDGE 4.3. Starting with r11.2 service pack 1, CA NSM supports CA SystemEDGE 4.3 and 5.0.

**Note:** For more information, see *Inside Systems Monitoring* in the CA NSM product.

## UNIX/Linux System Resources

The UNIX System Agent monitors the following resources:

- System
- OS Resource Parameters
- CPU
- Load Averages
- Memory
- Swap
- Quotas
- Directories
- File Systems
- Files
- Disks

- Processes
- Print Queues
- Network Interfaces
- Shared Memory
- Semaphores
- Message Queues
- Hardware/Programmable Watcher

### Windows Management Instrumentation Resources

Windows Management Instrumentation (WMI) is an implementation of the Web-Based-Management-Initiative. This movement provides an object-oriented structure for communication between web-based applications, similar to the structure provided by SNMP. WMI fits easily into the Agent Technology architecture, using classes, objects, and properties when describing object instances.

WMI uses namespaces to specify an object instance. A namespace is a string that uniquely identifies a group of instances and classes. It reflects a logical grouping containing specific classes and instances. For example, an operating system process object could have a namespace as follows:

```
root\CIMV2
```

The WMI Agent lets you monitor the content of your Windows Management Instrumentation (WMI) to solve problems in real-time. You can specify what resources to monitor as well as WMI events of interest. By providing this level of customization, you can receive alerts on vital areas of your WMI content without having distracting alarms from non-critical resources.

The Windows Management Instrumentation Agent queries your monitored resources for data. This data is compared with thresholds or applied to policies to determine the statuses of the monitored resources. Based on this comparison the agent sends alerts to agent managers in your enterprise.

### Windows System Resources

The Windows System Agent monitors the following resources:

- CPU
- Memory
- Logical Volumes
- Mounts

- Dfs Links
- Quotas
- Directories
- Files
- Processes
- Services
- Jobs
- Sessions
- Printers
- Network Interfaces
- Registry Entries
- Hardware/Programmable Watcher

### **z/OS Resources**

The z/OS system agent enables you to monitor key resources of your z/OS system and provides status, event, and configuration information. The agent can monitor individual resources as well as the health of an entire system, allowing you to quickly determine the cause of a problem. The z/OS system agent also monitors UNIX System Services (USS) resources.

The z/OS system agent puts you in control by allowing you to determine the warning and critical thresholds for each monitored resource. The agent monitors these resources and, whenever a user-defined threshold is exceeded, sends an SNMP trap.

The z/OS Agent runs in IPv6 environments.

## **Understanding Systems Management**

The Agent Technology component of CA NSM lets you closely monitor the critical computing resources in your enterprise by configuring the managed nodes to watch and answer queries regarding their resources. Agent Technology uses agents running on remote devices throughout the enterprise to collect data, such as information on file systems, memory utilization, and database performance. This data is distributed to one or more management servers that can interpret the data and take action based on custom configurations.

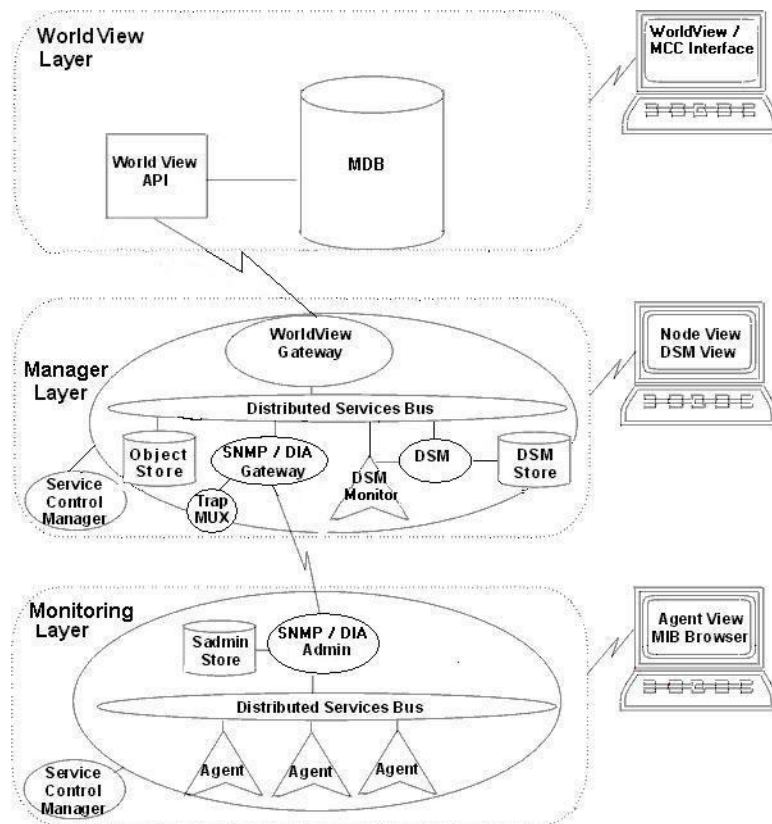
The Systems Management portion of this chapter explains briefly how each layer of the architecture works, but focuses on the management layer. It introduces you to the tools that allow you to manage, configure, and visualize your computing enterprise.

## Understanding the Architecture

CA NSM delivers an architecture consisting of the following three layers:

- Monitoring Layer--checks for abnormal behavior or inconsistencies
- Manager Layer--manages the enterprise
- WorldView--views the enterprise

This layered architecture delivers a powerful, distributed, and versatile management solution to accommodate large-scale, complex, and dynamic environments. Data collected on each node in the enterprise passes from the monitoring layer to the management layer to the WorldView layer, as shown in the following illustration.



### Communication Status (Monitoring Layer)

The Monitoring Layer collects data about your enterprise.

An agent is an application that supports network management. An agent typically resides on a managed software node, such as a Windows XP server, and provides information to a management application.

This information is interpreted according to a management protocol that is understood by both managers and agents. CA NSM agents use the following protocols:

**Communications Protocol**

User datagram protocol (UDP) of the transmission control protocol/Internet protocol (TCP/IP) suite.

**Network Management Protocol**

Simple network management protocol (SNMP) designed to run on top of TCP/IP.

Distributed Intelligent Architecture (DIA) designed to run on top of TCP/IP.

Both agent and management applications can view the collection of data items for the managed resource. This collection is defined by the *Management Information Base* (MIB). Each MIB describes attributes that represent aspects of a managed resource. The network management platform accesses MIB data using SNMP.

You can view the current statistics about monitored resources from various interfaces, such as MIB Browser, Agent View browser, and through the Management Command Center.

Every agent must be associated with at least one DSM. Through configuration, you can determine which machines in your enterprise report to a DSM. Each DSM can communicate with only one MDB, but a single MDB can accept information from multiple DSMs.

## Managed Objects

Each resource that an agent monitors is called a managed object, and every managed object has a state.

A managed object can represent a physical device, such as a printer or a router, or it can represent an abstraction, such as the combination of hardware and software components that constitute a network connection between two nodes. A managed object can be monitored and, in some cases, controlled with the use of one or more management applications.

CA NSM groups managed objects into classes. A class is a group of managed objects that share a common definition, and therefore, share common structure and behavior. By changing the behavior of a class, you can change the behavior of the managed objects that belong to that class.

The definitions for each agent class are in their individual policy files. For more information about policy files and their syntax, see the guide *Inside Systems Management*. For more information about a specific class of managed object, see the individual guide, such as *Inside Systems Monitoring*.

## States

A state is one of a set of predefined possibilities for the condition (for example, up, down, or unknown) of the managed object.

A change in state appears on the WorldView 2D Map or in Management Command Center, as a change in icon color. You can drill down through the network topology, to the machine name or IP address, to Unispace, and to the agent that communicated the state change. From the pop-up menu associated with that agent, you can view the current state of the agent according to the DSM (from Node View) or according to the agent (from Agent View).

Each layer of Unicenter interprets the state of a managed object differently, placing an emphasis on different aspects of the object. For example, a Windows 2003 server may be taken off-line for repairs. The Windows 2003 system agent stops sending information to its DSM. The state of the agent according to the DSM is Absent, indicating that no information is being relayed. However, the state of the agent according to WorldView is Down, indicating that the server is inaccessible.

## How the Threshold Breach Process Works

Knowing how the Monitoring Layer reacts to a breach in a predefined limit (threshold) will help you understand how communication moves throughout your enterprise and how information about monitored resources is conveyed. It will also help you respond to and resolve any problems when they occur.

The following steps outline what occurs architecturally when a threshold breach on device is detected:

1. A Unicenter agent identifies that a threshold has been crossed for a resource that is being monitored, such as CPU or memory. The agent passes this information to either the SNMP Administrator or the DIA Administrator by way of the Distributed Services Bus (DSB).
2. The SNMP Administrator takes the information from the agent, encodes an SNMP Trap Protocol Data Unit (PDU), and sends it to the agent's assigned DSM server, or the DIA Administrator encrypts the data, and sends it to the agent's assigned DSM server.
3. The SNMP or DIA Gateway receives the Trap, decodes it, and sends it to the DSM component by way of the DSB.
4. The DSM component determines whether the alert represents a change in status for the resource and, if so, passes a status update to the WorldView Gateway by way of the DSB.
5. The WorldView Gateway then updates the status of the monitored resource in MDB.

## Managing the Enterprise (Manager Layer)

The manager layer provides control of the day-to-day operations of your IT infrastructure. The Manager Layer consists of the Distributed State Machine (DSM) and its supporting components. The DSM interprets data collected and makes that information available to applications such as WorldView and Node View. The DSM can run on Windows and UNIX/Linux platforms.

The Manager Layer includes the following components:

- Service Control Manager (awsservices)
- Distributed Services Bus (aws\_orb)
- SNMP/DIA Gateways (aws\_snmp) and (aws\_agtgate)
- Object Store (aws\_store)
- DSM Store and Distributed State Machine (aws\_dsm)
- DSM Monitor (dsmmonitor)
- WorldView Gateway (aws\_wvgate)

These components run as processes that can be started and stopped independently, or as a group, by the Agent Technology Service Control Manager (awsservices). To view the Service Control Manager, see the section Tools to Configure Managed Resources. A brief explanation of each component follows.

For more information about any of these components, see the guide *Inside Systems Management*.

### Service Control Manager (awsservices)

The Service Control Manager, or the Agent Technology Services Manager, displays the status of each DSM component. You can access the Agent Technology Services Manager from its own window. For more information, see Tools to Manage Resources.

### Distributed Services Bus

The Distributed Services Bus manages communication between all the other Agent Technology components. All information exchanged is placed on and retrieved from the Distributed Services Bus.

## SNMP / DIA Gateways

The SNMP Gateway encodes requests into SNMP Protocol Data Units (PDUs), sends them to SNMP agents in the management domain, and receives and decodes get response and trap PDUs from the agents.

DIA (Distributed Intelligent Architecture) is a proprietary, encrypted, high-speed facility. You can use the DIA Gateway enterprise-wide only if all agents are running NSM r11, r11.1, or r11.2. Otherwise, SNMP is required for all other agents. DIA is covered more fully in the earlier chapters of this guide.

## Trap MUX

The Trap multiplexer (MUX) allows multiple management applications to listen for traps on the same trap port. For example, Enterprise Management and the DSM both listen to port 162.

## Object Store

Agent Technology provides a mechanism for the persistent storage of objects, called the object store. The Object Store stores class data on managed objects. Managed object class definitions are loaded into Object Store from DAT files. The DSM uses the class definitions in Object Store to discover agents and managed objects on the appropriate nodes.

## Distributed State Machine (DSM)

The Distributed State Machine controls the discovery of agents and monitored resources in the enterprise and maintains the status of these resources based on information received from remote agents. The DSM also converts trap data, poll responses, and user actions into current object states. The DSM interprets data collected by agents and then makes that information available to applications running on Windows and UNIX/Linux platforms.

The DSM evaluates information received from managed nodes by way of the SNMP/DIA Gateway. The DSM determines what managed object the information pertains to and what the information is saying regarding that object. The DSM has logic for each managed object, which it uses to determine if there has been a change in state.

The DSM stores managed objects and their current states in memory. When the DSM is stopped and restarted, it obtains the last reported state for each previously monitored object and the instance-level properties for those objects.



## DSM Store

The DSM Store contains DSM managed objects that represent network nodes, agents, and the resources that are being monitored. These managed objects are created each time the DSM starts up during the DSM discovery process. The DSM uses the managed objects to maintain the current status of monitored resources in its domain.

Each DSM managed object has associated property values assigned based on the class definition that is present for that type of object in the Object Store. An object's current state is one property value maintained by the DSM.

## DSM Monitor

The DSM is self-managing to ensure that your resources are under constant surveillance and that the health and load of each DSM can easily be determined. DSM Monitor uses various data collection methods to effectively monitor the DSM process, its impact on CA NSM, and its impact on the performance of the server on which it runs. You can use the historical data collected to fine tune the DSM-managed enterprise by balancing the number of managed objects and classes across multiple DSMs.

## WorldView Gateway

The WorldView Gateway communicates with MDB through the WorldView Application Programming Interface (API) to get information about any nodes that WorldView has discovered. The WorldView Gateway then filters the list of discovered nodes based on the contents of the DSM Configuration IP Scoping Table, and forwards the appropriate list of managed objects to each DSM.

For example, if DSM1 is configured to manage only devices with IP addresses 172.16.0.0 to 172.52.255.255, then the filtered list provided to DSM1 by WorldView Gateway would include only the addresses of any devices discovered within that range.

The WorldView Gateway also filters the list of nodes based on the node class. This information comes from the Class Scoping Table, which contains a list of the node classes that the DSM should monitor. Another task of the WorldView Gateway is to pass state change information from the DSM to the MDB.

## How DSM Discovers Your Resources

Knowing how the Distributed State Machine (DSM) discovers your resources and determines their state will help you understand how CA NSM communicates throughout your enterprise and how it conveys information about monitored resources. It will also help you respond to and resolve any problems when they occur.

During its operation, the DSM moves through the following steps:

1. DSM obtains the list of discovered nodes from WorldView
2. The list of discovered nodes gets filtered to create the domain list.
3. The DSM discovers Managed Objects within its domain.
4. The DSM creates a managed object for each running agent and child object, stores this information in DSM Store, and registers with its agents.
5. The DSM determines the current state of each managed object in its domain and loads it into DSM Store.

### View the Enterprise (WorldView Level)

The WorldView layer represents your entire CA NSM enterprise. This layer contains the WorldView 2D Map and Management Command Center, through which you can view all of your monitored resources and their relationships, such as which monitored resources appear on which nodes.

The WorldView layer includes the following components:

- WorldView Interface
- MDB (the common database)
- WorldView Application Programming Interface (API)

For more information about WorldView, see the appropriate chapters in this guide.

### Tools to Configure Managed Resources

When WorldView detects a problem in your network, someone must locate the specific resource that is generating the alert and resolve the problem. CA NSM provides a number of interfaces that let you stop and start services, view monitored resources, and modify their thresholds or properties.

### Agent View

Agent View provides an interface for configuring an agent. Agent View contains several windows that reflect the different sets of monitored resources. Within each window, you can set configurable attributes by adding and deleting resources to monitor, setting warning and critical threshold values, and setting resource polling intervals.

You can perform similar tasks in agent dashboards. For more information about dashboards, see the chapter "Managing On-Demand."

To access the Agent View window from Node View, right-click the bar containing its name, then choose View Agent from the pop-up menu.

## DSM View

DSM View displays the managed objects for an individual DSM. DSM View lets you find objects and manage properties associated with managed object classes in a management domain. You can create new properties for a managed object, as well as modify the values of existing properties.

**Note:** You can also use the DSM Wizard to modify a selected subset of property values for all discovered instances of specific agent object classes. To access the DSM Wizard from a command prompt, enter: *dsmwiz*.

You can access DSM View using any of the following methods:

- From Node View menu bar, choose Edit, Find Object
- From Node View, click the Find Object button (microscope icon)
- From a command prompt, enter the command, *obrowser*  
(For more information on the syntax of the *obrowser* command, see the online *CA Reference*.)
- From Management Command Center, select DSM View from the left pane drop-down list

## Event Browser

The Event Browser provides detailed information about the events that have affected the status of an object and includes the following information:

- State changes
- Reason for state changes
- Warning messages sent
- Creation of an object
- Deletion of an object

With this information, you can determine patterns of activity. For example, the Event Browser shows that an object moves from a NORMAL to a CRITICAL state once every hour. You can use this pattern to determine the cause of the problem with that object.

Because the Event Browser lists events in real-time, the display window changes continuously. You can freeze the display of the event log list temporarily to examine the details of a particular event. You can also sort and filter Event Browser information.

Access the Event Browser using any of the following methods:

- Right-click an object in WorldView Classic 2D Map and choose Event Browser from the pop-up menu.
- Right-click an object in Node View and choose Event Browser from the pop-up menu.
- Right-click an object in Management Command Center and choose Viewers, Event View from the pop-up menu.
- From a command prompt, enter the command, `ebrowser`

## MIB Browser

MIB Browser lets you view the agent MIB in a tree-like hierarchy. If you are familiar with a particular agent, you may prefer to use MIB Browser to set configurable attributes, such as threshold values. MIB Browser shows an agent's entire MIB as it appears in the file, whereas Agent View provides a graphical representation of the MIB.

You can access the MIB Browser by using one of the following methods:

- Right-click an object in WorldView Classic 2D Map and choose MIB Browser from the pop-up menu.
- Right-click an object in Node View and choose MIB Browser from the pop-up menu.
- From a command prompt, enter the following command, `mibbrowse`

After the MIB Browser appears, you can log into any MIB by using the Open Connection menu (or click the telephone icon). If you are using SNMPv3, you can perform a secure login from the Open Connection dialog.

## Node View

Node View builds an object tree from the information provided by the DSM. Node View recognizes status change from the DSM by changing icon colors. Status propagates up through the Node View tree—that is, the most severe status reported by a child object propagates horizontally to the parent object.

From the Guidance Window at the bottom of Node View, you can see the real-time recording of session activity, initial statuses of managed objects in the tree, status changes, acknowledgements of status changes, the syntax of commands triggered when using the Node View menu, error information, and so forth.

**Note:** You can change DSM policy to affect the way the DSM states are displayed in Node View.

You can access Node View by using one of the following methods:

From WorldView 2D Map, right-click an agent object and select Node View from the pop-up menu.

From Management Command Center, right-click an agent object, and select Action - View Node from the pop-up menu.

From a command prompt, enter the command: `nodeview`. (For more information about the syntax of this command, see the online CA Reference.)

## Remote Ping

CA NSM lets you poll a remote host from another host. The Remote Ping interface lets you indicate the IP addresses of the source and destination machines and to establish retries and timeouts for the poll. In addition, you can view the activity of the Distributed Services Bus where the poll originates.

You can request a Remote Ping from the Event Console, the 2D Map, or the command line.

For more information on polling a remote host, see the Remote Ping online help.

You can access Remote Ping using one of the following methods:

- Click Start, Programs, CA, Unicenter, NSM, Agent Technology, Remote Ping.
- From a command prompt, enter the command, `rping`

The Remote Ping dialog appears.

## Repository Monitor

Repository Monitor lets you monitor the various agent classes that are listed in the MDB. You can view and delete a complete list of objects for a specified agent class. Use this tool if you have discovered classes in your enterprise that you know you will not want to monitor.

**Note:** Advanced users can also delete the class name from the central list of classes being monitored.

You can access the Repository Monitor using one of the following methods:

- Click Start, Programs, CA, Unicenter, NSM, Agent Technology, Repository Monitor.
- From a command prompt, enter the command, *agtrmon*.  
The Repository Monitor appears.

## Agent Technology Services Manager

The Service Control Manager, or Agent Technology Services Manager, shows the status of each DSM component. It starts and stops all Agent Technology components and agents on a node in the correct order. Once this component is running, all other DSM components can be installed, started, stopped, or uninstalled.

**Note:** This executable is the only Agent Technology process that is installed as a Windows service.

You can access Service Control Manager using one of the following methods:

- Click Start, Programs, CA, Unicenter, NSM, Agent Technology, Service Control Manager
- From a command prompt, enter the command, *atscm*.  
The Service Control Manager dialog appears.

For more information about how to start and stop awservices from the command line, see *awservices--Service Control Manager* in the online *CA Reference*.

## SNMP Administrator

The SNMP Administrator checks the community string and Internet Protocol (IP) address of get, get-next, and set requests to ensure that these requests come from authenticated management applications. This component forwards trap messages to the appropriate destinations. The SNMP Administrator also stores configuration sets and MIB definitions for each agent on the node in memory.

While the DSM is discovering the monitored systems in its domain, the DSM registers its own IP address with each system's SNMP Administrator. At that point, each monitored system knows which DSM to send traps to. Consequently, those remote monitored systems do not require individual configuration for trap destinations.

**Note:** For fail-over purposes, a monitored system should have more than one trap destination.

Access the SNMP Administrator by right-clicking the AWSadmin object in Node View and selecting View Agent from the pop-up menu.

The SNMP Administrator View - Summary dialog appears.

## Configuring Managed Nodes

This section describes the benefits of configuration sets, how to create and load a configuration set, and methods of distributing the configuration sets.

### Benefits of Configuration Sets

If you change the configuration of an agent during runtime through an application such as Agent View, those changes take effect immediately and remain in force until the attribute values are reset. If you stop and restart the agent with no configuration set in use, your runtime changes persist. However, if you start the agent with a configuration set read from the SNMP Administrator Store, changes made online during the agent's last execution are not restored. Therefore, using a configuration set for similar managed nodes ensures that the managed nodes are being monitored according to your defined policy.

When you implement a custom agent policy, you need to do so in a documented, controlled, consistent, repeatable, and efficient manner. The agent configuration set (configset) provides this capability. It allows you to export an agent policy to a text file. You can then distribute the policy in a manner that ensures the policy is uniform across the enterprise.

Configuration sets let you set the initial values for the agent's MIB attributes (including threshold values), call-back references, community definitions, and trap destinations, and load these values into the SNMP Administrator store. When a configuration set is specified during the startup of the agent, information about the resources being monitored is taken from the SNMP Administrator store.

## Defining a Configuration File Name

When an agent configset is initially created, it is saved to a file referred to as an agent configuration file. The recommended location to maintain agent configuration files is in the folder located at `\ccs\at\agents\config`. There is no fixed naming standard for these files; your organization can establish its own naming standard.

**Note:** Though not required, it is a good business practice to keep agent configuration files in the `\ccs\at\agents\config` directory or a similar central location.

You write a configset using a prescribed syntax, typically taking advantage of the `mkconfig` utility. Therefore, detailed knowledge of the syntax is not necessary. If you do need to create or significantly modify a configset, information about configset syntax can be found in the guide *Inside Systems Management*.

## Using a Configuration File

You can create agent configuration files using the `mkconfig` (`mkconfig.exe`) utility. `Mkconfig` retrieves an agent's current configuration information from its MIB and delivers its output in configset format to a specified configuration file or `STDOUT`. All configuration files created with the `mkconfig` utility have a single configset, which is assigned an internal name `bootstrap`.

An example of the `mkconfig` command follows:

```
mkconfig caiW2k0s@NSMSrvr5 > caiW2k0s.cfg
```

## Using Adaptive Configuration

Some important challenges with agents are the initial configuration and the ongoing configuration adjustment. The Adaptive Configuration service for the UNIX Operating System Agent, the Windows System Agent, and the Active Directory Services Agent supports these areas with the following services:

- Detection of resources that should be monitored
- Providing suitable thresholds for these resources
- Automatic configuration of the agent with an optimum monitoring policy

After startup the Adaptive Configuration service provides a predefined configuration. If the predefined configuration does not match your specific applications you can customize the service to meet your needs. You can influence the Adaptive Configuration service, for example, by specifying threshold policies or including or excluding specific resources.



By default the Adaptive Configuration service is installed along with the Active Directory Services Agent, the UNIX System Agent, and the Windows System Agent. The Log Agent partially supports the Adaptive Configuration service. The Adaptive Configuration service moves through the following modes:

- Self-Configuration Mode

Rapid and automatic configuration of an agent when it is first deployed to its target environment with no other form of a predefined configuration.

Duration: about 3 minutes

- Initial Self-Adaptation Mode

- Self-Adaptation Mode

Ongoing refinement and adjustment of an agent's existing configuration. In this mode of operation, the Adaptive Configuration process provides an ongoing learning and training exercise conducted over a number of weeks or months.

You can access Adaptive Configuration through the Unicenter Configuration Manager, which is described in a subsequent section.

See the guide *Inside Systems Monitoring* for more information about running the Adaptive Configuration service on a specific host.

## Loading a Configuration File

Agent configuration files are loaded into the sadmin store using the `ldconfig` (`ldconfig.exe`) utility. `ldconfig` can load a configset into a local or remote sadmin store. It can also remove a configset from a local or remote sadmin store.

The `-h <host>` parameter is used to specify the remote node if the `ldconfig` remote load capability is used. `ldconfig` connects to the remote DSM and loads the configset into sadmin store using the SNMP Administrator. The `agentctrl` command can be used to verify that the configset exists in the sadmin store.

An example of the `ldconfig` command is as follows:

```
ldconfig -h NSMSrvr5 caiW2k0s.cfg
```

## Distributing Configurations

You can use the `ldconfig -h <host>` parameter to distribute agent configsets, either individually or in batch mode. A better approach to applying a configset to many similar servers is to use a software delivery based solution. Configuration files are usually distributed to the folder `\ccs\at\agents\config` of specified managed nodes.

Central configuration is provided by the Unicenter Configuration Manager. From Unicenter Configuration Manager you can create, modify, and distribute agent configurations in your enterprise. Within Unicenter Configuration Manager, agent configsets and Adaptive Configuration profiles become agent profiles and are deployed to remote hosts within configuration bundles. You should also use Unicenter Configuration Manager to centrally distribute other files that configure your environment, such as `atservices.ini`, `atmanager.ini`, `aws_sadmin.cfg`, or `aws_sadminV3.cfg`.

For more information about Unicenter Configuration Manager, see the section [Understanding Configuration Manager](#).

## Configuring a DSM Environment

DSM Configuration is a tool that lets you visually configure your DSMs to manage their domain environment (their scope).

You want to be able to configure what each DSM is monitoring and how often the DSM communicates with your devices in order to proactively keep all your systems running smoothly with the least number of problems or downtime.

DSM settings that in previous versions could be modified only within files on each remote DSM are now accessible from the following two interfaces:

- Management Command Center under Tools—DSM Configuration
- DSM Wizard

## Understanding DSM Configuration

You can decrease the work required of your DSMs, and therefore make them work more efficiently, if you specify only the classes and IP addresses you know reside within your enterprise and if you fine tune the pollset values and the managed objects that appear in WorldView. By customizing the DSM settings, you can decrease the time it takes for a DSM to discover all the resources it should monitor and to reload policy when restarted.

Use the five DSM Configuration utilities to set the scope of your DSMs at install time and to modify, delete, or add new entries while managing your environment.

## Agent Class Scoping

DSM Agent Class Scoping displays the names of all the agent classes that the DSM manages. This list tells the DSM which agent classes to look for and which policy to load. You can define which Agent Classes your DSMs manage by checking the associated Managed Status field.

## Discovery Community Strings

Based on information in the DSM Agent Class Scoping list, you can configure the DSM to use different community strings and ports to communicate with specific node classes and agent classes.

Why is it beneficial to modify this list of community strings? Because community strings provide the authentication notification with which the devices and DSM communicate, you may want to modify them at intervals, depending on your security requirements.

## Discovery Pollset Values

The DSM Discovery Pollset Values pane displays the poll intervals, the timeouts, and the number of retries that a DSM is using to discover agents and devices. Once discovered, the devices are assigned these poll values for use in future communications.

From this list you can quickly see what intervals, timeouts, and retries a DSM uses, or whether certain DSMs use different intervals, timeouts, and retries to communicate with certain node classes or agent classes. For example, maybe your printers are pinged less frequently than your nodes, which are pinged less frequently than the routers.

Why is it beneficial to modify this list of poll intervals? Because you can customize the intervals that the DSM uses to check all the devices, based on the type of resource or type of server being monitored. If DSM queries all devices too often, the DSM and network are busier than necessary; if a DSM doesn't query the devices often enough, you may not be notified promptly of a problematic situation.

## IP Address Scoping

Each DSM has a list of nodes that it manages; this is referred to as the DSM Domain. In the past CA NSM controlled the DSM Domain at the DSM layer using a file called gwipflt.dat, then later at the repository layer using the gwipfltii.dat file.

You want to modify the list of IP addresses reporting to a DSM when the historical data of the DSM Monitor suggests that the DSM server is having reoccurring problems.

By assigning your DSMs to manage specific IP addresses, you can distribute the load of monitoring your enterprise among the number of DSMs you deploy. You can also ensure that each DSM manages only those nodes which are relatively close to it within the network, to reduce the amount of network traffic.

### IP Address Scoping Example

The IP Address Scoping tool lets you change the range of subnets or hosts that report to a particular DSM, as shown in the following subnet examples:

Task	Subnet Example
Set up an entire subnet by using wildcards	172.28.*.*
Exclude specific IP address ranges from being monitored	-172.28.192.*
Specify a range of addresses within a subnet	+172.28.192.2-8
Add another subnet range for monitoring	172.30.4.*

The above entries define the scope for a DSM to manage all discovered nodes in the 172.28 subnets, except for the 172.28.192 subnet, but also manage all hosts with IP address in the 172.28.192.2 to 172.28.192.8 range, as well as all nodes in the subnet 172.30.4.

The DSM IPScope table, in conjunction with the setdsmname service, notifies each DSM which nodes to manage without having to restart the DSM process.

### Managed Object Scoping

The Managed Object Scoping pane contains a list of managed object instances that will be created and visible as icons in WorldView or in Management Command Center topology.

When a managed object in the DSM has a matching class definition in WorldView, an object is usually created in the database and is visible from WorldView. You can limit what gets displayed to objects of interest by defining which classes on which nodes should be made visible. Customize this list so that only those resources you are most interested in are displayed in WorldView and Management Command Center.

**Example:** You can create a Business Process View that contains a group of related resources that are managed objects (such as processes, or drives, directories, or CPU usage). You can then determine the impact that any warning or critical state will have on the rest of those resources in that Business Process View.

## Understand DSM Wizard

The DSM Wizard provides a simple and centralized method to manage DSM Scope objects. It is another tool to manage and set the scope of a DSM at install time and to modify, delete, or add new entries while maintaining your environment. You can perform the same scoping of the DSM within the pages of the DSM Wizard as is possible within DSM Configuration Tools in Management Command Center.

**Note:** Although the DSM Wizard can be executed from any host where a DSM is installed, we recommend that you designate one DSM machine from which to execute the DSM Wizard. This policy helps avoid confusion and duplicate work by reducing the risk that someone will incorrectly modify the managed object scope or the IP address scope and consequently change what resources are being monitored.

When working within any DSM Wizard page, select from a list of known DSM servers to apply the entries to. The word "ANY" displayed in the Select DSM field means that the entries should apply to all DSMs you may deploy. To insert the name of a new DSM that is not on your list, click New DSM and provide the name of the DSM server.

Before you exit the DSM Wizard, answer whether you want to Enable New Configuration. After you check this box and click Finish, the assigned DSMs immediately manage agent classes and update their managed hosts.

Access the DSM Wizard by using one of the following methods:

- Click Start, Programs, CA, Unicenter, NSM, Agent Technology, DSM Wizard
- At a command prompt, enter the command, `dsmwiz`

## Monitoring the Health of your DSM

The purpose of a DSM (Distributed State Machine) is to monitor your enterprise in a thorough and efficient manner so that you or your staff is promptly made aware of any possible problems or inconsistencies. Because the DSM is an important element in your enterprise, ensure that each DSM is always connected to an MDB and that the DSM is running at peak performance.

To ensure that each DSM is running efficiently, the DSM can now manage itself and report its own status. The DSM Monitor provides real-time monitoring of the following resources:

- Connectivity status of DSM to MDB
- Number of nodes, objects, and classes that a DSM is monitoring
- System performance impact by a DSM and by other services on which the DSM depends
- Message loads coming in to and leaving the DSM
- Test of system path to WorldView and the Event Console
- Historical data collection

For information about configuring each of your DSMs from a central location, see the section [Configure DSM Environments](#).

## DSM Monitor Interfaces

The DSM Monitor accesses the same kind of data, even if your DSMs are installed on different operating systems. You can view the data gathered by DSM Monitor from the following interfaces:

- DSM Monitor View
- Node View
- DSM Monitor dashboard

## DSM Monitor View

The DSM Monitor requests information from the system agent about certain resources that the DSM uses. Then it presents the data in a graphical user interface called [DSM Monitor View](#), which allows information technology managers to respond to trends or events that influence the performance of their DSMs.

[DSM Monitor View](#) looks very similar to other [Unicenter Agent View](#) browsers. But instead of displaying the metrics and the state of monitored resources for a remote node, [DSM Monitor View](#) displays data for only those resources that are important to the efficient running of the DSM itself.

The [DSM Monitor View](#) is a subset of the [DSM Configuration](#), which was described in [Configure a DSM Environment](#). After a DSM has been configured (for classes, community strings, pollsets, IP addresses and managed objects) at the [Management Command Center](#), you can view those resources in the windows of [DSM Monitor View](#) and determine the overall health of a DSM. If [DSM Monitor View](#) consistently indicates that a DSM server is impacted negatively when polling its assigned nodes and classes, you can make adjustments with [DSM Configuration](#) tools.

To access the Summary window and view the overall status of the DSM, access Node View, right-click DSM Monitor, and select Agent View from the list.

If you are connecting to a DSM server running SNMPv3, click File, Connect, which allows you to provide the SNMP Connection Parameters.

### DSM Monitor Node View

The DSM Monitor tracks and displays the following groups of metrics for the DSM Server:

- Object status
- System path to WorldView and Event Management
- Number of monitored objects
- Performance

Node View displays a horizontal tree structure that lays out the hierarchy of these three monitored groups, with all of their states being propagated to the dsmMonitor, then to the host. The most severe state overrides any less severe state.

### DSM Monitor Dashboard

DSM Monitor dashboard looks very similar to other dashboards. But the DSM Monitor dashboard displays data for only those resources that are important to the efficient running of the DSM. The DSM Monitor dashboard presents its summary data in multiple tiles. The tiles that appear in the dashboard depend on what tile collection you select when creating the dashboard.

The DSM Monitor dashboard, like DSM Monitor View, monitors the resources thresholds, based on the settings defined in the Management Command Center with the DSM Configuration tools that was described in Configure a DSM Environment. After a DSM has been configured using the tools in the Management Command Center, you can view those resources in the tiles of DSM Monitor dashboard and determine the overall health of a DSM. If DSM Monitor dashboard consistently indicates that a DSM server is being impacted negatively when polling its assigned nodes and classes, you can make adjustments by changing the scope of that DSM using DSM Configuration tools in the Management Command Center.

## Understanding Configuration Manager

Unicenter Configuration Manager provides a seamless interface for the reporting and management of configuration information for remote and distributed Agent Technology and Event Management components.

### Resource Model Groups

Resource model groups are logical groupings of managed hosts that contain managed resources or other groups. You can create groups of hosts by function, by role, by location, or any other grouping.

**Note:** You can nest groups inside of other groups.

### Create a Group

Groups are logical groupings of managed objects that contain managed resources. The Group is created as a Business Process View in the Common Database (MDB). You can create a group that contains hosts with managed resources that have similar configuration requirements, business process views, or dynamic business process views. You can then apply one profile to the group instead of applying the profile to each individual host or managed resource. Using groups ensures consistent management and maintenance across your environment.

#### To create a group

1. Select the Groups tab from the navigation bar.  
The group hierarchy tree appears.
2. Select the model type from the drop-down above the hierarchy tree.  
The hierarchy tree for the model you selected appears.
3. Select the Management Model or the Resources Model from the hierarchy tree.  
The Model pane appears.
4. Click New Group from the right pane menu.  
The New Group pane appears.



5. Complete the following fields for the new group.

**Group Name**

Defines the name of the new group.

**Limits:** up to 200 characters

**Description**

Defines the description for the object.

**Active**

Specifies the configuration bundle is active on the group or managed host when selected. You can temporarily suspend the configuration bundle from delivering to the group or managed host when the check box is cleared.

6. Complete the object filter criteria to add a host, and click Go.  
The results of the search appear in the Available Objects list.
7. Select the host that you want to add from the Available Objects list, and click Add.  
The host is moved to the Selected Objects list.
8. Click Save as Child or Save as Sibling.  
The new group is saved and appears in the hierarchy tree.

## Base Profiles

A profile contains a set of configuration data for a managed resource. You can apply different base profiles to different groups, managed hosts, or managed resources in the hierarchy tree. The lower level profiles override the higher level profiles in the hierarchy tree. Typically, a base profile contains configuration data that is common to a number of managed hosts. A differential profile can be applied to the base profile to create minor changes to the configuration data.

## Create a Base Profile

Base profiles contain a set of configuration data for a managed resource. You can create a base profile to define the configuration policies for each managed resource. A base profile can be created based on the options selected in the New Profile - Profile Properties pane.

### To create a base profile

1. Select the Groups tab from the navigation bar.  
The group hierarchy tree appears.
2. Select the model type from the drop-down above the hierarchy tree.  
The hierarchy tree for the model you selected appears.
3. Select the Profiles tab from the navigation bar to display the profile hierarchy tree.
4. Click New Profile from the right pane menu.  
The New Profile - Profile Properties pane appears.
5. Complete the following fields:

#### Profile Name

Defines the unique name of the profile.

**Limits:** Once defined, the Name field cannot be modified.

#### Resource Class

Specifies the resource class in which the profile is created.

#### Profile Type - Base

Specifies the base profile type. A base profile contains a complete set of configuration data for a managed resource group or host.

**Note:** You must select Base for the profile type to create a base profile.

#### Profile Location URI

Defines the name and location of the URI. The URI can be located on a file server or a web server.

**Limits:** The URI file must be an XML file with the scheme component defined as a file or http.

**Syntax:** scheme://authority/path

#### Examples:

- http://localhost:9090/profiles/profile1.xml (web server)
- file:///tmp/profile1.xml (Linux)
- file://c:/profiles/profile1.xml (windows)

**Get Config from Host**

Obtains the configuration data currently loaded in the managed resource or host and uses it to create the initial profile.

**Register Only**

Registers (saves) an existing XML profile with Unicenter Configuration Manager so you can use the profile to create other profiles.

**Note:** You must supply the exact URI location for the profile in the URI Location field in order to use the register option.

6. Click Next.

The New Profile - Get Configuration from Host Pane appears.

7. Select the appropriate host class and search criteria, click Go.

The search is processed and the results are displayed in the Available Hosts list.

8. Select the host you want from the list and click Finish.

The new profile is created and appears in the profile hierarchy tree.

## Differential Profiles

A differential profile can modify a base profile by overriding (adding, deleting, updating) configuration data. Differential profiles are applied to a base profile in the following order:

- Inherited Differential Profiles in the order they were applied.
- Locally applied Differential Profiles in the order they were applied.

**Note:** A base profile must be applied to the group or host in the hierarchy tree in order to use a differential profile.

You can create a differential profile that contains a threshold value of 5 that overrides the threshold value contained in the base profile for the specific group, host, or managed resource in the hierarchy tree.

## Create a Differential Profile

A differential profile provides a way of making minor modifications to the configuration information in the base profile without having to re-create all of the information in the base profile.

### To create a differential profile

1. Select the Groups tab from the navigation bar.  
The group hierarchy tree appears.
2. Select the model type from the drop-down above the hierarchy tree.  
The hierarchy tree for the model you selected appears.
3. Select the Profiles tab from the navigation bar.  
The profile hierarchy tree appears.
4. Click New Profile from the right pane menu.  
The New Profile - Profile Properties pane appears.
5. Complete the fields for the new profile and click Finish.

**Note:** You must select Differential for the profile type to create a differential profile.

The new differential profile is created and appears in the profile hierarchy tree.

## File Packages

A file package is a collection of files associated with a managed resource that is delivered to a target host through DIA file transfer mechanisms. A file package delivers one or more files from a location on the Unicenter Configuration Manager server to a target destination on the host.

## Create a File Package

A file package is a collection of files for a managed resource that is delivered from a location on the Unicenter Configuration Manager server to a target destination on the host. You can create a file package that when contained in a configuration bundle, is delivered to managed resources that require external configuration files.

### To create a file package

1. Select the Groups tab from the navigation bar.  
The group hierarchy tree appears.
2. Select the model type from the drop-down above the hierarchy tree.  
The hierarchy tree for the model you selected appears.
3. Select the File Packages tab from the navigation bar.  
The file packages hierarchy tree appears.
4. Click New File Package from the right pane menu.  
The New File Package - File Package Properties pane appears.
5. Complete the fields for the new file package and click Next.  
The New File Package - Files for Package pane appears.
6. Click Insert.  
A new row appears in the list.
7. Complete the Source Location field for the file you want to add to the file package and click Finish.  
The file is added to the file package and the created file package appears in the file packages hierarchy tree.

## Delivery Schedules

A Cron expression or calendar-based schedule, that when combined in a configuration bundle with a profile or file package, facilitates the audit and delivery of the bundled profile or file package.

## Create a Delivery Schedule

Delivery schedules are calendar or Cron expression-based schedules that, combined in a configuration bundle with a profile or file package, facilitate the audit and delivery of the bundled profile or file package. You can create a delivery schedule to meet the needs of your IT environment and to ensure consistent delivery of profiles or file packages to resource model groups.

### To create a delivery schedule

1. Select the Groups tab from the navigation bar.  
The group hierarchy tree appears.
2. Select the model type from the drop-down above the hierarchy tree.  
The hierarchy tree for the model you selected appears.
3. Select the Delivery Schedules tab from the navigation bar to display the delivery schedule hierarchy tree.
4. Click New Schedule on the right pane menu.  
The New Delivery Schedule pane appears.
5. Complete the fields to customize your delivery schedule, and click OK on the right pane menu.  
The new delivery schedule is created and appears in the delivery schedule hierarchy tree.

## Configuration Bundles

Configuration bundles are logical groupings of one Base Profile or File Packages together with a delivery schedule. You can also add Differential Profiles to a configuration bundle.

Base Profiles and Differential Profiles contain agent configuration data that is automatically loaded into the sadmin store after its delivery to the target server.

File Packages contain other files (scripts or configuration files) that are not loaded into the sadmin store and that have to be copied to specific locations on target servers.

Adaptive Configuration Profiles are contained in File Packages because they must be copied into specific directories on the target servers. Adaptive Configuration Profiles contain specific instructions for the Adaptive Configuration Service, which automatically creates agent configuration data on the target servers according to these instructions.


**Note:** You should not deliver a Base Profile and an Adaptive Configuration Profile for the same agent to the same target server. The Base Profile can overwrite the configuration data that was created by the Adaptive Configuration Service.

## Create a Configuration Bundle

You can create a configuration bundle to apply profiles and file packages to the group of managed hosts or resources.

**Note:** To create a configuration bundle you must create a group, a base profile or file package, and a delivery schedule.

### To create a configuration bundle

1. Select the Groups tab from the navigation bar.  
The group hierarchy tree appears.
2. Select the model type from the drop-down above the hierarchy tree.
3. The hierarchy tree for the model you selected appears.
4. Select the managed resource that you want from the hierarchy tree.  
The Group - Applied Configuration Bundle pane appears.
5. Click New Configuration Bundle.  
The New Configuration Bundle - Select Resource Class pane appears.
6. Select the resource class delivery schedule from the drop-down, and click Next.  
The New Configuration Bundle - Add File Packages pane appears.
7. Click Go from the file package search.  
The results of the search appear in the Available File Packages list.
8. Select the file packages that you want to assign to the configuration bundle from the Available File Packages list, and click Add .  
The file package is moved to the Selected File Packages list.

**Note:** You are not required to add a file package to the configuration bundle before moving to the next step.

9. Click Next to add a base profile.

**Note:** You can also click Finish to complete and save the new configuration bundle.

The New Configuration Bundle - Select Base Profile pane appears.

10. Click Go from the base profile search.

The results of the search appear in the Select Base Profile list.


11. Select the base profile that you want to assign to the configuration bundle by selecting the Select to Add column and click Next.

**Note:** You can also click Finish to complete and save the new configuration bundle.

The New Configuration Bundle - Add Differential Profiles pane appears.

12. Click Go from the differential profile search.

The results of the search appear in the Available Differentials list.

13. Select the differential profile that you want to assign to the configuration bundle from the Available Differentials list and click Add .

The differential profile is moved to the Selected Differentials list.

14. Click Finish.

The configuration bundle is created and appears in the applied configuration bundles list.



## Reporting Feature

The reporting feature of Unicenter Configuration Manager lets you configure and generate reports. The reports provide an audit trail for Unicenter Configuration Manager and include the user ID, date, and time on which the object was last configured. The following reports are available:

### **Configuration Bundles Audit Report**

Displays a list of configuration bundles that were created, updated, or deleted during the specified time range.

### **Configuration Objects Audit Report**

Displays a list of base profiles, differential profiles, and file packages that were created, updated, or deleted during the specified time range.

### **Resources Model Audit Report**

Displays a list of resource models that were created, updated, or deleted during the specified time range.

### **Delivery Schedules Audit Reports**

Displays a list of the delivery schedules that were created, updated, or deleted during the specified time range.

### **Delivery Forecast Reports**

Displays a list of deliveries scheduled in the future within a specified time and date range.

### **Delivery Status Reports**

Displays the status of deliveries during the time range selected and whether the delivery was successful or failed.



# Chapter 8: Administering Critical Events

---

This section contains the following topics:

[Event Management](#) (see page 275)

[Alert Management System](#) (see page 312)

## Event Management

Event Management, the focal point for integrated message management throughout your network, can monitor and consolidate message activity from a variety of sources. It lets you identify event messages that require special handling and initiate a list of actions for handling an event. Through support of industry-standard facilities, you can channel event messages from any node in your network to one or more monitoring nodes. You can centralize management of many servers and ensure the detection and appropriate routing of important events.

For example, you may want to route message traffic to different event managers:

- Event and workload messages to the production control event manager
- Security messages to the security administrator's event manager
- Problem messages to the help desk administrator's event manager

By filtering messages that appear on each console, you can retrieve specific information about a particular node, user, or workstation.

Wireless Messaging provides alternate channels for operator input in situations where the operator cannot access a CA Event Console. The supported messaging protocols are email and pager. Using the SMTP/POP3 mail messaging protocol, you can send and receive pager messages from two-way pager devices. An incoming message can trigger any series of actions you define for Event Console to perform in response to it.

Successfully implementing Event Management involves the following activities:

- Establishing date and time controls for automated event processing
- Trapping important event messages and assigning actions
- Putting Event Management policies into effect
- Monitoring message traffic
- Controlling access to messages
- Providing Wireless Message Delivery
- Using SNMP to monitor activity
- Implementing maintenance considerations

**Note:** For more information about Event Management, see the guide *Inside Event Management and Alert Management*.

## Events

An event is a significant situation that indicates a change in the enterprise. It can be positive, negative, or just informative. It can indicate a significant problem or just describe a situation. It can be a warning of conditions that indicate a possible future problem, or it can tell of the success or failure of certain things. When an event occurs, a message is usually sent. Event Management processes it using its Event Management policy.

## Event Management Policies

The Event Management service reads all message records and message action policy definitions from the Management Database (MDB) during startup and stores them in memory cache. It also writes a copy of this policy to a local file called Decision Support Binary (DSB) to be used when the MDB is not available. Any changes to message records and actions take effect only after the Event Management service refreshes those policies in memory. You accomplish this by using the `opreload` command or by shutting down and restarting Event Management.

The `opreload` command directs Event Management to refresh the active message record and message action lists immediately with the definitions stored in the Management Database. Any requests for automated message processing remain in the queue until the refresh of the active lists is complete.

## Event Agent

The Event Agent is a lightweight solution that provides all Event Management functions with very little overhead. Servers running the Event Agent do not have a Management Database or the administrative GUI. The agent gets Event policy from an Event Manager server or a local DSB file.

When you install the Event Agent, you indicate whether to load the policies from a local DSB file or from a specific remote Event Manager. When you start the Event Agent or run `opreload`, message records and actions are copied from the local DSB file or the specified Event Manager, and an in-memory version is created on the agent computer.

## Non-Root Event Agent

On UNIX/Linux you can install a non-root Event Agent. This agent increases security because it enables a specified non-root user to start and stop the agent and run Unicenter processes. Without the non-root Event Agent, Event Management must be started and stopped by the root user.

The daemons on the non-root Event Agent run under the ID of the non-root user, except for the CAICCI processes, which run as root.

**Note:** For information about CAICCI (CA Common Communications Interface), see the *Administration Guide* and the online *CA Reference*.

Message actions that run on the system run under the non-root user's ID unless `sudo` is used. The shareware utility `sudo` enables the non-root Event Agent to run programs and commands on behalf of any user. You can download `sudo` from [www.courtesan.com/sudo](http://www.courtesan.com/sudo). Follow the instructions on that site for installing, configuring, and using the utility.

## sudo and Message Actions

The message actions `UNIXCMD` and `UNIXSH` submit a command to a spawned UNIX/Linux shell. The following example shows how to integrate `sudo` into a message action for a non-root Event Agent. The `sudo` command is used to execute `/usr/bin/touch` as the specific user `userNAME`.

Command: `/usr/bin/touch myfile`

Command: `sudo -u userNAME /usr/bin/touch myfile`

The `sudo` command can be in the message action text field or in a shell script executed by the message action.

## Configure sudo

The sudoers file lets you configure sudo. See the sample file on [www.gratisoft.us/sudo](http://www.gratisoft.us/sudo). To modify the sudoers file you must use the visudo command. You need to add users and permissions to the file:

- Add authenticate parameters for each user so that sudo does not prompt for a password:

```
Defaults:user    !authenticate
```

- The following excerpt from the sudoers file (sudo configuration file) gives the user unimgr permission to execute the /usr/bin/touch file as root or opsuser on server1:

```
unimgr server1 =(opsuser) /usr/bin/touch, (root) /usr/bin/touch
```

## How Event Agents Are Implemented

Usually one Event Manager provides message records and actions for several Event Agents. A typical implementation consists of a manager on a Windows server and agents on UNIX/Linux servers.

Setting up a manager and agents involves the following tasks:

1. Install the full Event Management on a Windows server. Select the installation wizard for CA NSM.
2. Install the agent on UNIX/Linux servers. Select the installation wizard for Event Agents.
3. Verify the installation.
4. Add the agent machines to the administrative configuration on the manager server.
5. Create message records and actions on the manager server for the remote agents.
6. Activate the message records and actions on the agent machines.

## Configure the Event Agent

When Event Agents are installed, configuration settings are applied automatically based on your installation choices. You may need to change some settings later, however. This topic provides information for Windows and UNIX/Linux.

### Windows

Run the `cautenv` utility from the command line of the Event Manager or Event Agent.

- On the Event Manager:

```
opr cmd -n agent-name cautenv setlocal envname value
```

- On Event Agents:

```
opr cmd cautenv setlocal envname value
```

**Note:** The user running the commands must be listed in `CA_OPR_AUTH_LIST` (Users Authorized To Run Commands) on the agent computers.

The following are examples of settings that can be changed. Substitute *envname* with one of the following environment variables.

#### **CA\_OPR\_USEDDB (Load from Management Database?)**

Specifies whether the Event daemon should use the Management Database. Set this to N because Event Agent installations have no database.

#### **CA\_OPR\_PROXY (Event Agent Proxy Node)**

Indicates the name of the Event Manager server that provides policy to the Event Agent. If no value is specified, policies are loaded from the local DSB file.

#### **CA\_OPERA\_NODE (Console Daemon Node)**

Specifies the name of the server where event messages are forwarded. You may want to set `CA_OPERA_NODE` to the local agent computer so that it processes its own events. You may need to use Event policies to forward some events to the manager for processing.

### UNIX/Linux

To change Event Management settings on UNIX/Linux, edit the configuration files `$CAIGLBL0000/opr/scripts/envset` and `$CAIGLBL0000/opr/scripts/envusr`. For example, the following variable is set in the `envset` file based on the response to an installation question.

#### **CAI\_OPR\_REMOTEDB (Event Agent Proxy Node)**

Indicates the name of the Event Manager server that provides policy to the Event Agent.

## Dates and Times for Automated Event Processing

Determining a course of action based on when an event occurs can be critical to its proper handling. You can define calendars through the Enterprise Management GUI or the cautil command line interface. The GUI has a significant advantage over command line entry because it simplifies date and time specifications. You can click specific days, months, and times, or you can click a specific day and drag the cursor to the last day, "rubber-banding" the range of days to be set. You can define as many calendars as you require and store them for easy reference.

**Note:** The use of calendars is optional with Event Management and Alert Management.

Since the calendar object is a common object, and calendars may be shared by any of the Enterprise Management functions, we recommend that you use a naming scheme to identify the primary function of your calendar. You can create calendars for use by any of the Enterprise Management functions. For example, you can use a single holidays calendar for your company whenever company holidays need to be considered. After creating and saving the holidays calendar, the next time you create a calendar, you can associate your holidays calendar with the new calendar by selecting Options Configure commands. When you want to combine the dates set in this calendar with the new calendar you need only to click the appropriate command button.

## Automatic Responses to Event Messages

Through Event Management message record and action profiles, you can identify events that are important to your operation and define the special processing that CA NSM performs when encountering them. CA NSM components, system components, utilities (such as cawto), and user programs generate event messages and send them to an Event daemon for processing. The default processing, in the absence of any policy, is to write events to the Event Console log file.

To define an event processing policy to filter events received by the Event daemon and define specific actions to take, you begin by defining message records that describe which events need processing. The message record identifies matching criteria. Each field in the event message has a corresponding field in the message record. Message record fields may contain wildcards to match a wider range of event messages. A message record has associated message action records that describe the action to take when a message matches the message record.



## Event Sources

Event Management receives events from a variety of sources:

- The `cawto` command, which sends an event to the Event Console.
- The `cawtor` command, which sends an event to the Event Console and waits for a reply. It appears in the held messages pane and will not be deleted until the operator replies.
- The `oprcmd` command, which sends a request to execute a command to the designated target machines.
- The `careply` command, which lets you use any terminal to reply to an event being held by the Event Console.
- Enterprise Management components, which generate events directly to the Event Console.
- SNMP traps that are generated by various devices, such as switches or printers, and other software components. `catrapd` (an Event Management component), collects, formats, and routes these traps to the Event Management daemon on the local or remote node.
- The Windows Event Logs, which store events generated by the Windows operating system, device drivers, or other products. The Event Management log reader collects these events and forwards them to the Event Management daemon.
- The syslog daemon on UNIX/Linux platforms, where messages are routed through the syslog daemon to the Event Console. Events issued through the logger utility are included as they also use the syslog daemon. These events may have originated on a platform not running CA NSM.
- Agent Technology agents, policies, and DSM.
- Any CA or client programs that use the CA NSM SDK.
- API functions, such as `EmEvt_wto`, which issue events to Event Management.

For additional information about the `cawto`, `cawtor`, `oprcmd`, `careply`, and `catrapd` administrator commands, see the online *CA Reference* and the *CA SDK Reference*.

## Message Records

You identify events that require special handling by creating message record objects. You then specify the special handling requirements by creating message action objects that are associated with a particular message record object. Once defined, message records and message actions become an event handling policy that identifies events with special handling requirements and the tasks to perform when they are detected.

Event Management provides two categories of message records to identify important events:

- **Message** - Represents the output text string received by Event Management and displayed on the Event Console.
- **Command** - Represents the text string input by someone operating the Event Console. (You can enter commands at the command field of the console, use customized buttons to automatically issue commands, or enter them as command line arguments provided to the `opr cmd` utility.)

Command output can be a source of text to substitute into the message text in Management Database message records during the message matching process. For example, the string ``pwd`` in the MDB record message text field causes the current directory to be inserted into the message text.

## Message Actions

Message actions specify what Event Management should do when it detects a match between an input event message and a message record. Possible actions range from simply highlighting messages on the console display to replying to messages, opening problems, or executing commands or other programs.

For example, to ensure that a message catches the attention of the person responsible for monitoring the console, you can use either or both of these methods:

- Route the message to a held area of the console GUI where it remains until acknowledged by the console operator.
- Assign an attribute, such as highlighting or blinking, to make a message more noticeable on the Event Console.

You can use several types of actions in any sequence or combination to thoroughly automate processing of an input or output message. For explanations of these action keywords, see the `cautil DEFINE MSGACTION` control statement in the online *CA Reference*.

## Message Activity Distribution

Event Management lets you distribute message and action activity across multiple servers and their clients. You can:

- Create a central event log from which all servers can be monitored.
- Send selected messages to another server for processing.
- Manage special functions such as security or tape management on dedicated consoles.

Whenever the Management Database is loaded, it checks the EVALNODE of every message record against its own node name. If its node name matches the EVALNODE of the message record, the record and all associated message actions are read into memory. If there is no match, the message record is ignored. The set of message records and message actions read into memory constitute the Event policy for the current execution of the Event daemon until the policy is reloaded by a restart or the opreload command.

## Message Action Policy Definitions and Servers

Through message record and action policies, you select a message based on content and then define the appropriate action. A possible message action is to instruct Event Management to perform an action on a specific (and potentially remote) machine. You can easily define an action such as sending the message to a remote machine or initiating a command on the remote machine. For example, you can identify a network security event or a tape mount event for routing to an alternate machine simply by defining Event Management policies to that effect.

Actions to be performed on remote nodes can be done synchronously or asynchronously. Before proceeding to the next action, action processing waits for the remote action to complete and return a completion code. This means that the completion code received from the remote action can be tested as part of a message action and used to control subsequent processing. Remote actions are not attempted if the target node is known to be unreachable.

## Test Policy by Simulating Messages

Event policy can be difficult to test because messages have so many variables that are not easily duplicated unless generated by an actual event. The cawto command, which sends a message to the console, lets you test message policy by sending a message without creating the associated event.

To test policy by simulating messages, open a command-line interface, and enter the cawto command using the options for simulating messages.

## Message Routing to Remote Hosts

The best way to route messages to a remote machine through message/action policies is to create a message record that traps those events and a message action that uses the FORWARD action keyword and specifies the remote node to which you want the messages sent. Since one message can have many message actions, it is easy to send messages to multiple machines.

You can specify the name of any machine that is currently defined for CAICCI remote communication. The configuration is set up during installation.

On UNIX/Linux platforms, one source of event messages is the Berkeley syslog daemon, which can route messages to a CA NSM server for processing, even those originating from servers not running CA NSM.

Event Management takes advantage of the powerful messaging facilities provided by the syslog daemon to:

- Select from several priorities, levels, and facilities of messages.
- Route messages by level or priority to different devices.
- Route messages by level or priority to different hosts.
- Receive messages from other hosts for local display.

To instruct the syslog daemon to route all messages to a remote machine, edit the syslog daemon's configuration file and insert the remote host name in the action part of the line, prefixing the host name with a single at sign (@).

**Note:** If you use both the Berkeley syslog daemon and specific message action policies to reroute the same messages to the same remote machines, those messages will display twice on those remote machines as they were sent there twice, once by the Berkeley syslog daemon and again by Event Management.

## Message Action Restriction

On UNIX/Linux platforms, Event Management lets you restrict the nodes and RUNIDs authorized to send the message actions COMMAND, UNIXCMD and UNIXSH to your local host.

During installation, if setup detects that Event Management was previously installed at that node, a message appears informing you of the new message action restriction feature and the default setting that disables message action restriction. You have the opportunity to override the default and enable message action restriction.

If you accept the default response n to the prompt for message action restriction, setup creates the actnode.prf configuration file for you with a single entry of -n=\*,\*,E to allow all RUNIDs from all nodes to submit these message actions.

If you instead respond to the prompt for message action restriction, setup creates the `actnode.prf` configuration file with a single entry of `-n=*,*,D` to deny all RUNIDs from all nodes the ability to submit these message actions.

When setup detects that you are installing Event Management for the first time on the node, a message appears informing you of the new message action restriction feature and the default setting that disables message action restriction. You are given the opportunity to override the default and enable message action restriction at that time.

If you accept the default response `n` to the prompt for message action restriction, setup creates the `actnode.prf` configuration file for you with a single entry of `-n=*,*,E` to enable message action submission for all RUNIDs from all nodes.

If you instead respond `y` to the prompt for message action restriction, setup creates the `actnode.prf` configuration file with a single entry of `-n=*,*,D` to disable all RUNIDs from all nodes from submitting these message actions.

You can change this rule at any time after installation by executing the `caevtsec` utility located in the `$CAIGLBL0000\bin` directory. The utility only allows the `uid 0` user to maintain the file and preserve the file permissions. The file may also be maintained using a UNIX/Linux text editor. For more information about using the `caevtsec` utility, see the online *CA Reference*.

The `actnode.prf` configuration file is located in the `$CAIGLBL0000/opr/config/hostname` directory. You can use this file to maintain policies that specify how message action restriction is to be enforced based on the submitting node and RUNID. The file must be owned by `root` and only a `uid of 0` may have write access to it. An individual entry in the file has the following format:

```
-n=nodename,runid,flag
```

**nodename**

Specifies the node from which the `COMMAND`, `UNIXCMD` or `UNIXSH` message action is initiated; it may contain a trailing generic mask character.

**runid**

Specifies the node from which the `COMMAND`, `UNIXCMD` or `UNIXSH` message action is initiated; it may contain a trailing generic mask character.

**flag**

Specifies `D` for disable (feature is active; disallow the message action submitted by RUNID from `nodename`), `E` for enable (allow the RUNID from `nodename` to submit the message action), or `W` for warn (check the rule but allow the message action submission to occur).

For example:

```
-n=*,*,E
```

is the default rule in effect if, during installation, you elected not to activate message action restriction. The rule states that for all nodes and all RUNIDs, COMMAND, UNIXCMD and UNIXSH message action submission is allowed.

```
-n=*,*,D
```

is the default rule in effect if, during installation, you elected to activate message action restriction. The rule states that for all nodes and all RUNIDs, COMMAND, UNIXCMD and UNIXSH message action submission is disallowed.

```
-n=*,*,E  
-n=*,root,D
```

enforces a message action restriction on RUNID root and allows all other RUNIDs to submit the message actions.

```
-n=*,*,E  
-n=mars,*,D  
-n=*,root,W
```

allows all RUNIDs to bypass message action restriction unless the request comes from the node mars. In that case, message action restriction is enforced for all RUNIDs. The last entry sets a warning type restriction rule for RUNID root if it comes from a node other than mars.

Event Management scans the entire configuration file for a best match and uses that rule. It uses the node field as a high level qualifier when searching for a best match. For example if the following are the only two entries in the file, any request coming from the node mars uses the disallow rule. The user root only uses the warning rule if the request comes from a node other than mars.

```
-n=mars,*,D  
-n=*,root,W
```

**Note:** On Windows, to execute a command a user must be defined in the Users Authorized to Issue Commands configuration setting.

## Environment Variables for Messages and Actions

When Event Management invokes a command or script as a result of a COMMAND or UNIXCMD action, the new process is created with a list of new environment variables that contains information about the event that you may find useful. Programs, Windows .bat files, or UNIX/Linux shell scripts can reference these variables, but they cannot be altered.

- EVENT\_CATEGORY - Category field from the event
- EVENT\_DATEGEN - Date (yyyy/mm/dd) event was generated
- EVENT\_DEVICE - Device associated with the event
- EVENT\_JOBNAME - Event jobname (if event came from a Job Management Option job)
- EVENT\_JOBNO - Event job number (if event came from a Job Management Option job)
- EVENT\_JOBQUAL - Event job qualifier (if event came from a Job Management Option job)
- EVENT\_JOBSET - Event jobset (if event came from a Job Management Option job)
- EVENT\_LOGRECID - ID of the last record written to log when the command is invoked (possibly the current event if it is not suppressed)
- EVENT\_MSGNUM - Message number field from the event
- EVENT\_NODEID - Node origin associated with the event
- EVENT\_OPUSER - User name under which Enterprise Management is processing this message action
- EVENT\_PID - ID (decimal) of the process that generated the event
- EVENT\_PROGRAM - Program that generated the event (for example, cawto.exe)
- EVENT\_REPLID - Reply ID returned if the event was WTOR
- EVENT\_SEQNO - Sequence number of the action that invoked this command
- EVENT\_SEVERITY - Event severity (usually I, W, E, S, or F)
- EVENT\_SOURCE - Event source
- EVENT\_STATION - Event station (usually associated with a Job Management Option job)
- EVENT\_TAG - Platform tag associated with the event (WNT, HPUNIX, and so forth.)
- EVENT\_TEXT - Full text of the event
- EVENT\_TIMEGEN - Time (hh:mm:ss) event was generated

- EVENT\_TIME8 - Time (hh:mm:ss) command was invoked
- EVENT\_TOKEN - Token number of message record that matched this action
- EVENT\_TYPE - Type of event: MSG/CMD/REPLY/WTOR
- EVENT\_UDATA - User data (value of the CA\_UDATA environment variable when the event was generated)
- EVENT\_USERID - User origin associated with the event
- EVENT\_YYYYMMDD - Date the action was invoked

### Message Enhancement

Event Management enhances messages by automatically providing the source or origin of each message along with the message text. You can customize the message text to meet the specific characteristics of your enterprise.

Use the following action keywords to control the message text that appears on the Event Console:

- EVALUATE
- FORWARD
- HILITE
- SENDKEEP
- SENDOPER
- WAITOPER

For more information, see the `cautil DEFINE MSGACTION` control statement in the online *CA Reference*.



## Event Correlation

Often a single event coming across the Event Console is not important unless seen in context with other events. By constructing a series of message records and actions, you can be notified and take action if two or more events occur that together have more significance to your enterprise than any one of the events may have when it occurs in isolation.

For example, assume you have two PCs in your accounting department. If one goes down, it is a problem and you probably have established policies to deal with such an occurrence. However, should the second also go down, the problem suddenly becomes critical. The action you want to take in this situation may be quite different.

A solution is to define message records to trap events coming to the Event Console informing you that Accounting PC #1 and Accounting PC #2 are coming down. Then, for each message record, define message actions that test for the occurrence of the other event. As a result, you will be automatically notified of the critical situation that exists in the Accounting department.

## Event Console

Event Management gives you a visual window into event activity that lets you view and immediately respond to events as they occur. The Event Console provides two areas to view event messages written to the console log.

- The held messages area displays messages that require a response. These messages are often critical and require immediate attention or require an operator reply. Held messages that require an operator reply are WTOR (Write To Operator with Reply) messages. When no held or WTOR messages exist, the area reserved for messages of that type disappears. If a reply pending message (WTOR) has been sent, and either the Event Manager or the entire system goes down while the message is still pending, the message is queued and activated automatically (appears to be still active) when the Event Manager is brought back up.
- The log messages area displays all logged messages (including held messages). Through message records and actions, you can further highlight these messages with a variety of colors and attributes to make them more noticeable on the console display.

You can narrow the focus of the display so you can concentrate on events pertinent to your immediate situation. For example, you can apply filters that limit the number of messages displayed. You can also append comments to messages that come across the console.

## Security for Console Log Viewing

Security Management support for console log viewing lets you restrict message access to authorized users and user groups. By defining console view objects to the Management Database, you can filter messages from the console logs, thereby limiting access to sensitive messages.

Security Management provides two asset types for defining console log message access:

- CA-CONVIEW--Event Management treats each console view record as an asset of this type. You can define rules to control user and user group permissions for creating, modifying, deleting, or listing these console view objects.
- CA-CONVIEW-ACCESS--When users select the console log, access granted or denied by use of this asset type determines the view. You can define this asset type through Event Management or through Security Management user and user group profiles. By attaching a CA NSM calendar to definitions of this asset type, you can restrict console log access to specific times. For example, assume that USER1 is limited to VIEW1 Monday through Friday during regular business hours, but only to VIEW2 outside of normal business hours; you could enforce these restrictions using a CA NSM calendar.

For detailed descriptions of these asset types, see the Asset Types Table in the online *CA Reference*.

**Important!** To enforce your access rules, you must define users in FAIL mode. The only Enforcement mode that results in access being denied is FAIL mode, whether set explicitly in the user profile or implicitly by referring to a System Violation Mode of FAIL in the user profile.

After defining console view access rules, you can execute the commit process to put them into effect.

Users accessing the console log can choose from a list of console views associated with their user IDs. If no console view access rules exist for a user, the entire console log appears. When a user is removed from the console view definition, that view is no longer available to the user.

## Console Log File

The console log provides both a real-time and an historical online interface into all event activity. Created daily, the console log is a file containing all events written to the log on a particular day. You can view the log of the previous day or next day (relative to the date of the log you are currently viewing), or select a log created on a specific date.

Only Event Management can read these files. You can set the directory where these files are stored with the Console Files Directory (CAI\_CONLOG) environment variable. See the online *CA Reference* for more information about this environment variable.

You should consider a policy for archiving outdated console log files to tape for off-site storage.

On UNIX/Linux and HP NonStop systems, you can choose to view all events in the log file or you can download the last portion of the log file. For example, you can limit your view to the last two hours. By default, you are prompted for a choice when you start console log. After making an initial selection for a partial retrieval, you can request additional retrieval by choosing View Retrieve. You can also click the Retrieve Log button on the toolbar.

## Event Logs and Mobile Devices

Unicenter for Pocket PC contains a lightweight interface that provides access to console logs on other servers. To better fit the form factor, the interface differs from the conventional console GUIs by emphasizing filtering of the log records you want to receive and display.

As an extension to the filtering functionality available with the Java and Win32 consoles, the Pocket PC menus display intuitive event filters that allow you to retrieve events without creating complex filters. For example, to retrieve all events related to CA NSM Security, selecting Security from the Filter menu retrieves all security-related events. In addition to these convenient pre-defined filters, event filtering is fully customizable on the Pocket PC.

Unicenter for Pocket PC allows for the dispatch of events to specified Event Managers with the conventional WTO (Write To Operator), WTOR (Write To Operator with Reply), and OPRCMD commands. For more information on OPRCMD, see the online *CA Reference*.

## Multiple Remote Console Logs

By default, servers and administrative clients are set up to view the Console log of only one machine at a time. This configuration is suitable for sites that have one server running Event Management and one active Console log. On Windows, the variable `CA_OPERA_NODE` (Console Daemon Node) on the Configuration Settings, Event Management page identifies this computer.

You can, however, view multiple remote console logs from one machine by doing one of the following:

- Adding the node name to the server map `UNISHARE$`.
- Adding the node name to the `CONSOLE.TAB`. This file does not exist automatically; you must create it. A sample is provided in `%CAILOCL0000%\CONSAMPL.TAB` (where `CAILOCL0000` is a Unicenter environmental variable whose value you can view by running the `"cautenv CAILOCL0000"` command).

For more information about multiple log file support, see the topic [Event Management Console Table \(console.tab\) File](#) in the online *CA Reference*.

## SNMP Traps

Simple Network Management Protocol (SNMP), a widely used standard in network Event Management, identifies objects on a network and provides a method to monitor and report on their status and activities.

An SNMP trap is usually an unsolicited message that reports on one of two types of events:

- Extraordinary events indicate something is wrong, or an error has occurred.
- Confirmed events provide status information, such as a process ending normally or a printer coming online.

Many SNMP agents are available, including those provided through Unicenter Agent Technology. Although they vary in purpose, complexity, and implementation, all SNMP agents can:

- Respond to SNMP queries
- Issue an SNMP trap
- Accept instructions for routing an SNMP trap (accept a setting for a trap destination)

**Note:** On some UNIX/Linux computers, an `snmptrapd` system might be running, occupying port 162. If so, `catrapmuxd` stops `snmptrapd` and starts it at port 6164, which frees port 162 for `catrapmuxd`. When `catrapmuxd` is shut down, it stops `snmptrapd` and restarts it, listening to port 162.

## Support for SNMP Version 3 Traps

Event Management supports receiving and processing SNMP version 1, version 2c, and version 3 traps. On UNIX/Linux, the trap manager supports version 1, 2c, and 3 traps by default. On Windows, the default is to receive and process only versions 1 and 2c traps using the Windows SNMP service. To receive and process version 3 traps, in addition to versions 1 and 2c, you must install and configure the CA Trap Multiplexer (TRAPMUX) by entering the following command:

```
catrapmuxd UniConfig
```

The CA Trap Multiplexer also supports IPv6. Therefore, make sure that you use `catrapmuxd` instead of the Windows SNMP service if using IPv6 on Windows versions earlier than Windows Vista, because the Windows SNMP service does not support IPv6 on these versions.

For more information about `catrapmuxd`, see the online *CA Reference*.

TRAPMUX requires port 162 to be available. On Windows, if port 162 is in use, `catrapmuxd` issues an error message. You must free port 162 before attempting to run `catrapmuxd` again. On a Windows server system with the Windows SNMP service installed, the Windows SNMP service is probably using port 162. You can configure the Windows SNMP service to run on a different port. TRAPMUX can forward traps to that port if the Windows SNMP service is still required.

**Note:** When TRAPMUX forwards a trap to the Windows SNMP service, the trap loses its original embedded address. From the perspective of the Windows SNMP service, the trap originated on the local node with `catrapmuxd`. This also applies to any third-party SNMP service manager configured to receive traps from TRAPMUX.

### To support SNMP version 3 Traps

1. Shut down the `snmp` and `snmp-trap` services.
2. Open the `%system%/drivers/etc/services` file.
3. Change `snmptrap 162/udp` to `snmptrap xxxx/udp`, where `xxxx` is a port not currently in use, for example: `snmptrap 5162/udp`.
4. Save and close the services file.

After freeing port 162, if the Windows SNMP service is still required, follow these steps:

1. Restart the snmp and snmp-trap services.
2. To enable the Windows SNMP service to receive traps from TRAPMUX, enter the following command:

```
catrapmuxd add snmptrap:xxxx
```

where xxxx is the port to which snmptrap was moved in the services file, for example: catrapmuxd add snmptrap:5162.

### Authorize SNMP Version 3 Users for CATRAPD

To identify SNMP version 3 users from which you want to receive traps, manually modify the authorization file located, in a default installation, at: CA\SharedComponents\CCS\CommonResourcePackages\Misc\snmpv3.dat.

The format for adding SNMP version 3 authorizations is:

```
h:c:cn u:sl:ap:a:pp:p
```

where:

**h**

Specifies the host subnet or range for every authorization. This is a required field.

**c**

Specifies the agent classname.

**cn**

Specifies the contextname, or instance.

**u**

Specifies the username. This is a required field.

**sl**

Specifies the snmpSecurityLevel. This is a required field.

```
noAuthNoPriv
```

```
AuthNoPriv
```

```
AuthPriv
```

**ap**

Specifies the authProtocol. This is a required field if sl is not set to noAuthNoPriv.

```
MD5
```

```
SHA
```

**a**

Specifies the authentication password. This is a required field if `sl` is not set to `noAuthNoPriv`.

**pp**

Specifies the `privProtocol`. This is a required field if `sl` is `AuthPriv`.

DES

**p**

Specifies the privacy password. This is a required field if `sl` is `AuthPriv`.

**Examples**

- Set all hosts in the range to have minimum SNMP version 3 security (no authentication required) if the user is `evans`:  
`172.24.111.5-15:*:* evans:noAuthNoPriv`
- Set all hosts in the range to have `AuthNoPriv` security using MD5 protocol and an authentication password of `evansa` if the user is `evans33`:  
`172.24.111.5-15:*:* evans33:AuthNoPriv:MD5:evansa`
- Set all hosts in the range to have `AuthPriv` security using SHA protocol, an authentication password of `AJHa0123` and a privacy password of `AJHp0123`, if the user is `AJH3`:  
`172.24.111.5-15:*:* AJH3:AuthPriv:SHA:AJHa0123:DES:AJHp0123`
- Remove the node from SNMP version 3 security, and let it default to SNMP version 1/version 2 security:  
`-172.24.111.11`

**Note:** You must recycle `catrapd` for the updated authorized user information to take effect.

## Encrypt the snmpv3.dat File

When testing SNMP version 3 authorizations, we recommend that you use a clear text version of the file. When satisfied that everything is working, encrypt the file to be used in the production environment. We do not recommend leaving unencrypted SNMP version 3 authorization files in the production environment.

For example:

1. Test and validate snmpv3.dat.
2. Encrypt snmpv3.dat.
  - `aw_enc -i snmpv3.dat -o snmpv3.dat.crypt`
  - move the clear text snmpv3.dat to some archive area
  - `ren snmpv3.dat.crypt snmpv3.dat`

For more information about authorization of SNMP version 3 agents, see *Agent Technology Support for SNMPv3*.

## Trap Destinations

Traps should be routed to a destination where an action can be taken. Many vendors provide facilities for setting a system-wide default trap destination through an SNMP configuration file. For example, some UNIX/Linux platforms set their trap destination in the `/etc/snmpd.conf` file. This path and file name may be different for your system.

After a trap destination setting is accepted, there must be something at that destination to receive and process the trap. An Event Management agent, CA trap daemon (catrapd), automatically receives and processes traps directed to the destination (machine) on which it is executing.

catrapd receives an SNMP trap, unpacks (decodes) it, and sends it to other Event Management components for processing. As part of this decoding, character representations, or strings, can be assigned to substitute names for the Enterprise IDs that are part of the SNMP trap. CA NSM provides the following translation files for that purpose:

- `%CAIGLBL0000%\WVEM\DB\enterprise.dat` on Windows platforms
- `$CAIGLBL0000/snmp/dat/enterprise.dat` on UNIX/Linux platforms

On Windows, to enable enterprise name translation, go to EM Settings and modify the 'Enterprise OID displayed as' setting to 'NAME.' Recycle catrapd so that the change takes effect.

**Note:** On Windows you can update the enterprise.dat file with the command `catrapd update`.

For more information about catrapd, see the online *CA Reference*.



## How catrapd Formats Traps

If automatic formatting of SNMP traps is enabled, catrapd uses the TRAP and MIB tables in the CAITRPDB database to determine which traps to format. A TRAP record identified by a unique Eid (Enterprise ID), Generic, and Specific key contains information needed to format one trap. A MIB record identified by a unique Name key contains information about one MIB. All TRAP and MIB records have Enable columns that control which traps to translate. A MIB is a collection of related traps generated by one software program or hardware device. One company or organization usually compiles a MIB.

- If an incoming trap matches the Eid (Enterprise ID), Generic, and Specific trap fields, and the corresponding TRAP and MIB records are enabled, CATRAPD translates it to readable form.
- catrapd uses the Format, Args, and AlarmName columns of the matching TRAP record and the Variable Bindings (VarBinds) from the incoming trap to create the final text in the following format:

```
%CATD_I_066, AlarmName: unique variable-text ...
```

Where AlarmName is from the TRAP record and *unique variable-text* is created by substituting selected VarBinds into the Format column of the corresponding record.

- catrapd sends the formatted traps to the Event Management daemon for further processing.

## Enable Automatic Formatting of Traps

A catrapd option lets you enable automatic formatting of SNMP traps. The formatted text contains the most important information from the trap in an easy-to-use form. Predefined formatting is provided for thousands of traps from many companies.

### To enable automatic formatting of traps

1. Select Start, Programs, CA, Unicenter, NSM, Enterprise Management, EM Settings.  
The EM Settings window appears.
2. Click the Component Activation Flags tab at the bottom and the Client Preferences tab on the right.
3. Set the SNMP Trap Server Activated option to YES, and click Yes on the confirmation message.
4. Click the Event Management tab at the bottom and the Client Preferences tab on the right.

5. Set the Format traps using provided tables option to YES, and click Yes on the confirmation message.
6. From the Settings menu, click Exit.
7. Restart catrapd.

Your settings are in effect.

### How catrap Issues Traps

The catrap command can issue SNMP traps to any destination in your network and supports all operands accepted as Open System standards for an SNMP trap command. You can use it interactively, through UNIX/Linux shell scripts or in Windows .bat files, or as part of automated event handling policies defined to Event Management.

- The operands provided as destination and information data to catrap convert automatically into the appropriate Open Systems standard datagram.
- catrap sends the operands to the designated trap destination.

**Note:** This command makes it simple for user applications, shell scripts that are part of production jobs, or Event Management policies to issue their own SNMP traps, simply by executing the command and passing it the appropriate arguments. Unlike some other SNMP trap commands, catrap does not restrict itself to any particular set of ISO or Enterprise MIBs and is totally open for use with any MIB or pseudo-MIB with no dependencies on any third-party network management components.

For more information on catrap, see the online *CA Reference*.

### Binary and Hex Octet String Varbinds

The Windows SNMP service treats octet string varbinds containing binary or hex string data in traps as printable strings unless one of the following is true:

- None of the characters in the octet string is printable
- The conversion to printable data truncates data length

If either of the preceding cases is true, the octet string is displayed in hex.

Octet string varbinds containing binary or hex string data in traps when using catrapmuxd with v1, v2c, and v3 SNMP support are converted to printable strings with a potential for truncation of data in the Console when the octet string contains non-printable data. If this occurs, you can modify the `aws_snmp.cfg` file to specify that certain varbind OIDs are always displayed in hex. This ensures that the octet string data is displayed fully, in hex, on the Console.

## TRAP and MIB Table Manipulation

The TRPCNTRL command lets you manipulate all or selected records in the TRAP and MIB translation tables.

**Note:** This command applies only to Windows.

### Disable/Enable TRAP Records

To disable a single TRAP record, use the following syntax:

```
TRPCNTRL disable trap e=Eid g=Gen s=Spc
```

For example, the following command disables the automatic formatting of the Link Layer Operational trap:

```
TRPCNTRL disable trap e=1.3.6.1.2.1.10.16 g=6 s=1
```

To enable the TRAP record shown previously, enter the following command:

```
TRPCNTRL enable trap e=1.3.6.1.2.1.10.16 g=6 s=1
```

Enable or disable multiple TRAP records by using wildcards in the keywords. For example, the following command disables all TRAP records that have an Eid column beginning with 1.3.6.1.4.1.199.1:

```
TRPCNTRL disable trap e=1.3.6.1.4.1.199.1.*
```

### Notify CATRAPD of Changes

After modifying TRAP or MIB records, issue the TRPCNTRL refresh command to notify CATRAPD of the changes.

### List MIB or TRAP Records

To list MIB or TRAP records of a specific MIB or group of MIBs, use the following syntax:

```
TRPCNTRL list mib: m=<mib-name/mib-mask>
```

```
TRPCNTRL list trap m=<mib-name/mib-mask>
```

For example:

```
TRPCNTRL list mib: m=RFC1157
```

```
TRPCNTRL list trap m=RFC*
```

```
TRPCNTRL list trap enabled=N
```

For more information about TRPCNTRL, see the online *CA Reference*.

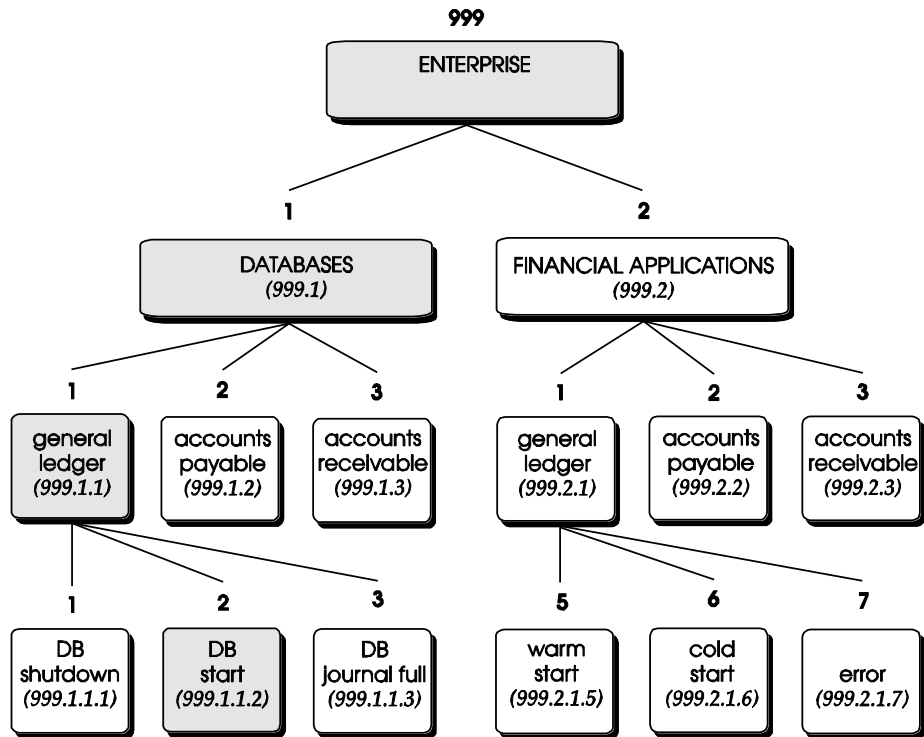
### MIBs

A MIB (Management Information Base) is the numeric code that identifies an event and includes other data as necessary to describe the object affected by the event. It is essential that no two vendors use the same MIB number to describe different events, so standards exist to organize MIBs into one of three broad categories.

- Industry Standard MIBs are sanctioned and published by the International Standards Organization (ISO).
- Enterprise MIBs are assigned by the Internet Assigned Numbers Authority (IANA) to a given organization and are reserved for the exclusive use of that organization.
- Pseudo-MIBs are not sanctioned or assigned by the IANA but can be just as meaningful and useful as an ISO or Enterprise MIB. Pseudo-MIBs often piggy-back on an Enterprise MIB of another organization and take advantage of many of the defaults available on a given platform.

### Sample Pseudo-MIB

The following sample pseudo-MIB describes an event tree. Each element represents information that can be sent when specified as a variable on the catrap command.



Sending a trap of 999.1.1.2 is equivalent to sending the message "The Enterprise Database server that handles the General Ledger database has been started."

A trap of 999.1.1.3 indicates that the General Ledger database has encountered a journal full condition. A trap of 999.2.1.5 indicates that the General Ledger financial application has resumed processing after a temporary outage (warm start).

Taking the example further, assume CA NSM is executing on several nodes, but you want to direct all SNMP trap traffic to a single monitoring machine running on the server Earth. The server Earth receives the SNMP traps. Event Management records and acts on them.

The server Mars runs production financial applications. The General Ledger production application running on Mars terminates with an error.

Testing the return code issued by the General Ledger production executable, the shell script detects an exit code indicating a problem and issues an SNMP trap to alert the server Earth by executing the following command:

```
catrap earth "" "" 6 0 22 999.2.1.7 integer 128
```

where:

**catrap earth**

Sends the identified trap information to the server Earth.

**"" and ""**

Instructs catrap to take the default Enterprise code and the default agent address respectively for this node.

**6**

Indicates that this command is sending a specific trap.

**0**

Identifies the specific trap number for this example.

**22**

Specifies an arbitrary number selected as a timestamp indicator.

**Note:** The following operands identify the variable binding (varbind) information for the trap.

#### **999.2.1.7**

Identifies the object about which information is being sent. In the event tree illustrated earlier, this object refers to an error in the Enterprise financial application, General Ledger.

#### **integer 128**

Provides additional information about the event. In this example, it could mean send an integer value of 128 to node Earth, assuming 128 is an error code that has meaning to the General Ledger application; or it could be the exit code that the shell script detected as indicating an error.

When received at the trap target server Earth, catrapd decodes the event and performs automatic actions in response. The event tree shows other types of events that could be sent, such as 999.1.1.1, indicating that the database of the Enterprise data server for the General Ledger system has shut down.

When combined with other CA NSM capabilities, the possibilities expand. For example, you can use Event Management to intercept error messages from any application and automatically execute customized catrap commands in response. The detection of key events can result in traps being sent in response to files becoming available for processing or applications completing their processing. Security violation attempts can result in other SNMP traps being sent.

On the receiving side of an SNMP trap, you can use Event Management message handling policies to:

- Send warning messages in human readable form to other consoles or terminals
- Issue additional traps to one or more other nodes

For more information on catrap, including an example of how to use it to issue an SNMP trap, see the online *CA Reference*.

## **Event Policy Packs**

Event Management provides preconfigured event policy packs for:

- Message record/action
- Advanced Event Correlation

## Message Record and Action Policy Packs

Message record/action policy packs provide actions to highlight and color-code various agent-related events according to severity. Use `cautil -f scriptname.cautil` to load the policy packs into the MDB. For information about `cautil`, see the online *CA Reference*.

- `caiOraA2_msgrec.cautil` policy provides critical and warning traps/polls for the Oracle database agent.
- `caiSybA2_msgrec.cautil` policy provides critical and warning traps/polls for the Sybase database agent.
- `caiSqlA2_msgrec.cautil` policy provides critical and warning traps/polls for the Microsoft SQL Server database agent.
- `caiAdsA2_msgrec.cautil` policy provides critical and warning traps/polls for the Active Directory agent.
- `caiUxsA2_msgrec.cautil` policy provides critical and warning traps/polls for the UNIX system agent.
- `hpxAgent_msgrec.cautil` policy provides critical and warning traps/polls for the Performance agent.
- `caiLogA2_msgrec.cautil` policy provides critical and warning traps/polls for the Log agent.
- `caiWinA3_msgrec.cautil` policy provides critical and warning traps/polls for the Windows 2003 system agent.

These policy packs are on the installation DVD:

- `DVD\Windows\NT\Policy Packs for Windows`
- `DVD/policypacks for UNIX/Linux`

## Advanced Event Correlation Policy Packs

Advanced Event Correlation policy packs, preloaded into the MDB during installation, provide correlation policies to detect combinations/associations of events and generate correlation events. For catastrophic events, the policies generate alerts into the appropriate alert queue. Each policy contains root cause logic for scheduled outages and node failures.

- AdsAgent (Active Directory Agent) policy has rules to generate an alert into the CA-Applications queue for a domain controller failure.
- caiUxsA2 (UNIX System Agent) policy has these rules:
  - File system failure rule suppresses symptomatic quota, directory, and files on the same file system.
  - CPU rule shows process-specific trap/poll as the root cause and suppresses general CPU traps/polls.
  - Memory rule shows process-specific trap/poll as the root cause and suppresses general memory traps/polls.
  - CPU spike rule detects five critical events within a time period.
- caiWinA3 (Windows 2003 System Agent) policy has these rules:
  - File system failure rule suppresses symptomatic quota, directory, and files on the same file system.
  - CPU rule shows process-specific trap/poll as the root cause and suppresses general CPU traps/polls.
  - Memory rule shows process-specific trap/poll as the root cause and suppresses general memory traps/polls.
  - CPU spike rule detects five critical events within a time period.
- caiW2kOs (Windows 2000 System Agent) policy has these rules:
  - File system failure rule suppresses symptomatic quota, directory, and files on the same file system.
  - CPU rule shows process-specific trap/poll as the root cause and suppresses general CPU traps/polls.
  - Memory rule shows process-specific trap/poll as the root cause and suppresses general memory traps/polls.
  - CPU spike rule detects five critical events within a time period.



- caWmiAgent (Windows Management Instrumentation agent) policy has example rules to determine what is possible with Advanced Event Correlation and the caWmiAgent:
  - Terminal services rule correlates the number of sessions and users to virtual memory.
  - Locked-out user rule correlates locked-out users to application failures. Applications may fail due to incorrect or obsolete credentials.
  - Device failure rule shows a fan failure as the root cause of other device failures.
- Ora2agent (Oracle database agent) policy has these rules:
  - Catastrophic failure rule generates an alert to the CA-Database queue.
  - Memory rule correlates Oracle agent memory monitoring to Windows/UNIX/Linux system agent memory monitoring and shows the more specific Oracle trap/poll as the root cause.
  - Disk space rule correlates the Oracle agent tablespaces events to Windows/UNIX/Linux system agent disk space and shows the more specific Oracle trap/poll as the root cause.
  - Multiple database failure rule looks for three or more failure of any kind on a particular database instance.
- Sqla2agent (Microsoft SQL Server database agent) policy has these rules:
  - Catastrophic failure rule generates an alert to the CA-Database queue.
  - Memory rule correlates Microsoft SQL Server agent memory monitoring to Windows/UNIX/Linux system agent memory monitoring and shows the more specific Microsoft SQL Server trap/poll as the root cause.
  - Disk space rule correlates the Microsoft SQL Server agent tablespaces events to Windows/UNIX/Linux system agent disk space and shows the more specific Microsoft SQL Server trap/poll as the root cause.
  - Multiple database failure rule looks for three or more failures of any kind on a particular database instance.
- Db2agent (DB2-UDB database agent) policy has these rules:
  - Catastrophic failure rule generates an alert to the CA-Database queue.
  - Memory rule correlates DB2 agent memory monitoring to Windows/UNIX/Linux system agent memory monitoring and shows the more specific DB2 trap/poll as the root cause.
  - Disk space rule correlates the DB2 agent tablespaces events to Windows/UNIX/Linux system agent disk space and shows the more specific DB2 trap/poll as the root cause.
  - Multiple database failure rule looks for three or more failures of any kind on a particular database instance.

- syba2agent (Sybase database agent) policy has these rules:
  - Catastrophic failure rule generates an alert to the CA-Database queue.
  - Cpu rule correlates Sybase agent memory monitoring to Windows/UNIX/Linux system agent memory monitoring and shows the more specific Sybase trap/poll as the root cause.
  - Disk space rule correlates the Sybase agent tablespace events to Windows/UNIX/Linux system agent disk space and shows the more specific Sybase trap/poll as the root cause.
  - Multiple database failure rule looks for three or more failures of any kind on a particular database instance.
- Job Management Option policy looks for correlations within the Job Management Option with these rules and generates an alert to the CA-Scheduling queue for critical events that go unresolved for a time period.
  - Job submission problems rule generates an alert to the CA-Scheduling queue for jobs submitted but not started and for jobs started but not completed within a certain interval.
  - Autoscan problems rule detects an autoscan (and/or pre-scan) started but never completed.
  - Predecessor warnings rule highlights warnings that are uncorrected after a time period.
  - SQL errors rule generates an alert to the CA-Scheduling queue for critical SQL errors.
  - Multiple failures rule detects multiple failures on a given Job Management Option.

## Wireless Message Delivery

The Wireless Messaging System provides a way to communicate with operators who are away from an Event Console. Two types of messaging protocols are available: email and pager.

You can send and receive pages from two-way pager devices using the SMTP/POP3 mail messaging protocol. An incoming message can trigger a series of message actions that you specify.

You can define Event Management policies for sending and receiving pager and email messages by using the Wireless Messaging Policy Writer GUI on Windows. The Policy Writer lets you do the following:

- Specify the text of the incoming message that triggers the pager or email response. The message can include environment variables like &NODID and substitution variables like &1, &source, and &severity.
- Define up to three pager or email messages to be sent during the Event Console action sequence.
- Define the text of the pager or email message and up to six possible replies to that message.
- Administer the Wireless Messaging address database and set the path used for storage and workspace.

To secure operations, warning messages from the Event Console are assigned a unique identifier and must have a correctly formatted reply before any action is taken. When a response is received or the page has timed out, the identifier is expired and cannot be reused.

### How Wireless Messaging Works

To send a message, the host machine must have access to SMTP and POP3 servers and the Wireless Messaging server must be running. This is how Wireless Messaging works:

- The Wireless Messaging client executable, `capagecl`, uses two files to generate outbound messages and interpret replies. The message file lists responses that are sent with a notification message and gives directions to the Wireless Messaging client. The configuration file applies a numerical return code to each active response sent to the Event Console.
- A message is created on the Wireless Messaging client from the message file, sent to the Wireless Messaging server, and then to a remote pager.
- The pager operator chooses a response from the list, adds the ID# of the original message to it, and replies to the message.
- The Wireless Messaging client receives the response from the Wireless Messaging server and translates it to a return code defined in the configuration file.

## Wireless Messaging Client - capagecl

The Wireless Messaging client executable, `capagecl.exe`, formats messages and sends them to the server. When a reply is expected, the client executable waits for a reply, and matches the ID number of the page it sent with the text of a reply.

This executable has four command line options that specify the following:

- Message address
- Text of the message
- Message file (`.rep`) containing replies to be sent to the remote device
- Configuration file (`.cfg`) used to interpret responses

The Wireless Messaging client performs some additional formatting and administrative tasks. These are set by entries on the command line or by directives in the message file.

For detailed descriptions of command line options for `capagecl`, see the online *CA Reference*.

For instructions about sending one-way and two-way messages from the command line, see the online help.

## Message File

Wireless Messaging creates messages from information in the message file. Messages are composed of fixed text, environment variables like `&NODEID` and substitution variables like `&1`, `&source`, and `&severity`.

The message file may include a list of pre-formatted replies expected from the message recipient. The processing of replies, however, is independent of the message file contents. The recipient may send additional or alternative replies, and these replies are resolved into actions if the replies are included in the configuration file and if policy actions specify how to handle the additional return codes.

Besides directives to the Wireless Messaging client (`set xxx=yyy`), text written to the message file is appended to the message specified on the command line and included in the sent message. The format of the reply text depends on the device to which the message is sent. The wireless messaging client recognizes replies delimited by three underscore characters, as in "`___Reply___`" though this formatting may be transparent on remote devices.

For information about formatting the commands embedded in the message file, see the online *CA Reference*.

## Configuration Files

Configuration files store addresses, audience information, message groups, and a list of replies with their associated return codes.

When a message arrives at the CA NSM mailbox, the message server opens it and searches for an ID code. If this code matches the code expected by the Wireless Messaging client, the server passes the message to that client. The client processes the text of the message and looks for a reply. If a reply is found, the client checks the appropriate configuration file to find the code it should return to the calling application.

The Reply Information Configuration file then maps responses to return codes. In the following sample configuration file, Acknowledge is mapped to return code 97.

```
#
SendKeep Text= 30
Page Someone Else= 31
Banner Message= 37
Acknowledge= 97
```

**Note:** You can define message actions that send responses based on the return code received in the message.

The Reply Information Configuration file may include any or all responses sent with a message, and can include additional responses that were not sent as suggestions but may be useful to the remote user. Valid return codes range from 6 through 95 and can be unique to each client instance.

**Note:** Return codes 6 through 95 can be assigned to user replies. All 90 return codes can be in any configuration file, but configuration files and standard policy generated by the Wireless Messaging Policy Writer recognize only replies that you list when defining the policy.

Wireless Messaging can interpret other replies if they are added manually to the configuration file or a default file is used. Otherwise, return code definitions can be arbitrary (as long as they are unique) without affecting the behavior of the policy.

The reserved return codes (0-5 and 96-99) are used by the system. The following codes are significant.

Code	Description
03	Message not sent (rejected by calendar check)
96	Abort code (triggered by server termination or <code>capagecl -I <i>issueid</i> -X</code> command)

Code	Description
97	Acknowledge
98	Reply not found in the given configuration file
99	Wireless Messaging client timed out without receiving a reply

For information about the format of configuration files, see the online *CA Reference*.

### Settings for Command Messaging

The Wireless Messaging tab of the Configuration Settings GUI on Windows provides environment variables for Wireless Messaging.

- Set the Server security mode and Inbound command authorization code to let the Wireless Messaging server accept commands from remote devices. Messages sent to the server are processed and passed to the Event Console as oprcmds.
- Set Autoreply to unsolicited messages to YES if you want to acknowledge receipt of validated and rejected pages. This option also tells the server to produce a warning message when a reply arrives for which no client is waiting.
- Set Server security mode to L to send a command page using unformatted email. This option relaxes security. Command pager messages have the following format:
  - The first line specifies the layout of your message, and is used for the automatic answering process. The layout strings are the same as for the capagecl command.
  - The second line contains only the unencrypted password.
  - The third line contains the literal text of the oprcmd to be issued from the Event Console.

For information about setting environment variables, see the online *CA Reference*.

**Note:** For security reasons, the password is contained in an encrypted portion of the command message. Using the PageNet® Pagewriter application, this process is transparent, although the password must be entered at the time the page is generated.

## Wireless Messaging Policy Writer

The Wireless Messaging Policy Writer on Windows lets you create one-way and two-way message policies by building policy templates. The policies can contain the following information:

- The text of an inbound message, which can include all CA NSM substitution variables like &1, &source, and &severity.
- Triggers for remote messaging sequences
- Recipient addresses
- Up to six replies for each message
- Return codes that identify replies

The Wireless Messaging Policy Writer provides persistent storage, so that you need not define email addresses, message layouts, default timeouts, and standard groups of messages for each message.

**Note:** You can send messages to up to three recipients.

For more information about the Wireless Messaging Policy Writer, see the online Help.

## Template Files

Template files, which are modified Event Management script files, provide the basis for Wireless Messaging policy. An example of a template file is one named Single Notify.

These files contain command actions for starting the pager client and sequences of conditional GOTOs for trapping return codes. Many of the details, such as the capagecl command lines and condrc= numbers, are supplied to template files by your entries in the Wireless Messaging Policy Writer. Entries in the file that are specific to each policy are replaced with flag entries beginning with [DUMMY...]. For a full list of these substitution flags, see the information on template files in the online *CA Reference*.

Some policy templates are supplied with Wireless Messaging, and you can create new template files by copying from those files.

## View the Wireless Messaging Icon

When you start the Enterprise Management (EM) Classic Interface, you see icons for the various components like Calendar Management and Event Management. When you click the icon for Event, more icons appear for the functions of Event Management like Console Logs and Message Actions. Note that the Wireless Messaging icon appears only on local servers where an Event Manager is installed. You cannot see the icon on remote computers where only an Event Agent or remote admin client is installed.

## Alert Management System

The Alert Management System (AMS) is a tool for organizing and tracking the most important events in an enterprise or a logical segment of an enterprise. It lets you focus on and manage the highest severity IT events.

AMS provides tools for defining alert policy and multiple panes in the Management Command Center for viewing alerts. AMS also lets you link to the following CA products:

- Unicenter Service Desk, which is a customer support application that manages calls and IT assets, tracks problem resolution, and shares corporate knowledge.
- eHealth Suite, which delivers comprehensive fault, availability, and performance management across complex, heterogeneous systems and application environments. eHealth collects a wide variety of data from your network infrastructure to generate alarms and reports.

**Note:** For more information about AMS, see the guide *Inside Event Management and Alert Management*.

### What Are Alerts?

Alerts are sophisticated messages that provide important or critical information about the state of your enterprise. They are generated by Event Management policy, and although they are functionally similar to Event Management held messages, they can be organized in more complex ways and provide the following powerful capabilities to resolve problems:

- Alerts have many properties related to their importance, display characteristics, age, acknowledgement, and more. These properties let you organize and view alerts in ways that benefit your enterprise.
- Alerts provide automated actions and menu actions in context to the type of problem the alert represents, thus helping you respond rapidly.
- Alerts are automatically linked to affected WorldView objects and optionally linked to Service Desk requests or eHealth alarms to simplify tracking and resolving situations.
- Alerts can be annotated, and each time-stamped entry is locked after it is added to ensure the integrity of your operations.
- Alerts can be automatically consolidated into a single parent alert when they are similar or identical. Consolidation eliminates confusion and console clutter when the same problem is reported in multiple similar alerts.



- Alerts are automatically escalated based on several of their properties to ensure that they receive attention quickly. Escalation can increase alert urgency, transfer alerts to another queue, set alarms, send notifications, and more.
- Alerts have a detailed audit trail that includes information about automated and manual actions that affect them or are carried out on their behalf so that you always know the actions taken to resolve them.

## How Alert Management Works

The Alert Management System (AMS) works with Event Management System (EMS) and Advanced Event Correlation (AEC) to capture the most important events so that you can respond to them quickly. Alerts are displayed on the Management Command Center (Unicenter MCC).

This is how AMS works:

- Alert profiles assign properties to alerts when they are created, escalated, transferred, and so on. The first thing you should do is define the objects that assign the properties. These objects include:
  - Alert classes, which organize alerts and supply most of their initial properties. Classes specify which of the other profiles are associated with alerts.
  - Alert queues, which specify how alerts are grouped for viewing on the Unicenter MCC. Alert queues can be for departments like Accounting, Research and Development, and Marketing. Alert queues can also be for your WorldView business processes. Alerts in each queue are shown in separate viewers on the Unicenter MCC.
  - Alert Global Definition, which specifies properties that apply to all alerts and to the entire Alert Management System.
  - Display attributes, which specify how alerts look on the Management Command Center. Examples of attribute properties are text color, blinking text, and more.
- Event Management message policy determines the conditions that prompt alert creation. Message records and actions for alerts use the ALERT message action. You assign initial properties for alerts by indicating the alert class in the message action.

**Note:** Alerts should be a small subset of the events that occur. They should represent only events that require human intervention or provide information critical to continued normal operations. We recommend no more than 1,000 alerts per day. There are two reasons for this. First, if few alerts are generated, the operations staff can focus more easily on what is important. Second, AMS is a complex system and each alert consumes more computing resources than other events. By carefully designing your AMS configuration and policy, you can help ensure that you get the most benefit from AMS.

- Advanced Event Correlation can generate correlation alerts directly into the Management Command Center. The alert class is specified at the engine level of the policy and each rule can enable or disable alert generation to that class.
- Alerts are shown by alert queue or managed object in the Management Command Center:
  - The queues you have defined are listed in the left pane when Alerts is chosen from the drop-down list above that pane. When you select a queue in the left pane, the alerts in that queue are shown in the right pane. You can open multiple queues in the right pane, and lock them in place as you move to other areas of the Unicenter MCC.
  - A bar chart of alert statistics is displayed when you right-click a node in the left pane and choose Viewers Status. The chart shows the total number of alerts for the node broken down by queue and priority.
  - Alerts for a managed object in the Topology view are displayed when you right-click the object and choose Alert Viewer from the context menu. Periodically, an association daemon polls the alert table for unassociated alerts and links them to their origin node.
  - The context menu that opens when you right-click individual alerts in the right pane lets you acknowledge alerts, view their properties, transfer them to another queue, and more.
- AMS provides a connection to Unicenter Service Desk, which is a customer support application that manages calls, tracks problem resolution, shares corporate knowledge, and manages IT assets. You can open and resolve Service Desk trouble tickets without leaving the Unicenter MCC. Besides viewing trouble tickets from AMS, you can also view them for managed objects in the Topology view.
- AMS also integrates with eHealth Suite, which delivers fault, availability, and performance management across heterogeneous systems and application environments. Based on policy that you deploy, eHealth alarms and netHealth exceptions create alerts automatically. When an alert or alarm is closed by either AMS or eHealth, the corresponding alert or alarm is also closed. You can display eHealth At-a-Glance and Alarm Detail reports for a selected alert from the context menu or the My Actions menu in the Unicenter MCC.

## Understanding Alert Policies

Alert profiles specify properties that are assigned to alerts. The objects that you can define are:

- Alert classes
- Alert queues
- User data
- Alert global definition
- Escalation policies
- Action menus
- User actions
- Display attributes

**Note:** When you define profiles for alerts, start with display attributes and move up the list until you reach alert classes.

## Alert Classes

Alert classes organize alerts and specify their initial properties. Classes are groups of alert profiles like queue, escalation policy, and display attributes. Classes make it easy to define alerts because properties are automatically given to alerts in each class. You do not have to specify all alert properties manually.

Besides linking alerts to the other objects like queue and escalation policy, classes also specify properties not defined elsewhere. These include:

### **Urgency**

Urgency indicates how soon a technician or operator should try to resolve the situation that caused an alert. Because a situation can become more urgent or less urgent as time passes, you can change the urgency manually or with escalation policies after an alert is generated.

### **Impact**

Impact indicates how much an event affects your business. A consideration in determining the level of impact is how many users are inconvenienced by a situation.

### **Priority**

Priority is a value calculated by multiplying urgency and impact. Priority provides an additional way, besides urgency and impact, to evaluate the severity of an alert.

### **Consolidation**

Consolidation groups alerts that have the same alert class, alert queue, node of origin, and alert text. Consolidated alerts appear on the Management Command Center as one alert that has a number indicating how many similar alerts are grouped together. When alerts are consolidated, fewer messages appear on the Management Command Center. Also, if the alert creates a Service Desk request, only one request is opened.

### **Alarm**

Alarm is a property that indicates an alert should be acted on promptly. Alarmed alerts attract more attention than other alerts because an Alarm dialog is displayed on your desktop.

### **Calendar**

(Optional) A calendar indicates the dates and times when alerts in a particular class can be created.

### **Expiration Date**

(Optional) Expiration date is a date when alerts in a particular class will no longer be created. On this date, the class is deactivated.

## **Alert Queues**

Alert queues are groups of similar alerts. Queues are usually based on role, function, department, geographical location, or other category that is meaningful to your enterprise. For example, your queues may be for departments like Accounting, Finance, and Research and Development. AMS organizes alerts by the queues that you define, and the Management Command Center displays alerts for each queue separately in the right pane.

Alert queues are like console logs because they group messages. The difference is console logs group messages by day, whereas queues group alerts by whatever you specify.

## Other Alert Profiles

Besides class and queue profiles, alerts have other profiles that you define as objects.

### **Display attributes**

Display attributes specify how alerts appear on the Management Command Center. Available attributes include foreground and background color, blinking text, reverse video, and more. Display attributes give alerts a consistent look and help limit possible display combinations.

### **User Actions**

(Optional) User actions are designed to run commands automatically in response to alerts, and manually from the context menu in the Management Command Center. User actions can run scripts, executables, commands, and so on. User actions are optional.

### **Action Menus**

(Optional) Action menus are custom submenus for the context menu in the Management Command Center. These submenus list commands that you defined as user actions. By customizing the context menu, you can run commands quickly, without having to open a command prompt.

### **Alert Global Definition**

The alert global definition specifies properties that apply to all alerts and to the Alert Management System as a whole.

### **User Data**

(Optional) User data is site-specific information that is added to an alert. Examples of user data are customer name, contact information, location of a server, hardware owner, and so on. When an alert is created, a user exit function attaches the user data. All alert properties are available to the exit function to determine the content of the user data.

### **Escalation policies**

(Optional) Escalation increases the attention that an alert receives. You specify the situations that cause escalation, like age of the alert, time since acknowledgement, and number of duplicate alerts. You also specify what should be done when escalation occurs, like running a command, changing the appearance of the alert on the Management Command Center, and increasing the urgency of the alert.

## Message Policy for Alerts

Message records for alerts are the same as for any other type of Event Management message. You can define new message records for alerts or use existing ones.

Message actions for alerts, however, are different from actions for other events in the following ways:

- The Alert class id field ties the message action to AMS.
- The action on the Action page is ALERT.

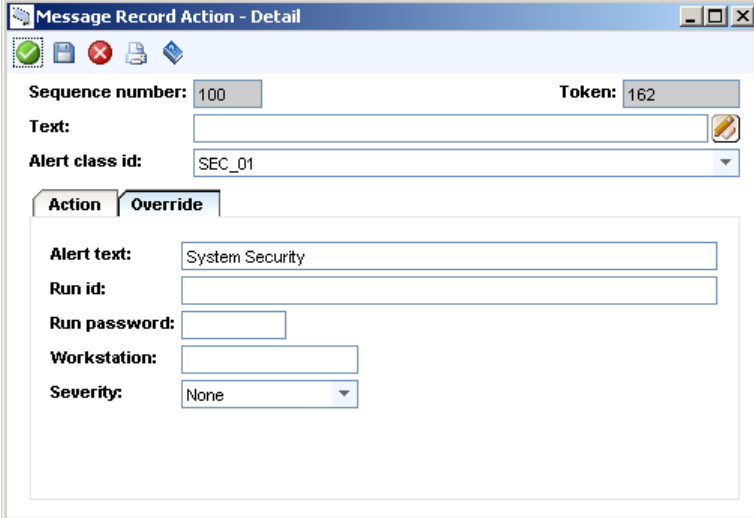
The screenshot shows a window titled "Message Record Action - Detail". At the top, there are fields for "Sequence number: 100" and "Token: 162". Below these is a "Text:" field with a text editor icon. Underneath is a dropdown menu for "Alert class id:" set to "SEC\_01". There are two tabs: "Action" and "Override". The "Action" tab is selected and contains several fields: "Action:" with a dropdown set to "ALERT", "Attribute:" with a dropdown set to "DEFAULT", "Color:" with a dropdown set to "DEFAULT", "Conditional opcode:" with a dropdown, "Conditional return code:" with a text box containing "0", and "Node:" with an empty text box. To the right of these fields is an "Options" section with five checkboxes: "Active" (checked), "Evaluate" (checked), "Simulate" (unchecked), "Quiet" (unchecked), and "Synchronous" (unchecked).

- An Alert text field becomes available on the Override page. This field is necessary if you want to consolidate alerts. Consolidation groups alerts that have the same alert class, alert queue, node of origin, and alert text. Consolidated alerts appear on the Unicenter MCC as one alert that has a number indicating how many similar alerts are grouped together. You enable consolidation for a class on the Alert Class - Detail, Limits page.

Alert text can be up to 80 characters of text, variables, tokens, or any combination of these. The variables and tokens are the same ones used for Event Management. See *Using Variables to Enhance the Current Action*. Examples of alert text are:

- "Workstation Agent" or "Eastern Region"
- "&nodeid" to consolidate all alerts for a node. All events for that asset are assigned to the same Service Desk request.

- "Issues for &9 assigned to &10" where the ninth word in the alert message is an office name and the tenth is the person assigned to resolve the critical situation. All alerts for that office and that person are consolidated.



The screenshot shows a window titled "Message Record Action - Detail". It contains the following fields and controls:

- Sequence number:** 100
- Token:** 162
- Text:** (empty text box)
- Alert class id:** SEC\_01 (dropdown menu)
- Action/Override tabs:** The "Override" tab is selected.
- Alert text:** System Security
- Run id:** (empty text box)
- Run password:** (empty text box)
- Workstation:** (empty text box)
- Severity:** None (dropdown menu)

- The Workstation field on the Override page provides a way to run a user action (command) that you defined in Alert Management. Specify the name of an Event Manager node that forwarded the alert. This lets you take corrective actions directly on the node that manages the node where the alert originated. For more information, see User Actions.

The Unicenter MCC shows the manager node in the Route Node column.

**Note:** Alert Management provides a way to define message records and actions quickly when you define alert classes. For more information, see Define Alert Classes.

## AEC Policy for Alerts

AEC can generate correlation alerts directly into the Alert Management component of CA NSM. The alert class is specified at the engine level of the policy, and each rule can selectively enable or disable alert generation to that class.

## Viewing and Responding to Alerts in the Management Command Center

The Management Command Center displays alerts by alert queue or for a managed object. This may occur automatically based on the current Unicenter MCC perspective or your previous Unicenter MCC activity. You can also display alerts manually by selecting a viewer from the context menu.

### By Alert Queue

Alert queues can be for departments like Accounting, Research and Development, and Marketing. Or, they can be for geographical areas like Eastern and Western regions. You define the queues that are meaningful to your business. Alerts in each queue are shown separately from alerts in all other queues.

When Alerts is selected from the drop-down list above the left pane of the Management Command Center, a list of alert managers appears. You can expand each alert manager to display alert queues. The left-pane list lets you display alerts in the right pane for each queue individually, or a status bar chart for a selected manager.

- Alerts are displayed in the right pane when you select an alert queue in the left pane. Alerts in that queue are shown in the Alert Queue Records right-pane view.
- A context menu lets you perform many actions involved in resolving alerts. By right-clicking an alert you can acknowledge it, transfer it to another queue, open a Service Desk request, and more.
- A bar chart of alert statistics for all queues appears in the right pane when you right-click a server in the left pane and choose Viewers, Summary. The chart contains bars that represent each queue. The length of the bars shows the relative number of alerts in each queue, and the color represents the highest priority for alerts in that queue, as indicated by the legend.

### For a Managed Object

Alerts associated with a managed object in the Topology view are displayed when you right-click the object and choose Alert Viewer from the context menu. Periodically, an association daemon polls the alert table for unassociated alerts and links them to their origin node.



## Integrating with Unicenter Service Desk

CA NSM provides a connection to Unicenter Service Desk (Service Desk), which is a customer support application that manages calls, tracks problem resolution, shares corporate knowledge, and manages IT assets.

The integration with Service Desk is through the Alert Management System (AMS) and the Management Command Center (Unicenter MCC) interface. The connection is installed automatically with CA NSM. Interaction with the Service Desk reduces the workload of the customer support staff because some manual tasks are eliminated.

### How the Integration with Service Desk Works

CA NSM creates Service Desk requests based on policies defined in the Event Management System (EMS), Advanced Event Correlation (AEC), and Alert Management System (AMS). This is how it works:

1. A situation is evaluated by either EMS message records and actions or by AEC rules. If the event is serious, an alert is generated.
2. AMS class or escalation policy determines if a Service Desk request (ticket) is appropriate, and creates one.

**Note:** AMS creates requests only for Service Desk resources that are active. This helps avoid flooding the Service Desk.

CA NSM comes with EMS, AEC, and AMS policy that can automatically create and close Service Desk requests. You can also write your own policy using message records and actions, correlation rules, and alert classes and escalation policies.

This is how CA NSM interacts with Unicenter Service Desk:

- Alert policy definitions specify that Service Desk requests be opened and closed during the life cycle of an alert:
  - Open a Service Desk request when an alert is created. Indicate this using the Alert Class Window.

**Note:** AMS does not open a request if an existing request has identical summary, description, and asset properties. This prevents multiple trouble tickets describing the same root problem.
  - Open a Service Desk request when an alert is escalated. Use the Escalation Policy Editor.
  - Close a request when the alert that opened it is closed or made inactive. Use the context menu in the Unicenter MCC to close an alert; use the Alert Class window Main page or Alert Properties dialog Status page to make an alert inactive.

- Alerts that are associated with Service Desk requests include the request reference number. Likewise, Service Desk requests created by alerts indicate that an outside application opened the request.
- The activity log of a Service Desk request is updated automatically with additional information from AMS when duplicate alerts are created.
- The context menu in the Unicenter MCC lets you interact manually with the Service Desk. You can view requests, open a request, and search the Service Desk Knowledge Tools. For example, when you right-click an alert, you can see requests associated with that alert. When you right-click a managed object in the 2D Map or Topology view, you can see requests for the selected node.

**Note:** When Service Desk requests are opened and closed, a message is sent to the Event Console.

## Scenarios

This section contains examples of situations that could trigger the creation of an alert and a Service Desk request.

### Scenario 1: System Agent on a Critical Server

1. An agent metric exceeds a threshold.
2. An Event Management System (EMS) event is generated.
3. EMS message record policy creates an alert for this event.
4. AMS escalation policy opens a Service Desk request because the alert is open more than 30 minutes.
5. A technician resolves the problem and closes the alert, and the Service Desk request is closed automatically.

### Scenario 2: Third-Party Software

1. Third-party software produces a series of events in the system log indicating a failure.
2. EMS captures the events.
3. AEC policy evaluates the events and creates an alert.
4. AMS class policy opens a Service Desk request immediately.
5. Operations staff resolves the problem and closes the alert. The request is closed automatically.

# Chapter 9: Correlating Important Events

---

This section contains the following topics:

[Unicenter Notification Services](#) (see page 323)

[Advanced Event Correlation](#) (see page 339)

## Unicenter Notification Services

Unicenter Notification Services lets you send wired and wireless messages using various protocols and devices to get the attention of operators or administrators, wherever they are, who must resolve problems or attend to emergencies.

Notification Services is different from Wireless Messaging, which is still available in Event Management. Wireless Messaging lets you send emails and pages.

The available protocols are:

### **Email - SMTP, POP3**

Simple Mail Transfer Protocol (SMTP) is used to send one-way and two-way email messages to various devices, including cell phones. Post Office Protocol version 3 (POP3) is used to receive emails from a mail server.

### **Wireless - WCTP**

Wireless Communications Transfer Protocol (WCTP) uses XML over HTTP and is designed for sending and receiving messages and binary data between wire-line systems and one-way or two-way wireless devices.

### **Page - SNPP**

Simple Network Paging Protocol (SNPP) is based on TCP/IP and offers one-way and two-way pages.

### **Page - TAP**

Telocator Alphanumeric Protocol (TAP) sends pages by modem, and is the oldest one-way paging protocol.

### **Short Message - SMS**

Short Message Service (SMS) is used to send text one-way to cell phones using HyperText Transport Protocol (HTTP).

### **Instant Message - Sametime**

IBM Lotus Instant Messaging and Web Conferencing (Sametime Instant Messaging - SIM) is used on Windows to send one-way, and two-way instant messages.

### Voice - TAPI

Telephony Application Programming Interface (TAPI) is used on Windows to send one-way voice messages that are synthesized from text using the Microsoft Speech Application Programming Interface (SAPI) text-to-speech (TTS) engine. The default speech is set in the Windows Control Panel. The messages travel by telephone line using a TAPI-compliant telephony device to a human recipient.

### Script

Third-party or customer programs or scripts can be used to send one-way messages. Scripts and command definitions are stored in the file UNSConnections.ini in the *install\_path/config* directory.

## How Unicenter Notification Services Works

Unicenter Notification Services keeps track of all notifications that you send. This is especially important for two-way notifications that must be matched with responses. Here is the process:

1. You create a notification message by using one of the following features:
  - User interface
  - Command line or script
  - Event Console by right-clicking a message
  - Event Management NOTIFY action
  - Alert Management escalation
  - Application using the Notification Services client SDK
2. Based on the recipient, provider, or protocol information in the request, the Notification Services daemon (unotifyd) selects a protocol-specific driver to send the notification.

**Note:** The daemon runs as a service on Windows and as a background process on UNIX/Linux.
3. The daemon assigns a tracking ID, which it returns to the command or program that sent the notification.

**Note:** If the daemon stops and then restarts, it also restarts the outstanding notifications stored on disk.
4. If a response was requested, the daemon checks for it periodically from the service provider.

5. The daemon stores information about the notification on disk, and updates that information throughout the life cycle of the notification. This is called checkpointing. Updates include:
  - The request is created.
  - The service provider received the notification.
  - The provider delivered it.
  - The recipient read it.
  - The recipient sent a reply.

## Features of Unicenter Notification Services

Unicenter Notification Services provides an SDK, a user interface, a recipient and provider registry, the `unsutil` command line utility, reports, configuration, and commands.

### Notification Services Client SDK

The Client SDK provides libraries, include files, and samples using C and C++. It contains functions and methods to send one-way and two-way notifications, retrieve status, and administer recipients and providers. The Recipient plug-in SDK lets you create plug-ins to resolve recipient information in your directory services. For more information, see the *Programming Guide* and online *CA SDK Reference*.

### User Interface

Notification Services provides windows and dialogs in the classic Windows GUI that let you send messages, view message status, and define recipients and providers. To use the classic GUI, install the Administrative Client for Enterprise Management on a Windows machine.

### Event Message Actions

Event Management provides the message actions NOTIFY and NOTIFYQ that let you define message policy for sending notifications and getting notification status.

## **Recipient and Provider Registry**

The Notification Services recipient and provider registry lets you enter information about recipients, recipient groups, and service providers. It contains recipient addresses, protocols, and connection information so that you need not enter everything manually with each notification. You just enter the name (alias) for the information you want to use. This saves time and may hide sensitive information.

The registry has files that you can edit with a text editor, the Windows GUI, or the `unsutil` command-line utility described later in this section.

Each file contains an explanation of the file's contents and includes sample templates to help you define your own recipients and providers. The files are:

### **`uns_provider.ini`**

Defines provider aliases with connection information for the protocols that service providers support. See the topic [Connection Information](#) for details about what is required for each protocol.

### **`uns_recipient.ini`**

Defines recipient aliases and their default providers. Each recipient has one default provider.

### **`uns_recipient_group.ini`**

Defines recipient groups.

### **`uns_recipient_address_link.ini`**

Defines recipient addresses for each provider.

## **Recipient Plug-in**

External recipient registries can be queried by the provided LDAP recipient plug-in or user-developed plug-ins. The file `uns_source.ini` defines the recipient plug-ins available to the Notification Services daemon for recipient resolution.

When a recipient alias cannot be found in the recipient registry (`uns_recipient.ini`), and a recipient plug-in is configured, active, and successfully loaded, the daemon tries to resolve the recipient alias in the plug-in.

The file `uns_source.ini` provides samples for the default installed LDAP servers:

- `uns_rcp_ldap_eTrust.ini` -- CA eTrust Directory
- `uns_rcp_ldap_sun.ini` -- Sun Java Directory
- `uns_rcp_ldap_domino.ini` -- IBM Domino LDAP Server
- `uns_rcp_ldap_novell.ini` -- Novell eDirectory
- `uns_rcp_ldap_msad.ini` -- MS Active Directory on Windows 2003 Server

Before you activate the source, customize the corresponding configuration file according to your environments. The files have comments that explain the values you can change.

### **unsutil**

The `unsutil` command-line utility lets you define, alter, delete, and list recipients, groups, providers, and addresses. This utility provides facilities similar to the user interface, and the syntax is similar to `cautil`. For more information, see the online *CA Reference*.

### **Reports**

Statistical reports from the `unsutil` command-line utility display the following types of information: summary, provider, protocol, recipient, sender, and error.

### **Notification Services Daemon Configuration**

Some features of the Notification Services daemon can be configured to reflect the way you use Notification Services in your enterprise. For example, you can specify whether the daemon should create a transaction log and what that log contains; indicate whether to store information about notifications on disk (checkpointing); and enter a default sender name for several protocols. To customize the daemon, update the file `UNSDaemon.ini`. This file has comments that explain the values you can change. Also, see the procedure *Configure the Notification Services Daemon* in the online help and *CA Procedures*.

## Commands

The Unicenter Notification Services commands are:

### **unotify**

Sends a one-way or two-way notification message using the Notification Services daemon. If a reply or other status is requested, the command waits for the status and displays it when received. If the wait times out, you can use the uquery command for future queries of that notification.

### **unotifys**

Sends a one-way or two-way notification message on the local node without using the Notification Services daemon. This command lets you send notifications when the daemon is not running. The unotifys command does not store notification information on disk because the daemon is not running.

### **uquery**

Requests the status of one notification or all notifications from the Notification Services daemon. For a one notification, you can display the current status immediately, or wait until a requested status, like a reply, is received.

### **uquerys**

Requests the status of one notification sent by unotifys on the local node. It does not use the Notification Services daemon, so you can use this command when the daemon is not running.

### **unscntrl**

Starts, stops, and queries the status of the Notification Services daemon.

### **unconfig**

Encrypts and decrypts a configuration file. Some connection information in the file uns\_provider.ini requires a user name and password that you may want to protect. Only Notification Services applications can read an encrypted file.

**Note:** Before changing data in an encrypted file, decrypt it. After changing the file, encrypt it again.



## Configuration and Diagnostics

This section provides information that may help with the proper configuration and diagnostics of Notification Services protocol drivers and service providers. The following protocols are discussed:

- [Email - SMTP/POP3](#) (see page 330)
- [Wireless - WCTP](#) (see page 331)
- [Page - SNPP](#) (see page 331)
- [Page - TAP](#) (see page 332)
- [Short Message - SMS](#) (see page 335)
- [Instant Message - Sametime](#) (see page 331)
- [Voice - TAPI](#) (see page 336)
- [Script](#) (see page 338)

### Notification Services Usage of SSL

In order to support the secure SNPP, WCTP, and SMSHTTP protocols, Notification Services requires and installs the OpenSSL component (version 0.9.7d).

The SSL usage indication (Y/N) is included as part of the connection string for SNPP and WCTP.

**Note:** (Windows) Make sure that older versions of these libraries (ibeay32.dll, and ssleay32.dll) are not in the System or Windows directories.

## Email - SMTP/POP3 Protocol Issues

This section provides configuration and diagnostics information for the SMTP/POP3 email protocol. Consider the following items:

- Some firewall and antivirus software may block ports 25 and 110. The notification appears to time out. Check with your software provider to find out how to configure the software so that the Notification Services daemon (unotifyd) can access these ports.
- Because of site-specific mail server names, user IDs, and passwords, the EMAIL two-way notification is disabled after installation. You must configure POP3 in the UNS\_Daemon.ini file MAIL01 section. Uncomment by removing the semicolon for lines ;ResponseName=POP3 and ;ResponsePath=uns\_smtp. Also uncomment and enter the configuration information for your mail server at line ;ResponseConnectionInfo=servername.company.com:110:user:password.
- Set up a separate mailbox to receive responses to notifications. We do not recommend using a personal mailbox.
- For 2-way notifications sent to cell phones, we recommend using the -o sms parameter of the unotify command, the check box Cell Phone Device (SMS) on the Notification Window, Options page, or the SDK function nscli\_attr\_in\_protocol\_options. These items embed the unique message identifier in the message text instead of the subject line, where it is located by default. Because of the limited screen space on cell phones, the subject line may be truncated. In addition, a cell phone may not include the subject line in the reply. The Notification Services daemon uses the subject to associate a unique ID with the message, so if the reply does not contain the subject, the daemon will discard the reply.

These options cause a prompt asking the message recipient to include the ID in the reply:

```
Please include *2387* in reply.  
Service xxx is down, call...
```

**Note:** IDs are small integers and different daemons could possibly assign the same number to different messages. Therefore, each daemon should monitor a dedicated mailbox to avoid mismatched replies.

## Troubleshooting

- You can diagnose errors with SMTP/POP3 using the smtp\_session.log and pop3\_session.log files in the Notification Services log directory.
- Use "TELNET host PORT#" to verify if ports 25 or 110 are blocked.

### Wireless - WCTP Protocol Issues

This section provides configuration and diagnostics information for the WCTP wireless protocol. Consider the following items:

- The WCTP protocol has been tested only with Skytel and Arch.
- Use of international character sets is limited by the pager hardware.
- Multiple Choice Response options appear on a pager depending on the hardware. Responses are not limited to the MCR choices sent.

### Instant Message - Sametime Protocol Issues

This section provides configuration and diagnostics information for the Sametime instant message protocol. Consider the following items:

- The Domino/SameTime server and the Notification Services daemon (unotifyd) cannot be separated by firewalls or proxies.
- The default SIM port is 1533.

### Page - SNPP Protocol Issues

This section provides configuration and diagnostics information for the SNPP paging protocol. Consider the following items:

- By default, most firewalls block the default SNPP port (444). Check with your software provider to find out how to configure the software so that the Notification Services daemon (unotifyd) can access these ports.
- Some SNPP providers may use other ports for SNPP. For example, Skytel uses port 7777 for regular paging and 7778 for SSL secured paging. Check with your provider for details.

## Page - TAP Protocol Issues

This section provides configuration and diagnostics information for the TAP paging protocol. Consider the following items:

- Because TAP is modem-based, it has to operate with the different modems available now, and may require special Hayes AT initialization strings to operate correctly. When to use an initialization string depends mainly on the quality of the phone line, baud rate, and how well the modem handles the phone service (with its current settings). A modem may work fine on one phone line with no initialization string but not on another. The higher the baud rate the greater the chance an initialization string is required.
- For modem parity the TAP specification recommends 7 bits, even parity, 1 stop bit (7E1), however the specification also notes that not all providers adhere to this. The most likely alternative is 8 bits, no parity, 1 stop bit (8N1). The TAP specification does not give average numbers, but the most likely maximum message length ranges from 80 to 256 characters. We recommend that you check with the provider.
- The following US providers have been tested with Notification Services TAP:
  - SkyTel
  - Arch
- The following modems have been tested with Notification Services TAP:

**Note:** ATI3 is the modem driver version and ATI6 is the chipset type. Modems listed with an init string require that the string be set for a successful connection.

  - Boca Modem 33.6  
ATI3: V2.05C-V34\_ACF\_DS1  
ATI6: RC336DPFSP Rev. 44BC  
Init string: AT&F&C1&D2&G0-C0%C0%E2\N391=13
  - Boca Modem v.34 28.8  
ATI3: V1.000-V34\_DS  
ATI6: RC288DPi Rev 04BC  
Init string: ATQ0E1F1N1W1\K5S37=11S82=128S95=47X4
  - Hayes Accura 56K + FAX  
ATI3: V1.120HY-K56-DLS  
ATI6: RC56DPF L8570A Rev 35.0/34.0

- LASAT SAFIRE 288  
ATI3: LASAT Safire 288 V1.43C  
ATI6: RCV288DPi Rev 05BA
- Lucent LT Winmodem  
ATI3: LT V.90 Data+Fax Modem Version 6.00  
ATI6: 6.00,0,19,11C1,0448,1668,2400
- MICA (V.90/K56FLEX/FAX/V.110) Hex Modem Module (installed in a Cisco AS5300 series router)

ATI3: Cisco MICA Hex Modem Module Product Information

- Country Code           001  
V.90, K56FLEX 1.1, V.34+, V.32terbo, V.22bis, V.42, MNP2-4,  
V.42bis, MNP5,  
Fax, V.110, SS7\_COT, TRACE, VOICE
- HEX modem index       00  
CP code revision        2.7.2.0  
CP revision date        May 30 2000
- SP code revision       2.7.2.0  
SP revision date        05/30/2000 (MM/DD/YYYY)

ATI6: Returns error, instead looked at ATI4

ATI4: Cisco Mica V.90/K56FLEX/FAX/V.110

- Unknown name internal modem  
ATI3: V1.301-V34\_DP  
ATI6: RC288Dpi Rev 05BA

**Note:** The modem works fine on UNIX/Linux. On Windows it works fine when the TAP protocol uses TAPI, but has problems when TAP uses the COM port directly. Direct COM port access works only when previous use of the modem worked with TAPI, but stops working when the modem is reset.

- ZOOM V.92 External Modem (3049)  
ATI3: Zoom ACF3\_V1.801A-V92 -C Z201  
ATI6: RCV56DPF-PLL L8571A Rev 50.02/34.00

Init string:

```
ATZ&F~AT+IBC=0,0,0,,,,,0;+PCW=2;+PMH=1;+PIG=1~ATE0V1S0=0
&C1&D2+MR=2;+DR=1;+ER=1;W2~ATS7=60S30=0L1M0+ES=3,0,2;
+DS=3;+IFC=2,2;BX4
```

**Note:** The init string was taken from the modem's Windows driver init string. This init string and no init string provided successful connections most of the time, but not always. The ~ character represents a carriage return.

## Troubleshooting

- If the protocol driver reports a timeout error while dialing or after connection, specify a modem initialization string. If the initialization string does not resolve the problem, or if no string is available, lower the baud rate. A higher baud rate may cause handshake problems and more sensitivity to line noise.
- If an error occurs during modem initialization, make sure the modem is connected properly and detected by the operating system:
  - On Windows, choose Control Panel, Phone and Modem Options. On the Modems tab, highlight the modem being used and click Properties. On the Diagnostics tab, click the Query Modem button. The modem should respond with response codes and not return an error. If it returns an error, the modem needs to be configured properly on Windows before it can be used with the protocol driver.
  - On UNIX/Linux, open a terminal program such as minicom. When the program has access to the modem device being used, enter the command ATQ0 and press Enter. The modem should respond with OK. If it does not return anything, the modem needs to be configured properly on UNIX/Linux before it can be used with the protocol driver.
- If your phone network requires a prefix number to dial out, the phone number used in the connection information must begin with this number and a comma. For example, if the dial-out prefix is 9, the phone number would be: 9,18005555555.

## Short Message - SMSHTTP Protocol Issues

This section provides configuration and diagnostics information for the SMSHTTP short message protocol driver. Consider the following items:

- Service providers do not follow any protocol standards for SMS messages via their web form interfaces. The content is specific to each provider. The UNS\_SMSHTTP library is designed to read provider data from UNSConnections.ini for flexibility. The provider data holds information like the web site URL, POST data string, maximum message length, and expected success and response codes. The major challenge is keeping this data up to date. Because service providers may update their web site with new content or move their interface to a different URL, the setting may need to be changed.
- This library has been tested with many service providers such as Cingular, T-Mobile, Sprint, and Verizon in the US, and O2 in the UK.

## Troubleshooting

If there is an error trying to connect to a service provider and you are required to go through a proxy server, make sure the proxy server information and credentials (if required) are set properly. If you are behind a firewall and do not use a proxy server, the port number (usually 80) required by the service provider's website must be opened.

## Voice - TAPI Issues

This section provides configuration and diagnostics information for the TAPI voice protocol. Consider the following items:

- TAPI is used for the telephony services that Windows provides. A program must initialize TAPI and shut it down when finished. TAPI can be initialized many times in the same program or another program. For all the TAPI initialize calls on the entire machine, there must be an equal number of TAPI shutdown calls. When TAPI is no longer in use by any program, the final shutdown does not free all resources in the TAPI service. If UNS\_VOICES is the only program that uses TAPI on the machine, each first initialization and final shutdown combination leaks memory.

Microsoft has acknowledged this problem. The company determined that the changes required would impact the current design and therefore no fix was made. The problem happened on a Windows 2000 Server SP4 machine. Microsoft says that this memory leak does not exist on Windows XP and will not exist in future Windows versions. See the Microsoft knowledge base article titled "A memory leak occurs in the Svchost.exe process that hosts a TAPI in Windows 2000."

A workaround does not fully resolve the leak, but it decreases the rate at which the leak grows. Each time the UNS\_VOICES library is loaded, TAPI is initialized. TAPI is shut down when the library is unloaded and is where the TAPI leak occurs, provided it is the last program that shuts TAPI down. While the library is loaded no leak occurs.

- Telephony cards are highly recommended over voice modems because they can usually process voices and analyze call progress. Notification Services supports the Intel Dialogic D/4PCI telephony card. Any telephony card that is Microsoft TAPI v2 compliant may work as well.

In order for the Intel Dialogic D/4PCI telephony card to work properly with UNS\_VOICES, the Dialogic Generation 2 Telephony Service Provider (TSP) and WAVE drivers must be installed after the card is installed. For information about installing the drivers, see *Install Drivers for Telephony Cards (Voice protocol)* in the online help.

- Windows provides a TAPI UNIMODEM Telephony Service Provider (TSP) that supports voice modems (except NT4, unless the manufacturer supplies their own TSP). Therefore as long as a voice modem works with TAPI, UNS\_VOICES should work. UNS\_VOICES depends on TAPI to report that a connection is established. When a call is made, most voice modem drivers are configured to report a connected state back to TAPI immediately, regardless of whether the call is answered or still ringing. After UNS\_VOICES receives the connected event, it starts speaking the message. Therefore, if a recipient picks up the phone a few seconds later, they either miss the entire message or hear the end of it, depending on the message length.



There is a workaround, but it is not straightforward and the interaction between UNS\_VOICES and the recipient is not seamless. When TAPI makes a call, it initializes the voice modem with settings that are defined by the modem's INF registry setting. Most modems use the voice command settings +VRN and +VRA (although some use #VRN and #VRA):

- +VRN is the "Ringback Never Appeared Timer" whose value is the time the modem waits for the first ring tone to occur before assuming the call has been answered. Voice modem manufacturers usually set +VRN to 0, which means the modem does not wait for a ring tone. This value needs to be greater than 0. A recommended value is 10, but this requires trial and error depending on the phone system.
- +VRA is the "Ringback Goes Away Timer" whose value is the time the modem waits for silence between the one ring tone and the next before assuming the call has been answered. We recommend that you increase this value if it is 0.

You can find descriptions of +VRN and +VRA at [http://www.cisco.com/en/US/products/sw/accesssw/ps275/prod\\_command\\_reference09186a00801f6327.html](http://www.cisco.com/en/US/products/sw/accesssw/ps275/prod_command_reference09186a00801f6327.html).

To change these settings:

1. Launch regedit and expand the following registry key:  
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E96D-E325-11CE-BFC1-08002BE10318}`  
Key enumerations for each modem on your system appear, starting with 0000.
2. Expand the key for the voice modem you are going to specify to UNS\_VOICES.  
**Note:** The key may be hard to find. Try reading the FriendlyName key value for each enumeration.
3. Select the VoiceDialNumberSetup key.
4. Find AT+VRN=0, and change it to 10. If AT+VRA=0, change that value, too.

Voice modems are not designed to detect a human voice, therefore it is not possible to wait until the recipient answers the phone and says something. With this resolution, when a recipient answers the phone, the message may not be spoken immediately because the voice modem uses timers to determine when a ring-tone has not occurred. TAPI reports a connected state only after the voice modem determines there are no further ring tones.

A program cannot set these values because after control is returned, TAPI reads the settings in the registry and therefore overrides any changes a program may have made.

- The following device is not supported:
  - US Robotics Sportster Voice 33.6 Faxmodem with Personal Mail  
FCC ID: CJE-0375  
FCC REG#: CJEUSA 20778-MM-E  
ATI1: D869  
ATI3: USRobotics Sportster Voice 33600 Fax RS Rev. 2.0  
ATI9: (1.0USR0007\\Modem\Sportster 33600 FAX/VOICE EXT)FF

## Troubleshooting

- If an error occurs while using a voice modem or telephony device, make sure the device is connected properly and detected by Windows.
  - For a voice modem, choose Control Panel, Phone and Modem Options. On the Modems tab, highlight the modem being used and click Properties. On the Diagnostics tab, click the Query Modem button. The modem should respond with response codes and not return an error. If it returns an error, the modem needs to be configured properly on Windows before it can be used with the protocol driver.
  - For another telephony device, check the device's diagnostic or sample program (if any) to determine if the device is working properly. Resolve any problems before using the device with the protocol driver.
  - For an Intel Dialogic D/4PCI telephony card, run the sample program installed with the Intel Dialogic System Release software, talker32.exe (Program Files\Intel Dialogic System Software\Sample Programs\TAPI). This must work before you can use the device with the protocol driver.

## Script Protocol Issues

This section provides configuration and diagnostics information for the script customer program protocol. Third-party or customer programs or scripts can be used to send one-way messages. Consider the following items:

- All scripts must be explicitly configured to run with Notification Services. By default, the protocol does not allow execution of arbitrary commands from Notification Services.
- Configuration information for the SCRIPT protocol driver is located in *install\_path/config/UNSConnections.ini*. The file contains a section [UNS\_SCRIPT], which defines the general properties of the protocol driver. Also refer to UNSConnections\_ini\_\*.txt for information and examples.

## Advanced Event Correlation

Advanced Event Correlation (AEC) integrates seamlessly with Event Management to provide a powerful event correlation, root cause, and impact analysis capability. When used with existing CA NSM features, AEC can increase the quality *and* reduce the quantity of the information reported on the Event Console, which is used to automate certain operational tasks.

In simple terms, event correlation is a way to group associated events together for the purpose of further processing. Grouping events in this way lets you do simple but powerful forms of processing, such as event suppression, reformatting, aggregation or consolidation. For example:

- Suppress events according to source, duplication, transient states (for example, a flapping link), frequency, thresholds associated with field values, and so on.
- Combine (aggregate) information spanning multiple events into one event.
- Extract data from events that may be difficult to extract using existing tools, making it available for further processing through automation.
- Reformat events for easier processing or to be more readable for operators.
- Detect the absence of scheduled events, such as a Backup Complete. Event correlation can also facilitate more powerful contextual forms of processing, such as *root cause analysis* and *impact analysis*.
- Enrich events with properties from external sources, such as WorldView managed objects.

Root cause analysis lets you clearly differentiate the root cause event associated with an event stream from the non-root cause or symptomatic events that may not require a direct response. Root cause analysis helps you to reduce the number and frequency of events seen by console operators, eliminate message flooding, and reduce false notifications.

Symptomatic events can provide valuable information about the impact of the root cause problem on the overall system, and, therefore, should not be discarded in all cases. The impact analysis function helps you alert users to an impending problem, thus reducing the load on your help desk. It also helps you to initiate failover or recovery procedures for the dependent systems, or alert operations staff that they need not address a particular problem.

**Note:** On non-Windows platforms, AEC is installed with the Event Manager and Event Agent. On Windows, AEC is a separate component.

## Why Use AEC?

CA NSM reports messages generated by the failure of managed components to the Event Console. Within the Event Console, message records trigger actions for the failure messages reported. However, some of the messages that the Event Console receives can be misleading or unnecessary. These messages can trigger unnecessary actions to fix false or secondary failures.

Examples of misleading or unnecessary failure messages include the following:

- An expected but normally inappropriate state, such as services taken offline for repairs
- Service failures cause dependent object failures
- System failures cause agent failures
- Repetition of a previously received message

These false failure messages cause problems because message records and actions erroneously generate notifications and trouble tickets, and, therefore, important messages may be lost in all the erroneous, secondary, false messages.

Using AEC, you can do the following:

- Distinguish between primary and secondary failure messages
- Determine the root cause of the failure
- Provide an impact analysis of a failure
- Diagnose and filter unwanted messages
- Respond to dynamically changing environments

## How AEC Works

AEC extends the functionality of Event Management. To use AEC, you must first identify a set of events that you want to monitor and correlate, and specify any actions to perform if correlation exists or does not exist. The events to be monitored are reported to the Event Console as messages that act as input to AEC, and are intercepted and analyzed. You configure AEC to act on the input messages it receives to produce the desired output, which are the messages that are actually sent to the Event Console.

AEC uses *correlation rules* to analyze the input messages in relation to each other and to identify the *root cause messages* from those incoming messages. A correlation rule performs the following functions:

- Describes patterns so as to recognize those incoming messages that are related
- Defines timings on when to report root cause messages to the Event Console
- Captures the logic of cause-and-effect relationships of all related messages
- Describes formatting of the root cause messages reported to the Event Console

AEC processes events as follows:

1. Listens to all incoming messages.
2. Uses patterns in the rule to identify the incoming messages that match.
3. Triggers the correlation rule when a matched message is detected.
4. Listens to incoming messages to see if any more messages match the patterns in the rule.
5. Uses timing, specified in the rule, to determine continuation of monitoring.
6. Stores the logic of cause-and-effect relationships of different messages.
7. Identifies which incoming messages are root causes, based on the cause and effect logic.
8. Applies the formatting specified in the correlation rule.
9. Reports the resulting message to the Event Console.

## Alert Management Integration

AEC can generate correlation alerts directly into the Alert Management component of CA NSM. The alert class is specified at the engine level of the policy, and each rule can selectively enable or disable alert generation to that class.

## Event Definitions

Understanding event definitions is critical to understanding AEC, configuring it, and using it correctly.

You can define the two types of events in AEC: input events and output events. Input events are the events that define patterns used by AEC to match messages coming in to the Event Console. Output events are events generated by AEC and sent to the Event Console.

You define these events in the correlation rules when you configure AEC. Each event that you define has a key field, called the *message string*, which describes the event. The message string can contain regular expressions and tokens.

## Configure AEC

Configuring AEC consists of defining correlation rules and saving them in the Management Database (MDB). You can create correlation rules using either the Integrated Development Environment (IDE), which is a Windows application, or the browser-based Policy Editor.

Before you use AEC in a production environment, you must typically work through the following procedures:

1. Define the correlation policy.
2. Deploy the correlation policy.
3. Test the correlation policy.
4. Save the correlation policy in the MDB.

**Note:** You can import CA NSM 3.x rca files into the Policy Editors and then save them to the MDB.

After you define your correlation policies, you can deploy them to a test Event Agent (preferably a non-production machine) using deployment dialogs within the editors. The Windows IDE editor also provides a real-time testing environment by reading the Event Console messages and applying the rules you have defined.

**Note:** You need only use the Policy Editors when you are defining, deploying, and testing rules. After you are satisfied that your new AEC policy is working properly, you can use the Deploy Policy dialog to deploy it into a production environment. See Implement AEC.

The policy editors have a real-time status capability to let you see the following:

- What rules were triggered.
- How many instances are running, and the values of tokens in each instance. For more information, see Tokens.
- The times the rule processing started.
- How much time is left for maturity and reset of the rule. For more information, see Timing Parameters.

**Note:** If Security Management is running, and AEC policy is intended to create alerts in the Alert Management System (AMS), the user defining the policy must have permission to the CA-AMS-POLICY asset type. Without this permission, an access violation message appears.

For more information about CA-AMS-POLICY, see the online *CA Reference* topic Asset Types for Windows and UNIX/Linux. It is under Security Management, Executables: Security Management, until Security Management Control Statements, Control Statements for ASSETTYPE.

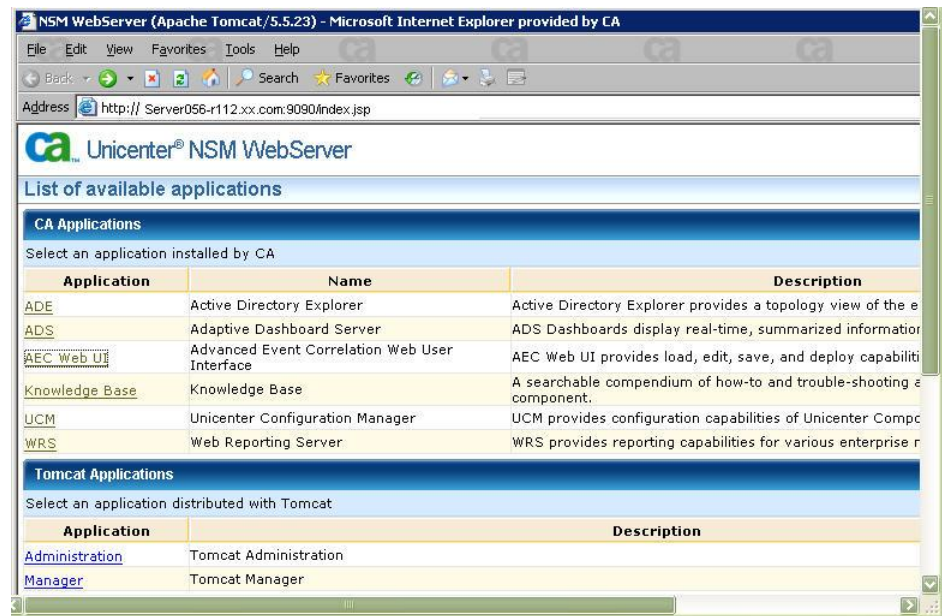
### Start the IDE Policy Editor

To start the AEC Windows IDE Policy Editor, locate the Windows Classic Enterprise Management GUI, navigate to the Event icon, and then navigate to the AEC Policies icon. You can use the resulting list container of AEC policy names to open the IDE for a single policy. You can also open the IDE from the Alert Management Class notebook.

## Start the Web Policy Editor

You can launch the AEC Web Policy Editor a few different ways:

- Locate the AEC Policies tree item in the Management Command Center. The Web Policy Editor opens in the right pane of the Management Command Center.
- Open the Web Policy Editor from Alert Management using the Alert Class - Detail window, AEC tab in the Unicenter MCC.
- Launch through the CA Web Server menu, which displays a list of all CA web applications, including the AEC web policy editor.



Launching the Web Policy Editor independently, outside of Management Command Center, requires some basic configuration. The Web Editor automatically displays a Configuration tab that prompts for the name of the Distributed Intelligence Architecture (DIA) Knowledge Base and the Event Manager host.

**Note:** Launching from within Management Command Center does not require this configuration because Management Command Center already understands these values and passes them automatically to the web policy editor within the right pane of Management Command Center.



## Windows and Dialogs

The Web Policy Editor provides policy wizards as an effective way to create and modify rules without having to manually configuring the many flags and options. You can also deploy policies to Event Agents, determine if the AEC Engine is running, and configure the host names of the DIA knowledge base and Event Manager.

The window contains the following notebook pages:

[New](#) (see page 345) - Create a new policy

Open - Open an existing policy

Edit - Edit a new or existing policy

Save - Save policy to the database

Deploy - Deploy policies to Event Agents

Import - Import an XML policy for editing or saving

Export - Export a policy from the database to an XML file

Status - Refresh and view the status of AEC policy

Utilization - Refresh and view the AEC Engine Utilization Report

Configure - Enter or edit host names of your DIA knowledge base (UKB) and Event Manager

## New Policy Tab

Every configurable AEC object has a configuration wizard associated with it, which guides you through the steps necessary to configure the object. The New Policy tab displays all available policy wizards and lets you select from a list of preconfigured rules. Click the radio button of any wizard you want to work with and then click Select so that you can start defining the incoming event.

The window contains the following Policy Wizards and descriptions:

#### **Missing Event Wizard**

Detects the absence of an important event.

**Example:** When a Database Backup Started event is detected but the Database Backup Completed event is not detected within a specified time range, an alert is sent indicating that the backup failed.

#### **Down for Maintenance Wizard**

Suppresses messages from systems that are down for scheduled maintenance.

**Example:** If software patches that require a reboot are scheduled for a particular machine, you can select an event that indicates that machine is down for maintenance. All messages coming from that machine during the specified time are suppressed.

#### **Transient Event Wizard**

Eliminates spike alarms when a resource has acceptable periods of peak activity.

**Example:** A web server is known to have surges in activity every time new content is posted. The Transient Event rule suppresses alerts caused by these surges.

#### **Suppression of Duplicates Wizard**

Suppresses repeated similar events.

**Example:** A host IP device failure causes DSM to repeatedly generate ping failure events. This type of rule suppresses the redundant events, allowing only the initial failure to trigger a trouble ticket.

#### **Dependency Event Wizard**

Raises an alert for a component based on events raised by other components.

**Example:** A web application requires both a database and a file server. An alert for the web application is sent if either resource reports a problem.

#### **Dual Dependency Event Wizard**

Raises an alert for a component based on two events raised by other components.

**Example:** A web application runs on a cluster consisting of two cluster nodes. An alert is sent if both cluster nodes report a problem within the specified time period.

**Event Threshold Wizard**

Detects the number of times a specific event occurs within a time range.

**Example:** When CPU usage exceeds its threshold five times in two minutes, an alert is raised.

**Root Cause Wizard**

Raises an alert for a component based on events raised by other components and provides information on the event that initiated the issues.

**Example:** A switch failure causes a ping failure, which causes an agent failure.

**Missing Heartbeat Wizard**

Detects the absence of a heartbeat message within a specified time range.

**Example:** A heartbeat message is sent from a server to communicate that it is online. The absence of the heartbeat event from the server indicates that the server is offline.

**User Defined**

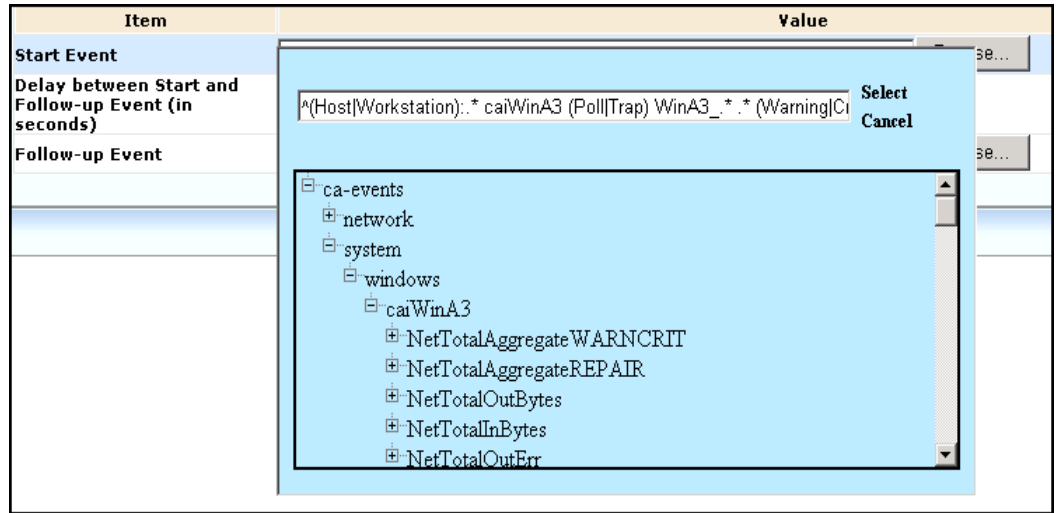
An empty rule list; it lets you manually create a customized rule.

**Event Pick List**

When you enter match events it is sometimes difficult to remember the syntax of an agent trap, security message, or job scheduling event. The event pick list provides a comprehensive listing of CA events, together with their regular expressions. The list is available when you enter any of the input events in the Policy Editors.

### Web UI Pick List Example

The following graphic shows what type of syntax is provided when you use the Browse button to select an incoming event from the drop-down list.



**Note:** For more information about Advanced Event Correlation, see the guide *Inside Event Management and Alert Management* and the online help for any of the AEC help systems.

### Impact Analysis

You can configure AEC rules to generate, in addition to root cause messages, messages associated with the events impacted by the root causes.

AEC analyzes input messages to determine the impact a failure has on a component of a system. AEC responds by sending out impact analysis messages based on its rules. These messages can contain specified substrings from both the root cause and the impacted message. In addition, these impact messages can be sent to the Event Console in the form of an aggregate report, one for each non-root cause.

AEC recognizes a dependency of event A on event B (which is defined in the correlation rules), so you can use it to report impact messages like the following:

- A is impacted because B went down
- B has impacted A
- B has impacted [A1, A2, A3, A4]

For example, an operator shutdown on US-NY-01 has caused a ping failure and an agent failure.

You can use impact analysis to do the following:

- Provide the operators with complete and intelligent information, enabling them to understand and provide notification of the failures in the enterprise.
- Use impact messages to notify the repair personnel to fix the real failures. These messages can also be used to notify users that "false" failure of components has impacted their hardware or software.
- Make infrastructure changes that affect the impacted components, after receiving impact messages.

For example, a router failure has caused a group of applications to fail because they have been disconnected from the database server. After receiving the impact messages, provide an alternate route or use a failover router from the applications to the database server that would bypass the failed router, thereby reducing the downtime of these applications.

- Provide system administrators with a way to measure the impact of failures of hardware and software throughout the enterprise—measuring not only the downtime of the failed component, but the impact of the failure on all affected components.

For example, a failure on a router that is connected to two less critical workstations may not necessitate a repair until hours later. However, a failure on a router that supports hundreds of servers that house an enterprise's main applications, which are accessed in real time by its clients, requires an immediate fix.

## Implement AEC

After you have created, deployed, and tested rules, you can put them into production. The correlation process runs unattended by deploying policy to the AEC Engine, which is installed with every Event Agent and Event Manager.

The AEC Engine runs in the background and processes the events received at the Event Console.

By default, Event Management passes all events to this Engine before doing any of its own processing (that is, message actions, writing to the log, and so forth). You can also configure AEC to process events sent directly to the console log, such as the message action SENDOPER.

**Note:** The Windows IDE Policy Editor processes events *after* they are sent to Event Console, whereas the Engine processes them beforehand. So, when AEC is configured with reformatting or suppression features, these features work only when using the Engine.

## Deploy Policy

You can deploy policy using either the Deploy Policy dialog or the `ace_reload` command line utility. Both facilities let you select any combination of policies from the MDB and deploy them to any combination of Event Agents.

**Note:** Do not use the Windows IDE Policy Testing function and deploy policy to the AEC Engine at the same time to process rules. If both the IDE Engine and the Engine are running at the same time on the same machine, duplicate messages appear in the Event Console.

## Check the AEC Engine Status

You can check the AEC Engine status across multiple nodes and platforms from a single location by using the Policy Editor's Deploy Policy dialog. The same dialog lets you pause and resume any combination of Event Agent engines.

## Check Policy Status and Utilization

Diagnostics in the engine track each policy, the number of times it is triggered, when it is triggered, the rules within the policy that are triggered, and a number of other parameters. This diagnostics data is stored internally and available for reporting later. This provides a convenient way to monitor the engine and to ensure that policy is being invoked.

Utilization statistics are also tracked and reported. For better performance, you can optimize those policies that are frequently invoked and remove policies that are never invoked.

## Event Log Player

The Event Log Player utility plays back prior Console Logs so that you can optimize the AEC policy. The player resembles a video cassette recorder. You must enter the log directory and starting date and times as well as a target Event Management node. Clicking Play causes all prior log messages that meet the criteria to be resent to today's console on the designated node. In this way, you can use historical events to provide simulation for the AEC rules. Other features allow for a real-time playback mode in which the messages are sent with their original frequency.

The Event Log Player installs with the Windows AEC only. To start the Event Log Player, enter the following at a command prompt:

```
C:\Program Files\ca\sharedcomponents\ccs\wvem\ace\EmEvtLogPlayer.exe
```

## Understanding the AEC Components

This section briefly explains the components within AEC that help you analyze events to identify the true issues within a series of related events. For a more detailed understanding of these components, see *Inside Event Management and Alert Management* and the online AEC help systems.

### Components of a Correlation Rule

The two types of correlation rules that you can add to your policy are as follows:

- An *Event pipeline rule*.
- A *Boolean logic rule*, which supports complex nested Boolean logic conditions, helps you analyze and identify complex patterns of behavior.

**Note:** For more information, see Boolean Logic in AEC Rules.

### Event Pipeline Rule Components

The components of an event pipeline correlation rule are as follows:

- Pipeline
- Root events

#### Pipeline

The pipeline is where most of the logic of cause-and-effect relationships of messages is defined. Each correlation rule has one pipeline listing the pipeline items, each of which contains descriptions of similar messages. Each pipeline item deals with only one message type. You group pipeline items to form a pipeline that has a cause-and-effect relationship among the items. The order of the items in a pipeline is important, as any item is considered to be the root cause of all the items below it.

When AEC receives many messages that are matched by different pipeline items, it chooses the highest item and determines that message to be the root cause. For example:

```
Pipeline Item # 1: Ping Failure on Server  
Pipeline Item # 2: Service Critical on Server
```

The Promote/Demote feature lets you modify the order.

The main components of pipeline items are as follows:

#### Match Event

This component indicates conditions (message string and node name) under which the item triggers the rule.

### **Local Correlation Event and Local Reset Event**

The Local Correlation and Local Reset Events describe the message strings that are sent to the Event Console at maturity (reset) of the correlation rule. In this way they are similar to the Root Correlation and Root Reset Events.

However, you can configure AEC to use either the Local or the Root Correlation (Reset) message by setting one of two flags. The flags Use Root Correlation Event and Use Root Reset Event let the Root Correlation and Root Reset messages override the Local Event Messages.

The advantage of setting the Root Correlation (Reset) message is that, for example, you must configure the correlation (Reset) message at only one place. However, the disadvantage may be that, regardless of the root cause, AEC generates the same formatted message (although, using tokens, it can be specialized to reflect the root cause event in each case).

The disadvantage of setting the Local Correlation (Reset) message is that you must configure these messages at each of the individual pipeline items. This lets you configure different messages to be sent to the Event Console when you have different root causes.

### **Exclusion Event**

You can use the Exclusion Event with the Match Event to help restrict the events that match the matching element. For example, you could define a Match Event to match any events containing the text ABC but exclude any events also containing the text DEF.

If you defined the following events in a rule, all Application Failure events are matched except for those that refer to lab applications.

Match Event: `^Application .* has failed$`

Exclusion Event: `^Application LAB_APP1|LAB_APP2 has failed$`

You can also use the Exclusion Event to restrict the matching of any element of an event. For example, you could use it with the Match Event to match a given event from all servers except those specified in the Node field of the Exclusion Event.

### **Reset Request Event**

The Reset Request Event lets you reset an individual pipeline item.

Set the Enable Local Reset Request Event flag to reset an individual pipeline item. In addition, this lets you decrement the counter for the number of matching events associated with a pipeline item when a Reset Request Event is received. When you set this flag, a pipeline item is reset only when the counter is decremented to zero.



For example, suppose that you have five automated procedures that generate consecutive events to indicate that they have started and completed successfully. Using this flag, you can match the five start events and decrement the counter by assigning the Reset Request Event to the completion event. If the matching element has not reset at the end of the maturity period, one or more of the automated procedures must have failed to complete, and the rule can generate a Root Correlation Event to indicate that.

### **Local Reformat Event**

Configured at the rule or matching element level, you can use the Reformat Event to change the format of a matched event. The reformatted event can consist of the following:

- All, or any element of the original event (using &TEXT or &1 - &n, respectively)
- Any global or user-defined token value
- Static text

For example, suppose that you want to prefix any event that matches Pipeline Item # 1 with the string %AEC\_HOLD. This prefix could then be identified by a standard Event Management message record/action, resulting in the event being placed in the Held Messages queue.

### **Local Revised Correlation Event**

It is possible that a higher pipeline item can be matched after a correlation event has been generated (for example, where the rule matures before the highest pipeline item is matched). In that case, you may want to generate an event indicating that the previous correlation event has been superseded. A Revised Correlation Event can consist of the following:

- All, or certain elements of the original root cause event (using &TEXT or &1 - &n, respectively)
- All, or certain elements of the new root cause event (using &RCTEXT or &RC1 - &RCn, respectively)
- Any global or user-defined token value
- Static text

For example, if Event B was initially determined to be the root cause but was subsequently replaced by Event A, you could generate the Revised Correlation Event "Event A has been replaced by Event B as the root cause for Problem X" using the template "&TEXT has been replaced by &RCTEXT as the root cause for Problem X."

### **Reset Request Acknowledge Event**

The Reset Request Acknowledge Event can be generated whenever a rule or pipeline item resets in response to a Reset Request Event.

### Local Impact Event

If the rule is configured to generate impact events, the pipeline item Use Root Impact Event flag is set to false, and this is not the root cause item, this output event is generated to the Event Console after maturity to report the events impacted by the root cause.

## Root Events

You can define root events to override pipeline events. Individual root events are defined as follows:

### Reset Request Event

You can configure this input event to match incoming events that trigger the rule to reset automatically, rather than waiting for the duration of the reset period.

### Root Reformat Event

You can configure this output event to reformat events that matched the item if the pipeline item Reformat Matched Event is set to TRUE, and the Use Root Reformat Event flag is set to TRUE.

### Root Correlation Event

This component describes the message that identifies the root cause event to be sent to the Event Console at maturity of the correlation rule.

### Root Revised Correlation Event

This output event is generated to indicate that a new root cause has been identified in the following circumstances:

- The rule level Enable Revised Root Cause Event flag is set to TRUE.
- The new root cause pipeline item Use Root Revised Correlation Event flags is set to TRUE.
- The pipeline item is matched after maturity and is higher than the current root cause item.

### Root Impact Event

This component describes the message to be sent to the Event Console for each of the impacted messages. This message could contain components of the root cause message as well as the impacted messages. You can also use event-by-event impact messages, or aggregate impact messaging.

**Note:** For more information, see Impact Analysis.

### Root Reset Event

This component describes the message to be sent to the Event Console when the correlation rule is reset.

**Note:** For more information about resetting a rule, see Timing Parameters.

### Root Request Acknowledge Event

This output event is generated to acknowledge receipt of the request to the Console if the rule has been reset using a Reset Request Event.

## Boolean Logic Rule Components

The components of a Boolean logic correlation rule are as follows:

- Boolean operators
- Root events

## Boolean Operators

Each Boolean rule can have one or more nested Boolean operators, each with one or more pipeline items. These Boolean operators let you establish complex relationships among messages. When AEC receives many messages that are matched by different pipeline items, it performs the logical Boolean operations to determine if all conditions have been met.

For example, assume you define a rule that contains the following components:

```
Boolean Operator AND has been selected.  
Pipeline Item # 1: Disk I/O Usage Critical  
Pipeline Item # 2: Database Backup Starting
```

When AEC detects both events (Item # 1 AND Item # 2) occurring within the maturity period, it generates a correlation event. In this example, you may want to stop the database backup to save disk I/O.

The following components of Boolean operator pipeline items are the same as a pipeline rule:

- Match Event
- Exclusion Event
- Reset Request Event
- Local Reformat Event
- Reset Request Acknowledge Event

**Note:** The Local Correlation Event, Local Reset Event, Local Impact Event, and Local Revised Correlation Events are not available in a Boolean pipeline item, because all pipeline items must be considered together in a Boolean rule. Use the root versions to generate these output events.

## Root Events for Boolean rules

The following Boolean rule root events are the same as a pipeline rule:

- Reset Request Event
- Root Reformat Event
- Root Correlation Event
- Root Reset Event
- Reset Request Acknowledge Event

## Boolean Logic in AEC Rules

AEC provides support for complex nested Boolean logic conditions, which lets you use Boolean logic on events. In addition to defining patterns and capturing matched events, AEC can perform Boolean logic on whether matched events have occurred and make a corresponding determination on whether to send out correlated messages when the Boolean logic returns a TRUE condition.

Boolean logic can be applied by defining Boolean logic rules. Boolean logic rules are different than event pipeline correlation rules. Boolean logic rules do not have pipelines but instead correlate events using Boolean operations.

## Timing Parameters

Each correlation rule has time settings that specify when to report a correlation message to the Event Console once the rule is triggered, and when to stop processing the rule after it is triggered.

## Tokens

You can use tokens in correlation rules. A token is similar to a substitution parameter and can be recognized by the preceding ampersand character (&). For each Event field, any tokens are replaced by their actual correlation rule values (if they exist), otherwise they will be replaced by a word wildcard, that is, any value will match.

AEC supports the following types of tokens:

- Internal tokens
- User-defined tokens

## Internal Tokens

Internal tokens already exist in AEC and can be referenced without the need to define them. Internal tokens are also referred to as built-in tokens.

AEC internal tokens closely match the tokens available in Event Management message records and actions, and can be used to identify predefined event variables. In addition, there are tokens specific to AEC.

See the online help for descriptions of AEC internal tokens.

## User-Defined Tokens

User-defined tokens can be included in any field of an incoming matching message. User-defined tokens are defined as `&(..)`, with the name of the user-defined token in the parentheses.

User-defined tokens can be used to establish a relationship among messages that should be correlated. For example, if you want to relate a ping failure on one server to a service failure on that particular server, you can define a token such as `&(NODENAME)` in the matching string of the two messages.

An assigned token, such as `&(NODENAME)`, parsed from an incoming message, can be reused in an output message. For example, if you enter `&(NODENAME)` in the node field of a pipeline item match message, it is assigned to the first matching message and may be reused as part of an output message, such as a local correlation message.

User-defined tokens can facilitate template rules, that is, a new rule will be triggered for each unique user-defined token assignment. For more information, see [Template Rules](#).

The user-defined token value is assigned by the first matching message, and it does not change until the rule has been reset.

## Global Constants

A global constant is a constant that you define once — manually or by calling an external script, Dynamic Link Library (DLL), or executable — and then use throughout AEC policy. Global constants apply to all rules in a policy. These constants can be used to implement a static text substring in multiple rules. The substring can be changed globally, making it unnecessary to modify many rules manually.

Global constants can be either static or dynamic. The value of static constants can be determined using the fields of an event. A dynamic constant can be configured to use an external script, DLL, or executable to return the constant value. The DLL is loaded periodically, and the specified DLL function is called to retrieve the constant value. In this way, constants that reflect the current state of the dynamically changing enterprise can be assigned.

As with user-defined dynamic tokens, the script or executable invoked must return a string in the format:

```
[input-string]\n[output-string]
```

where *input-string* is the string substituted in input events, and *output-string* is the string substituted in output events.

If you want to write a DLL function then it must be in Microsoft MFC/ATL. The function declaration is as follows:

```
bool DllFunc(CStringList *lpParams, int *nBufSize, CString *lpReturnString);
```

where the parameters are as follows:

**CStringList \*lpParams**

In parameter. A cstring list of all parameters, specified during the creation of the global dynamic constant.

**int \*nBufSize**

Out parameter. Return the length of the lpReturnString here.

**CString \*lpReturnString**

Out parameter. Return the string here. This should be in the format Input\nOutput, as with the executables and scripts.

## Credentials

Dynamic constants and tokens can sometimes contain sensitive data in their command-line parameters, such as user names and passwords. To prevent clear-text passwords from being stored in the MDB or seen by a passerby, the policy editor provides a way to hide and encrypt passwords. You can add a Credential item containing the user name and password, and then reference it in the command-line arguments of the script, instead of the user name and password.

**Note:** The password is always stored and displayed in its encrypted form.

The reference to the Credential item has the following format:

```
&(CREDUSER:CRED1) and &(CREDPASSWORD:CRED1) .
```

You must create a credential before it can be referenced later as a way to hide and encrypt passwords.

## Calendar Support

As with correlation rules, dynamic global constants support the use of CA NSM calendars. If configured, the value of the dynamic global constant is only refreshed when the calendar is active.

## Template Rules

A template rule is a rule that acts as a generic template and lets multiple instances run. The rule should contain user-defined tokens that enable AEC to identify similar (but unrelated) events.

Tokens are set when events are compared against the rule; the tokens take the values of the corresponding items in the event.

When an event occurs that matches the match string in a rule but does not agree with the user-defined tokens, the new event invokes another instance of the rule. This new instance processes with its own token values, and, at the time of its maturity and reset, creates its own correlation and resets messages.

## Regular Expressions

AEC allows the matching of events based on patterns instead of fixed strings. These text patterns are known as *regular expressions*, and they are stored in the match event fields of AEC. AEC uses the Perl-compatible Regular Expression Library for the evaluation of regular expressions.

Regular expressions evaluate text data and return an answer of true or false. That is, either the input string matches or it does not. Regular expressions consist of characters to be matched as well as a series of special characters that further describe the data. The description specified by special characters (also called *meta* characters) can be in the form of position, range, repetition, and placeholders.

Within AEC, rules can be specified for the set of possible events that you want to match. Regular expressions can also be used to split the event apart in various ways. Regular expressions are also used to extract parts of strings and store the parts as tokens.

All fields of the Match Event accept regular expressions, including the following:

- Message Number
- Message String
- Node
- User
- Station
- Severity
- Device
- Job Management
- Process
- User Data
- Category
- Tag
- Source
- Hour
- Day of Month
- Month
- Year
- Day of Week
- Day of Year

AEC correlates only those messages whose node, user, station, message string, and so on, matches what is specified in the match event of a rule, and then triggers that rule.

For a list of regular expressions and their meanings, see the online help.



**Note:** For more information about Advanced Event Correlation, see the guide *Inside Event Management and Alert Management* and the online help for any of the AEC help systems.



# Chapter 10: Improving Systems Performance

---

This section contains the following topics:

[Analyzing Systems Performance](#) (see page 363)

[Performance Scope Usage](#) (see page 364)

[Working with Performance Trend](#) (see page 365)

[Effective Reporting with Performance Reporting](#) (see page 365)

[Charging for Resource Usage with Performance Chargeback](#) (see page 366)

[Data Fundamentals](#) (see page 366)

[Performance Architecture](#) (see page 368)

[Administrative Tools](#) (see page 375)

## Analyzing Systems Performance

IT leaders are striving to be more efficient by aggressively managing hardware, software, and technology resource costs. With so much emphasis placed on maximizing the return on investment (ROI) made in IT infrastructures, it is imperative for organizations to identify under-utilized and over-utilized resources and balance workloads effectively.

Servers often run at low utilization levels. As such, applying appropriate and effective management policies can greatly reduce costs while increasing performance levels. Effective performance information also assists IT managers in justifying new IT expenditures by clearly illustrating when current resources have reached maximum capacity and new resources are required to maintain expected service levels.

Your IT organization must also minimize the impact of unexpected performance problems during daily operations to maintain expected service levels. Effective performance management is needed to quickly analyze problems when they occur and improve responsiveness by isolating their causes. Consistent management policies and techniques across heterogeneous server infrastructures are needed to minimize complexity and reduce the need for specialized administration.

Systems Performance manages the performance of the servers that deliver business-critical IT services using consistent, fact-based management policies that reduce complexity. Costs are minimized by applying platform-independent management policies and techniques across heterogeneous server infrastructures, reducing the need for specialized administrators to manage such infrastructures.

With comprehensive systems platform coverage and support for industry standards, Systems Performance provides a flexible and extensible architecture that simplifies the management of the numerous systems and devices that make up today's complex infrastructures. Its facilities for collecting, analyzing, and reporting performance information simplify performance and capacity trend analysis, and increase IT responsiveness to unexpected problems, ensuring higher service levels. Prepackaged management policies and secure, centralized configuration further simplifies administration and increases IT efficiency, resulting in faster ROI.

## Performance Scope Usage

The Performance Scope application lets you monitor the current performance of your enterprise and perform real-time analysis of any performance problems or outages that occur. Each Performance Scope view provides both real-time and historical data, so you can check the performance of resources now and in the past.

A Performance Scope view displays the performance of several different resources. It also supports multiple concurrent views and allows you to monitor many machines at once.

Besides its graphical analysis features, Performance Scope lets you assign performance thresholds to resources, causing alarms to be generated and actions to occur when a threshold breach occurs.

In addition, Performance Scope has built-in support for both graphical views (for general analysis) and tabular views (for detailed analysis).

Finally, Performance Scope provides intelligent, correlation-based problem analysis to help you determine the root cause of problems.

## Working with Performance Trend

Use the Performance Trend application to observe patterns of activity and resource consumption, identify short-term and long-term usage trends, and analyze the impact of moving workload and traffic across different servers and devices.

Performance Trend uses spreadsheets for flexible reporting of preconfigured or ad hoc requests. Reports can be processed on-demand or in batch during off-peak periods, and you can save them as an .XLS file or publish them as HTML or through e-mail or hardcopy.

You can publish Performance Trend reports to the Unicenter Management Portal or Web Reporting Server for access anywhere.

Like Performance Scope, Performance Trend provides a correlation analysis tool that simplifies the process of determining problem causes. Its correlation capabilities sift through large amounts of performance data from disparate sources to uncover hidden relationships within the data. Color-coded reports clearly depict these relationships.

Performance Trend also includes prepackaged, customizable macros that extend its data analysis capabilities. One such macro determines the minimum, maximum, and average data values for performance baseline analysis. Another is used to continue the data curve to indicate when, in the future, the current performance trend may reach a critical threshold.

## Effective Reporting with Performance Reporting

The Performance Reporting application lets you view charts and tables of historical performance data through a standard Web browser, such as Microsoft Internet Explorer or Mozilla Firefox. You can generate reports dynamically to show any level of detail, from the average performance of a group of servers over an entire year to a detailed analysis of a given application's performance over a specific period in a given day. Using Performance Reporting, you can create either chart-based reports or tabular reports.

You can choose to view performance data for several different types of objects, such as selected applications, databases, servers and devices, or objects in a Business Process View. Performance Reporting also lets you generate executive reports - high-level summaries of the overall performance and health of your enterprise.

Performance Reporting provides powerful analytical capabilities, including baseline generation and threshold breach analysis. In addition, you can schedule reports to run automatically in the future, and save the finished reports in PDF or CSV format.

## Charging for Resource Usage with Performance Chargeback

Use the Performance Chargeback application to identify and report on usage of system resources, and charge end-users for the computing resources and services that they consume. Performance Chargeback is based on Microsoft Excel, so you can display resource accounting data in a format that satisfies many presentation needs.

You can display resource usage data in several chart styles, including bar, pi, line, scatter, and 3-D.

## Data Fundamentals

Systems Performance uses Performance Agents running on each monitored machine to collect data on a wide range of system and database resources, SAP resources, and SNMP-based resources. There are two types of agents:

- **Real-Time Performance Agent (prfAgent)**  
This agent is responsible for the real-time, transient collection of performance data, which it supplies to client applications such as Performance Scope.
- **Historical Performance Agent (hpaAgent)**  
This agent provides facilities for collecting, storing, and managing historical, time-banded data. Where necessary, it can act as a proxy to enable the monitoring of resources from SNMP-enabled hosts or devices that cannot support a Performance Agent directly, such as a network router.

For information on how to install and customize the Performance Agents, see the Inside the Performance Agent manual.

## Real-time Data Gathering

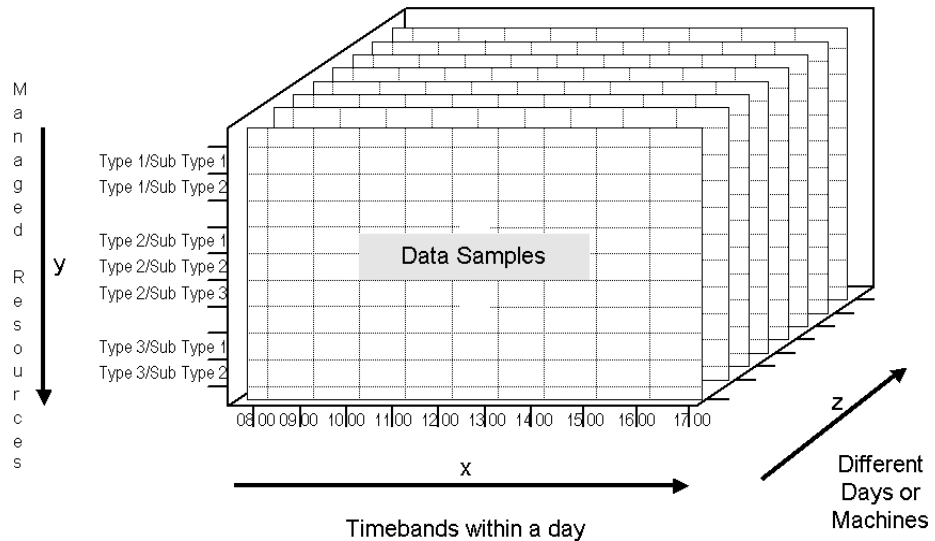
The Real-Time Performance Agent is entirely dormant except when responding to data gathering requests from client applications like Performance Scope. The agent collects the requested metrics only, and it can gather data at a frequency as low as one second.

The real-time performance data is transported over CA's own lightweight, connectionless messaging technology.

## Historical Data Gathering

The Historical Performance Agent collects historical data and stores it in highly compressed files called performance cubes. These cubes are initially stored on the file system of the machine that the agent is monitoring, and are then automatically transferred to a designated Performance Distribution Server.

Cubes conform to a three-dimensional data model, as illustrated below.



The three axes in this data model are as follows:

- Different resource metrics (such as Disk Reads/sec or % Processor Time) are represented on the Y axis.
- Time-bands across the day are represented across the X axis.
- Different days (Monday, Tuesday, Wednesday, and so on) or periods (average day within March, April, June) or machines (machine 1, machine 2, machine 3) are represented on the Z axis.

There are three types of performance cube:

#### Daily

These are a two-dimensional matrix of the Y axis (resource metric) and X axis (up to 24-hour timeband). You primarily use this cube to view how a resource is performing on a given day. Using this daily data lets you closely monitor resources on a real-time basis.

#### Period

These are similar to daily cubes except that they include the Z axis to track same-machine performance over multiple days. For example, you might use a period cube to monitor how a machine has performed over the course of a month.

#### Enterprise

These are just like period cubes except that the Z axis represents different machines during a single day rather than the same machine during different days. For example, you might use an enterprise cube to monitor the performance of a related set of servers.

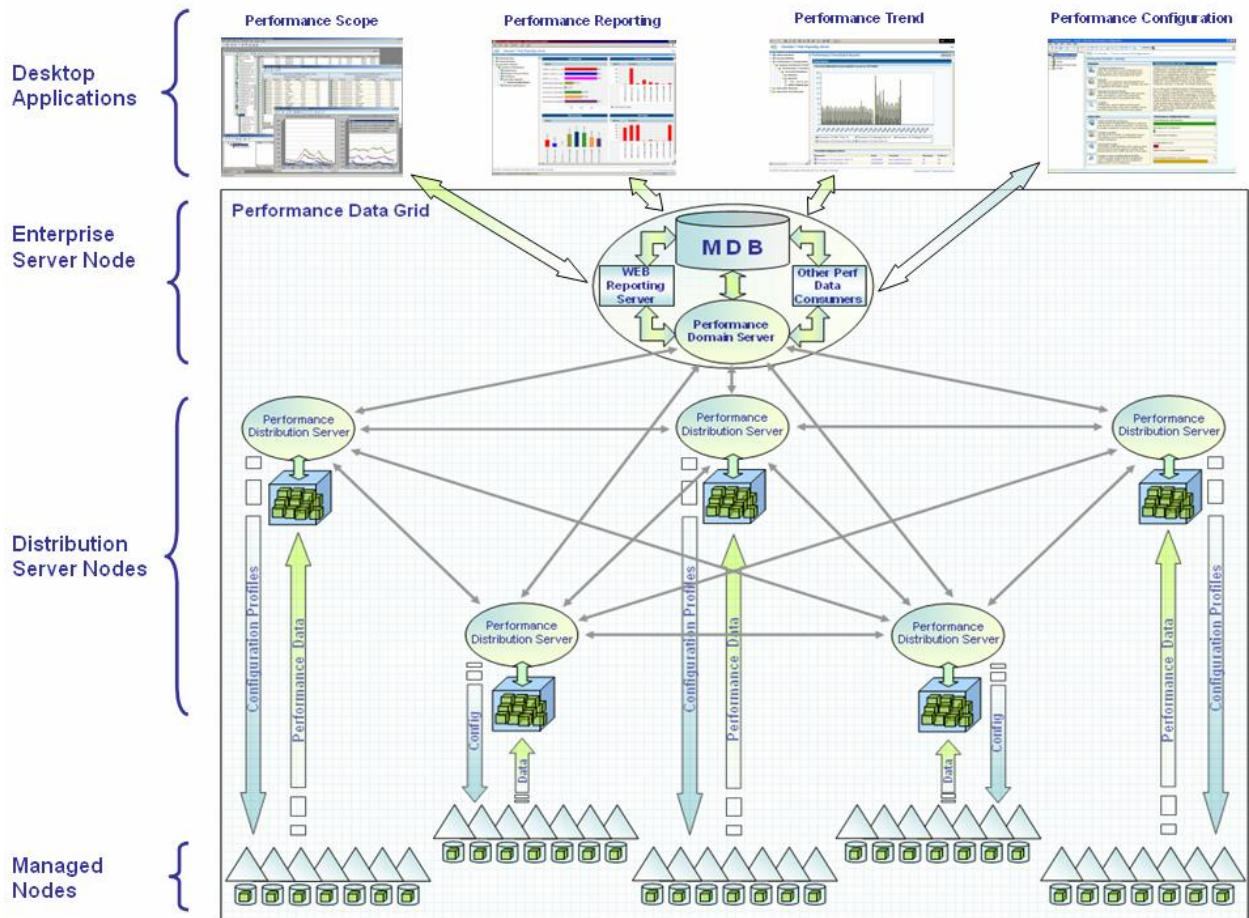
You can use the Performance Configuration application to set up your cube requirements.

## Performance Architecture

The computing facilities in many companies typically comprise a large and diverse collection of machines spread over a wide geographic area. The Systems Performance architecture supports such a distributed environment by providing high levels of scalability and enabling easy configuration across many thousands of machines. Furthermore, because it is often desirable to define logical groupings within the enterprise and manage each of these groups independently, the architecture implements the concept of multiple configuration domains.



The following figure shows the main components in the Systems Performance architecture. See the following pages for further details of these components.



The architecture of Systems Performance has two main functions:

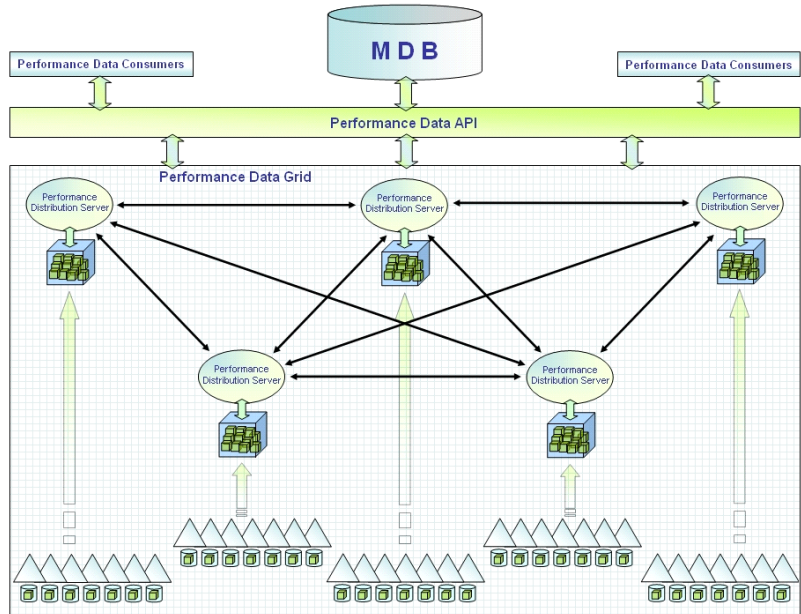
- To provide access to and management of the performance data gathered by the Performance Agents.
- To enable the configuration of the Performance Agents.

## Data Accessibility and Management by the Performance Data Grid

With each Systems Performance domain capable of supporting up to 100,000 monitored servers and workstations, there is a critical need for a mechanism that provides fast and reliable access to the vast amount of performance data that may potentially be gathered over time. In particular, it is important that the data held in performance cubes across the enterprise:

- Can be summarized and stored in the Management Database (MDB) so that you can correlate it with your organizational assets and other management data.
- Is readily accessible to management and analysis applications that need performance data.
- Is reliably marshaled and stored in a highly scalable way, but is centralized enough to support easy data management and recovery.
- Can easily have data management rules applied so that you can age and aggregate the data as it gets older.
- Can be changed into a form that has greater business relevance, such as "health" indexes, workload-based or application-based metrics, and customized metrics that you yourself can create.

The Performance Data Grid (PDG) in CA NSM Systems Performance (r11 and above) provides all these capabilities and more. In essence, the PDG lets you obtain performance information that covers any time period and at virtually any degree of granularity for any managed element in your enterprise (device, server, application, and more).



The PDG is formed from a network of orthogonal, distributed Performance Distribution Servers that form a grid to service data requests. This grid creates a single image of the performance data for the entire enterprise and grants you seamless access to the data. A notable feature of this design is that you do not need to know where the data you are requesting is physically stored, or which end-point is servicing the request; you simply place a query on the grid and obtain a response.

## Configuration Services

Performance Agents (running on the managed nodes) report to the Performance Distribution Servers, which in turn report to a Performance Domain Server. These Domain and Distribution Servers run as persistent services/daemons, so they can react immediately to registration requests from agents and service instructions from the Performance Configuration application.

## Main Performance Architecture Components

### Performance Domain Server

The Performance Domain Server holds all the performance configuration information for an entire domain. Multiple clients can simultaneously connect to the Performance Domain Server and engage in configuration operations.

You can perform configuration operations from any machine on which the Performance Configuration application is installed. And you can even use a single machine to administer multiple domains (but note that a client cannot connect to more than one Performance Domain Server at a time).

## Performance Distribution Server

Performance Distribution Servers request configuration data from the Performance Domain Server and deliver it to the Performance Agents. As this data is centralized at the Domain Server level, Distribution Servers operate without the need for any local persistent information. Therefore, you can install or re-install them without the need to backup and restore data.

A key function of a Distribution Server is to manage performance data for the machines for which it is responsible and maintain this data in its local cube store. Each Distribution Server communicates with the others and is aware of the contents of their cube stores. In fact, to operate efficiently, each Distribution Server must have knowledge of:

- All the performance cubes to which it has access in its local cube store.
- All the monitored machines for which other Distribution Servers have data.
- Which Distribution Server is the primary source of data for each monitored machine.

To achieve this, each Distribution Server examines its local cube store and builds up an index of the machines for which it has cubes, the cubes that exist for each machine, and the resource metrics for which there is data in the cubes. The Distribution Server then passes its index to the other servers with which it is registered so that they all have up-to-date information on each other.

When you use an application like Performance Scope to request performance data, the application submits the request to any available Distribution Server. The server examines its local cube store and then either returns the requested data or, if it is not the most appropriate server to handle the request, forwards the query to a more suitable server.

## Configuring Distribution Server Replication

Each Performance Agent is assigned a manager Distribution Server when it is installed, or first connects to the Performance Domain. One of the responsibilities of the managing Distribution Server is for the Performance Data Engine (PDE) portion to assemble a superset of cube data for machines it manages.

To allow for high availability, the Distribution Server PDE provides redundancy. This is achieved by configuring each Distribution Server to replicate its Performance cubes to another Distribution Server, so that if the primary Distribution Server fails or becomes unavailable, the backup Distribution Server will have all cube data required to answer queries.

Once replication has been configured, the agent automatically delivers the cubes to both the primary and backup Distribution Servers. Additionally, the primary and backup Distribution Servers exchange cube lists, allowing the primary and backup Distribution Servers to pull cubes from each other. This ensures that both the primary and backup Distribution Servers contain a full superset of Performance cubes for machines managed by the primary Distribution Server.

When configuring Distribution Server replication, you must ensure that you use the correct server name in the `cfgutil` command. Failure to specify the full server name will cause the `cfgutil` command to fail and replication will not occur. The following procedure is considered best practice.

### **To configure Distribution Server replication**

1. Issue the following command on the machine you want to use as the backup PDE:

```
camstat -n
```

This command returns the correct server name to specify in the `cfgutil` command.

2. Issue the following command on a Performance Domain Server to configure the backup PDE:

```
cfgutil -P <primary Distribution Server> -b <backup Distribution Server>
```

Replication is configured.

3. Issue the following command to verify configuration was successful:

```
pdectl -h <backup Distribution Server> status
```

Status details for the backup Distribution Server are returned. The Backup for value should be the primary Distribution Server.

## **Summary Cubes**

A Performance Distribution Server also builds a summary cube for each machine for which it is the primary source of performance data. This type of cube is designed to provide fast access to data that spans several days, weeks, or months. Each summary cube contains one year's worth of data for a single machine. However, the data is averaged to a granularity of 24 hours for each of the monitored resources, so a typical cube provides 365 samples per resource. All requests for performance data with a granularity of 24 hours or more are directed to the summary cube.

Note that summary data occupies approximately 10% of all storage space.

## Access to Metadata

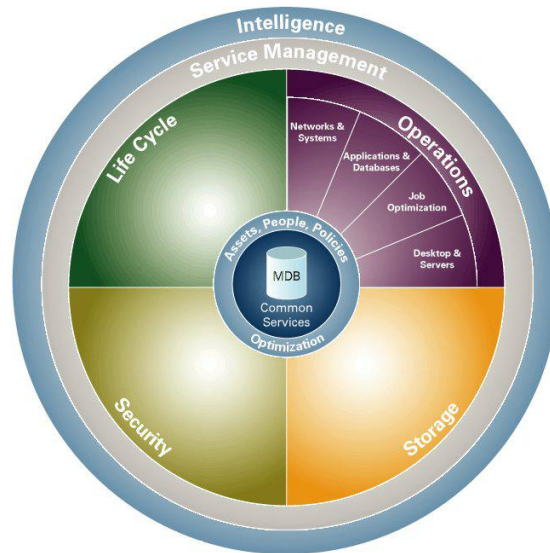
As well as providing access to performance data, each Performance Distribution Server also provides access to metadata-information about the performance data itself. Examples of this metadata include lists of:

- System types
- All the machines known to the entire PDG
- Machines filtered by system type
- Resources monitored on a particular machine

Using facilities such as Performance Reporting, you can submit a request for this metadata. The Performance Distribution Server that receives the request examines its local cube store and then either returns the required metadata or forwards the request to a Distribution Server that is better able to handle the query.

## Systems Performance and the MDB

All Unicenter management functions use the MDB to store information about managed objects, their properties, relationships, and the methods by which they are managed. Examples of managed objects are hardware devices, software applications, databases, and processes.



You can use the Systems Performance tools to publish performance, configuration, and asset data to the MDB.

The Performance Domain Server automatically publishes to the MDB historical performance data that it has obtained from the PDG. Systems Performance tables exist in the MDB schema for this information. Although the data is published to the MDB automatically, the content and granularity are configurable. In addition, the Domain Server enables the publishing of asset and configuration information to the MDB.

You can also use command-line utilities to retrieve performance data from either the PDG or one or more cubes and publish it to the MDB. For more information, see the *Inside Systems Performance* manual.

## Administrative Tools

Systems Performance provides a number of tools for easily and effectively performing configuration operations.

### Secure, Centralized Configuration with Performance Configuration

Performance Configuration is a client-server application that provides extensible configuration capabilities from a central point of control. Use Performance Configuration to create and apply configuration policies (called *profiles*) to Performance Agents. Once a Performance Agent has received a profile, it knows what performance data to collect, when to collect it, and where to send it.

Performance Configuration provides an intuitive Profile Editor with which you can easily create configuration profiles.

You can create and apply new profiles through a drag-and-drop mechanism for dynamic reconfiguration of individual systems or customizable groups.

Performance Configuration also includes a separate One-click Configuration wizard with which you can quickly configure the Performance Agent on an individual machine or device directly from the Management Command Center.

### Command-Line Utilities

Systems Performance also provides a number of configuration commands that complement the graphical Performance Configuration application. These commands include the following:

Command	Function
cfgutil	Communicates requests to Performance Configuration through a command-line interface.

<b>Command</b>	<b>Function</b>
configserver	Starts, stops, and displays the status of the Performance Distribution Server on the local machine.
cubespider	Fetches missing remote cubes.
hpaagent	Controls the Historical Performance Agent.
pcmtocsv	Converts performance cubes to CSV format.
pdectl	Controls a Performance Distribution Server.
pdtodb_m and pdtodb_u	Publishes performance data to a relational database.
pdtoxml	Converts performance data to XML format.
prfagent	Controls the Real-Time Performance Agent.
profilesserver	Starts, stops, and displays the status of the Performance Domain Server on the local machine.
rtpmon	Displays real-time performance data.



# Chapter 11: Creating Customized Reports

---

This section contains the following topics:

[Types of Reports](#) (see page 377)

[Report Templates](#) (see page 378)

## Types of Reports

Web Reports let you view all of your historical performance data through an Internet browser, in a way that is meaningful to you.

You can view Web Reports from one of three sources:

- WRS
- Unicenter MP
- Unicenter MCC

You will encounter three report types in Unicenter MP:

- Configured Reports
- Report Templates
- Published Reports

### **Configured Reports**

Configured Reports are out-of-the-box report templates that have not yet been executed. You can immediately use these to access meaningful information from supported products and view that information in a report. You can create and save these reports using the provided templates, or use the reports provided by the product. Configured Reports are listed in the tree according to product-specific classification.

### **Report Templates**

Report Templates provide a way to customize reports using option fields to fill in the criteria that you want to use to generate your report. They provide all of the possibilities of what you can define. Once you fill in a template, you can either publish the information into the tree as a published report, or simply add it to the list of configured reports in the tree.

For example, if you are concerned about x factor within a supported product, you can execute the x factor Configured Report, which provides summarized information on x factor activity for that product. But if you want to view more specific information on y within the x factor, you can fill in the y within the x factor template provided by that product to define your own configured report, or publish the report into the product tree so you can retrieve it. Web Reports provide several Report Templates and Configured Reports across supported products that let you see your data the way you want to see it.

Report templates are listed in the tree according to product-specific classification.

### **Published Reports**

Published Reports are the static contents of the results of executing reports in Unicenter MP. After you publish a report, you must reload the Knowledge Tree for the report link to become available. You may choose to delete items that are published in the tree after you have viewed them.

The key to Web Reporting is establishing a connection to the Web Reporting Servers using the WRS Catalog page in Unicenter MP. Establishing connections makes it possible for users to view reports running on these servers. The WRS Catalog lists all defined Web Reporting Services (WRS) connections and allows you to add, delete, or manage them.

## Report Templates

Report templates provide a blank slate of fields in which you can fill in the criteria that you want to view in your report. They provide all of the possibilities of what you can define. Once you fill in a template, you can either publish the information into the Knowledge tree as a published report or simply add it to the list of configured reports in the tree. Report templates are listed in the tree according to product-specific classification.

To get you started, Unicenter MP ships with several predefined reports. If needed, you can use the Report Configuration pages to edit these predefined reports to suit your needs.

# Chapter 12: Securing CA NSM Objects

---

This section contains the following topics:

[What is Security Management](#) (see page 379)

[How Security Management Works](#) (see page 380)

[How Security Management Is Implemented](#) (see page 382)

[Phase 1: Customize Security Management Options](#) (see page 383)

[Phase 2: Start Security in QUIET Mode](#) (see page 388)

[Phase 3: Create Rules for Production in WARN Mode](#) (see page 389)

[Phase 4: Set Options for Production, FAIL Mode](#) (see page 399)

[Security Management Reports](#) (see page 399)

## What is Security Management

Security Management provides a policy-based security system that protects against unauthorized access to vital CA NSM components and products.

To protect your management systems, you must augment physical security with software that can do the following:

- Prevent unauthorized individuals from gaining access (logging on) to your systems being managed by CA NSM.
- Ensure that only authorized personnel can access management data and resources.
- Protect against access abuse by users with administrative privileges.
- Provide a way to review and audit the use of data and resources.

Enhanced and simplified Security Management means reduced errors, greater responsiveness, and increased flexibility in meeting the needs of your organization. Most importantly, it means you can implement thorough and effective security policies without disrupting your work environment.

**Note:** The CA NSM Security Management components no longer provide file access authorization. If you need this type of additional security, you may want to evaluate eTrust Access Control. For more information, see *Integration with eTrust Access Control*.

## How Security Management Works

Security Management provides a policy-based layer of security that works with native operating system security to provide enhanced security, including the following:

- Controlled resource access by users with administrator privileges; you can delegate operator or support functions that require administrator authority, such as loading management policy, without allowing the user complete access to your policy records
- Restricted access to critical management systems such that important policy records are not removed or modified

One of the most important advantages of a policy-based security system over access control lists (ACLs) is that systems are protected, not by their physical attributes and ACLs, but rather by security policies you define. These security policies are stored in the MDB.

By configuring default DENY security policy, newly created management policy is protected automatically. This set-and-forget nature of policy-based security is the key to managing hundreds of users on a system as easily as you can manage a dozen.

## Security Policies

All of the asset access controls provided in Security Management are maintained through policies you define in the MDB. Once these policies are set, they are enforced until the security administrator changes them.

Additionally, all access violation attempts are routed to the Event Console Log, providing a real-time window into security activity.

The primary policy definitions used in managing security policies are as follows:

### **User Groups**

User groups logically group users and access permissions together, providing a role-based security policy. Defining user groups is optional and is not tied to the native OS user groups.

### **Assets**

Assets describe specific occurrences of a protected entity, such as an Enterprise Management calendar object. Users can be given access to an asset directly by granting permission to the user, or indirectly by granting permission to a user group of which the user is a member.

### **Asset Groups**

Asset groups describe multiple assets with similar attributes; for example, all the Enterprise Management components to which a user group has CONTROL access. As with assets, users may be given access to an asset group directly by granting permission to the user, or indirectly by granting permission to a user group of which the user is a member.

## How the Commit Process Works

The commit process puts that version of the policies you currently have defined into effect. It also generates the Decision Support Binary (DSB) files. Security Management requires that you execute the commit process during various phases of implementation. The commit process is typically required anytime after you make a change to your policies, that is, define, delete, or change a security access rule.

The commit process, when executed, performs the following tasks:

- **Security Database Access**--The first phase extracts the policies that you have defined in the Security Management Database. Therefore, that database must be online and active.
- **Security Administration (or SDM daemon) Active**--Another phase takes the policies extracted from the database and places them into effect on the designated server, where they are processed by the Security Management functions, which must already be running. If Security Management is not running, the commit process will detect this and issue an appropriate error message indicating that an attempt to place the new security policies into effect has failed.

To perform a commit process, issue the `secadmin` command with option `c` from the `cautil` command line, or select File, Commit from the menu bar in the Security Management GUI.

**Note:** For command syntax, examples, and additional information about warnings and commit customization, see the `secadmin` command in the online *CA Reference*.

## How Security Management Is Implemented

The implementation of a security policy for a server is the result of planning, implementing, testing, evaluating, and adjusting. Understanding the implementation process, including why specific procedures are performed in a particular order, is critical to the success of your Security Management solution.

Security Management is implemented in the following phases:

- **Phase 1**--Review your Security Management Options and Client and Server Preferences, to ensure that needed components are available. Customize options to your site's security requirements.
- **Phase 2**--Start Security Management daemons (service providers) in QUIET mode.
- **Phase 3**--Create security rules for users, user groups, access permissions, and asset groups in WARN mode and verify user access, adjusting rules based on results, then re-execute the commit process. Consider any optional features you may need. For more information, see *Optional Features*.
- **Phase 4**--Customize Security Management options for FAIL mode production operations.

This chapter provides detailed information and points you to procedures that allow you to accomplish each of the phases of implementation.

## Phase 1: Customize Security Management Options

Security Management options that affect enforcement modes, violation rules, and other controls are set globally and are used systemwide. These options have default values that can be customized to your site's security requirements.

Before continuing to Phase 2 of Security Management implementation (Start Security Management daemons in QUIET mode), you must customize the Security Management options so that they are consistent with your enterprise's security requirements.

### How You Modify Windows Security Management Option Settings

You can modify Windows Security Management option settings using the Enterprise Management GUI. For instructions about security option settings, see *Customizing Your Security Management Options* in the online *CA Procedures*.

After you have reviewed your options, see *Customizing Your Options for Phase 2* for platform-specific requirements.

### How You Modify UNIX/Linux Security Management Option Settings

UNIX/Linux platform Security Management options are located in the `$CAIGLBL0000/secopts` file. You can use the editor of your choice to modify the option settings in this file.

If you are upgrading to CA NSM from an earlier version, you may consider refreshing the `$CAIGLBL0000/secopts` file with the new version, `$CAIGLBL0000/secu/scripts/secopts`, so that options can be located by Security Management.

### Options to Consider for Your Operations

This section describes some of the options you can customize to support Security Management. Other Security Management options are described later in this chapter under the following topics:

- Additional Options for the Windows Platform
- Additional Options for UNIX/Linux Platforms
- Setting CA Standard Security Facility (CAISSF) Scoping Options

## Default Permission Option

The Default Permission option (DEFAULT\_PERMISSION for UNIX/Linux platforms) controls the action of Security Management when a user attempts to access an asset. This option can be set to DENY or ALLOW. The default is DENY.

### **DENY**

Directs Security Management to deny access to any asset not otherwise permitted for a given user. DENY results in every system asset being considered *protected by default*. Think of DENY as a way of saying, "no one gets access to anything unless specifically told otherwise."

### **ALLOW**

Directs Security Management to allow access to any asset not otherwise (explicitly) denied for a given user. ALLOW protects only those assets you tell it to protect. Think of ALLOW as giving access to *everything except* those assets that are specifically protected.

## System Violation Mode

The System Violation Mode option (SYSTEM\_MODE for UNIX/Linux platforms) determines whether and how Security Management responds when a user attempts to access an asset. This option can be set to QUIET, WARN, FAIL, or MONITOR (for UNIX/Linux only). The default is FAIL.

### **QUIET Mode**

Ignores user accesses to assets. Security Management invokes no enforcement action.

### **WARN Mode**

Allows access to an asset for which the user has no permission, logs the event to the Event Console Log, and sends a warning message to the user who originated the access violation. WARN logs the event.

### **FAIL Mode**

Denies access to an asset for which the user has no permission, and logs the event.

### **MONITOR Mode (UNIX/Linux only)**

Allows access to an asset for which the user has no permission, and sends a warning message to the Event Console Log. The user is unaware that a violation has occurred.



You can associate the Default Permission option with the System Violation Mode option to produce a specific level of asset protection. For example, a Default Permission of DENY combined with a System Violation Mode of WARN causes Security Management to log all unauthorized asset access, but not deny access. (Assets specifically permitted to a user will not generate an Event Console Log entry.) The combination of DENY and WARN is especially useful while you are in the process of implementing your security policies, because it produces an audit trail of all asset access through the Event Management component.

Alternatively, when the security option USE\_PAT is set to YES, a Default Permission option setting of ALLOW combined with a System Violation Mode of FAIL enables Security Management to protect only those assets specifically defined as protected, that is, *everything except* the protected assets will be accessible. Many security administrators prefer this approach because it quickly protects a set of defined assets without affecting the rest of the system.

### USE\_PAT Option (UNIX/Linux Only)

If you are implementing Security Management on a UNIX/Linux platform, you can effect security evaluation using the Protected Asset Table (PAT).

The USE\_PAT option affects security evaluation behavior, which is based on the Protected Asset Table, or PAT. The PAT contains asset type-asset ID combinations, which encompass the entire set of all security rules defined for your enterprise.

At evaluation time, requests are evaluated to validate user access to a specific asset type-asset ID-access mode combination, otherwise known as a tuple. For example, user Joe has access to asset type CA-MSGRECORD and access mode CONTROL. If no rule is found to match this tuple, the PAT is searched for the pair asset type-access mode within the defined rules. If a match is found, an access of DENY is returned by the evaluators for the request. If no match is found, the value of the Security Management option DEFAULT\_PERMISSION (either ALLOW or DENY) is returned.

This option can be set to YES or NO. The default is NO. Setting this option to YES maintains this level of security evaluation (search the PAT). A setting of NO bypasses PAT processing and will result in a more “get what you define” behavior. If a tuple is found for the user, the access associated with the tuple (ALLOW, LOG, or DENY) is returned; otherwise, the value of the DEFAULT\_PERMISSION option is returned.

### Authorized User List (for QUIET Mode)

The Authorized User List identifies the IDs that have authority to administer security policies. The default user IDs that are in effect for QUIET mode rely on the setting of the following Security Management option:

#### Authorized User List (Windows)

The default user IDs for the Authorized User List option are ADMINISTRATOR and CAUNINT. At a minimum, you must supply an administrator ID to continue with your implementation. If you add user IDs, separate them with a comma, as spaces are valid in user IDs.

#### SSF\_AUTH (UNIX/Linux)

The default user ID in the SSF\_AUTH option is **root**. Edit the IDs in the list using the text editor of your choice. If you add user IDs, separate them with a space.

### Remote CAISSF Return Codes

The following return codes are associated with Remote CAISSF:

#### EMSEC\_API\_REMOTE\_NOT\_SPECIFIED

The EmSecRemote\_ssfCheck API does not identify the remote node, or no remote node value was specified for the option, Remote Evaluation Server Path.

#### EMSEC\_API\_REMOTE\_NOT\_RESPONDING

The request has failed to reach the servers in the path.

### Rule Server Support (Windows only)

Rule server support lets you specify another server from which to obtain rules used for evaluation by commit processing. This server's rules are obtained in conjunction with the local server's rules. Rule server support lets you define global rules shared by two or more servers, in addition to enforcing your *local* rules. Specify the name of your global rule server as the value of this option.

### User Group Server Support (Windows)

User group server support lets you specify another server from which to obtain user groups used for evaluation by commit processing. User group server support is usually implemented on a server or current Backup Domain Controller (BDC) where users typically use domain user IDs to access assets. Rules defined for those nodes refer to the user groups defined on the domain controller. Specify the name of your user group rules server as the value of this option.

## Security Management Automatic Startup (Windows)

Unless you specified otherwise during the CA NSM installation process, this option specifies that Security Management is started during the activation of the other Enterprise Management services. You can change this option so that Security Management does not start automatically.

## Additional Options for UNIX/Linux Platforms

The options discussed in this topic are associated exclusively with UNIX/Linux platforms and are used to enable the following Security Management Node support feature.

### Node Support Implementation

Within Security Management, all asset definitions are typically *global*; that is, the definitions apply on all hosts that may share the Security Management Database.

CA NSM provides node support for the UNIX/Linux platform under which a definition can be associated with a specific node in a group. Node support is useful for sites that share an MDB among multiple systems and sites that would like to use Security Management evaluators in a distributed client/server environment.

The implementation of node support requires an understanding of asset definitions and policy evaluation before you can determine how you want to set the options that control this support.

### Asset Definitions

Asset definitions can contain an associated node value. Similarly, within any definition that accepts an asset specification, there can be a node specification associated with that asset. Upon completion of the initial Security Management installation, all node values are blank, which means that all definitions are *global*.

Asset definitions can be qualified by node. Asset definitions are associated with user, user group, and asset group definitions. When you include a node value in an asset definition, Security Management applies the policy only to that specific node. In the absence of a node value, the policy is global and applies to all nodes supported by the MDB.

When you commit Security Management policies to your system, Security Management looks for rules that are applicable to your system. It retrieves all rules that have a node value equal to the system identification on which you execute the commit process; it then retrieves rules that are global (have no associated node value). The Security Management evaluators use this set of rules to enforce the Security policy.

## Set Certain Options to Absolute Values

In addition to customizing Security Management options for your environment, you can set several options to absolute values before starting the Security Management daemons (service providers) in QUIET mode (Phase 2). Setting these options to absolute values lets you test your policies during implementation without locking out your user community.

### To set options to absolute values

1. Set the following options to their respective values:

Windows Option	UNIX/Linux Option	Value	Default Value
Overall Activation Setting	N/A	ACTIVE	ACTIVE
System Violation Mode	SYSTEM_MODE	QUIET	FAIL

2. Supply an Administrator ID for the Authorized User List (SSF\_AUTH for UNIX/Linux platforms) option, the user ID of the person who is authorized to modify your Security Management policies.

**Note:** For instructions about setting options, see Customizing Your Security Management Options in the online *CA Procedures*.

## Phase 2: Start Security in QUIET Mode

Phase 2 of implementing Security Management consists of starting the Security Management daemons in QUIET mode.

After starting the Security Management daemons in QUIET mode, you can test security policies without risk of adversely affecting your user community. Although you are in QUIET mode, an administrator ID is required; ensure that you have supplied an administrator ID for the Windows option, Authorized User List (SSF\_AUTH option for UNIX/Linux platforms), before starting the daemons.

**Note:** See the instructions for starting the Security Management daemons in the online *CA Procedures*.

## Phase 3: Create Rules for Production in WARN Mode

In Phase 3, you create security rules for users, user groups, access permissions, and asset groups in WARN mode and verify user access, adjusting rules based on results, then re-execute the commit process.

With the System Violation Mode set to WARN, you can test all of the rules you have defined and refine and adjust those rules without risk of impacting your users' ability to work.

**Note:** For instructions about configuring your system to operate in WARN mode, see Setting Options for Production, WARN Mode in the online *CA Procedures*.

Complete the procedures referenced in the rest of this section to create security rules for the following:

- User groups
- Asset groups
- Asset permissions

### Defining User Groups

Using Security Management, you can group users to simplify administration. User groups logically connect permissions to many users in a single operation. For example, if you define 50 users to the user group PRJTX, and PRJTX has permission to access Event message records, then all 50 users are automatically granted access to Event message records.

In addition, because CA NSM recognizes that a user may belong to more than one user group, your organization can define security from a variety of perspectives. For example, you could define the user USER01 to the following user groups:

---

User Group	Description
PRJTX	Project X Development Team

---

User Group	Description
LAB	Laboratory Team
PLANNING	Planning Analyst
HRSAPPL	Human Resource Systems
CLSIFIED	Clearance is Classified

USER01 has access to any CA NSM management component permitted to any of these user groups. This example illustrates that permissions can be logically assigned, based on USER01's department (PRJTX), the user's area within the department (LAB), or specific title (PLANNING). You can also base access on the application on which the user works (HRSAPPL), or even the information clearance level assigned to the user (CLSIFIED).

User groups exemplify the power of the Security Management architecture. If USER01 is promoted, changes jobs, or changes departments, permissions can be adjusted automatically by changing the user groups of which the user is a member.

**Note:** For instructions about defining your users, see *Defining User Groups* in the online *CA Procedures*.

## Nested User Groups

In addition to creating multiple groups for a user, you can nest groups (have a group that contains other groups) for additional flexibility in mapping how your organization typically works.

For example, a corporate division has a certain set of assets associated with it, such as being able to read management console records. Additional departments (such as personnel, payroll, and mailroom) also have special authorities. You can use nested user groups to give the departments automatic access to the corporate assets.

By making each of these departments a group within the corporate group, each department has access to their assets *and* any assets defined to corporate. They do not, however, have access to the assets permitted specifically to the other departments.

You can apply as many levels of nesting as you want. If, in addition to the corporate and department groups, you also have a company group, you can make the departments a member of corporate, and corporate a member of the company. Permissions apply as described in the following list:

- All members of departments can access their department's assets, the corporate assets, and the company assets.
- All members at the corporate level can access the corporate assets and the company assets.
- All members at the company level can access only the company assets.

## Defining Asset Groups

CA NSM refers to any resource as an *asset*, which is defined to CA NSM as a member of one or more *asset groups*. An asset group may have a single member or many, and access to an asset group can be permitted to either a user group or an individual user. You can also define nested asset groups.

An asset group is defined by assigning a name to the new asset group and specifying the assets that are to be members of the group. Asset groups can be nested within other asset groups and can be included as members of multiple asset groups.

Assets that are members of an asset group are identified by the following:

- Asset type
- Asset name, also referred to as asset ID

The *asset type* is the category into which the asset logically falls. You can define your own asset types; you also have available over 60 predefined asset types identified by the prefix, CA-.

For a list of supported asset types, see the asset type tables in the online *CA Reference* under the Security Management cautil ASSETTYPE command.

For procedures to define your own asset types, see Defining Dynamic Asset Types in the online *CA Procedures*.

The *asset name*, or *asset ID*, is used primarily to identify specifically and individually a particular instance of an asset type.

## Nested Asset Groups

In addition to creating nested user groups, you can nest asset groups within as many levels as you require.

For example, imagine a corporate asset group that contains various files and other general-purpose assets, one of which is the telephone directory. You want to associate these general-purpose assets with the more specific assets in three other asset groups that comprise payroll, administrative, and mailroom assets. You can accomplish the appropriate sharing of assets and avoid a duplication of asset groups by nesting the three more specific asset groups within the general-purpose asset group.

### Adding an Asset Group

As mentioned earlier, asset groups are collections of one or more assets, and are a key part of administering Security Management policies.

For procedures to create a new asset group, see Adding an Asset Group in the online *CA Procedures*.

### Asset Permissions

*Asset permissions* govern which protected assets a user can access and how they can be used after being accessed.

Permissions are created through the Security Management GUI or the cautil command line interface by specifying the name of the asset or asset group to which you want to give access to a user or user group. You can, conversely, provide the name of the user or user group to be permitted access to an asset or asset group. With this level of flexibility, you can manage security in the manner that is most comfortable for you.

**Important!** If a user is denied access based on asset permissions (or the lack of an asset permission to a protected asset), a violation results. It is important to remember that the presence of a violation does not necessarily mean that a user will be stopped from accessing the asset. The System Violation mode (QUIET, WARN, FAIL) and access type (for example, log) controls whether the user will actually reach the asset.



## Access Types

The *access type* specified in a definition determines whether the user will be allowed to access an asset. When defining access permissions, you can use the following access types to meet your specific control or audit needs:

### PERMIT

Allows access. The standard control is "Allow this user to access this asset."

### LOG

Allows access and logs the event. LOG is used in those cases where you want to maintain an audit trail of access to a critical asset. For example, you may want to record all updates to the CA NSM calendar BASE. Do this by using two access types—a PERMIT for READ and a LOG for WRITE. READ authority is allowed normally, while a WRITE request generates a record of the access to the Event Console Log. The end user will not be notified that this access has been logged.

### DENY

Denies access and logs the event. DENY is useful for creating exceptions to normal permission sets.

Whenever an asset is referenced (either explicitly or generically) as the subject of a PERMIT or DENY rule, it becomes "protected." This protection means that when you permit a user to access an asset, such as CA-CALENDAR, any other users who have not been granted permission to access this file (those in FAIL mode) will be denied access when the security option USE\_PAT is set to YES. On Windows, such protection is referred to as *implicit* DENY. USE\_PAT and Implicit DENY are disabled by default.

**Important!** Security Management considers the access mode when evaluating a rule. For example, if the access mode (READ, WRITE, and so on) of a permission does not match the requested access type, the permission is not used. For more information, see [Access Modes](#).

## Date and Time Controls

You can associate the name of a CA NSM calendar with any access rule. When a calendar is associated with a rule, that rule is considered applicable only when the dates and times specified as ON in that Calendar definition are met.

For example, calendars are often used in conjunction with DENY asset types to set special dates and times when access will be denied.

Understanding how access modes and calendars affect security evaluation make it possible to construct sophisticated security rules. For example, assume the staff in the systems management department has read and write access to the CA NSM calendar BASE, and you want to deny write access to members of that group on weekends. You could create a weekdays-only calendar and associate it to Payroll as a PERMIT access type for access mode UPDATE.

## Access Modes

The access policies of Security Management support several types of authority, or *access modes*. Think of access modes as ways a user may try to access an asset. An access policy can specify one, several, or all of the applicable access modes for an asset. At least one access mode must be specified.

These access modes include the following:

### **READ**

Controls READ access to the asset.

### **WRITE**

Controls WRITE access to the asset.

### **UPDATE**

Controls UPDATE access to an asset. UPDATE is only valid for CA NSM asset types. UPDATE access implies READ and WRITE authority.

### **CREATE**

Controls creating a new asset, such as a new calendar.

### **DELETE**

Controls deleting or removing the asset. For example, this mode would prevent a calendar erase or delete operation. DELETE is useful for preventing accidental deletions.

### **CONTROL**

Controls access CA NSM administration objects, such as STATION, CALENDAR, and so forth.

## Examples: Access Rules

The following examples present typical questions and answers regarding access rules. These examples assume the system is running in DENY mode. When your system is in DENY mode, the set of permissions you initially define may not grant your user community access to all the tools they require.

**Note:** On Linux and UNIX, system administrators (SYSADMIN group) have full access to most Unicenter NSM components. By default, the only users in the group are "administrator" and "root." If you want to grant another user access, add the user ID to the SYSADMIN group.

To grant a user access to the Event Console Log, define a PERMIT rule with the following information:

User ID:	USER01
Asset Type:	CA-CONLOG
Asset Name:	*
Access Modes:	CONTROL
Access Type:	PERMIT

## Defining Access Permissions

When Security Management makes a real-time decision about whether to permit a user or user group access to an asset, it must identify two entities--the user and the asset. Once identified, CA NSM must determine what policies (if any) have been defined that describe how this access attempt should be handled. Defining these policies typically uses one of two approaches--"user perspective" or "asset perspective."

Defining an access permission from the user perspective can be thought of in terms of "what asset does this user need to access, and where?" Defining an access permission from the asset perspective can be thought of in terms of "who requires access to this asset, and where?" CA NSM is uniquely capable of supporting both perspectives in granting access permissions to assets.

A common requirement of both perspectives is that the asset access rule have a specific access mode and a specific access type. For lists of access modes and access types, see *Access Modes*.

## Access Determination

When a user attempts to access an asset, Security Management looks at all rules associated with the user to find the ones that apply to the access in question. Security Management only considers rules that match the asset type, asset name, asset node, access mode, and any conditions determined by a calendar or criteria profile. If several permission rules match an asset name, the rule with the best fit (closest match to asset name) is used, and an access type DENY overrides LOG, which overrides PERMIT.

For step-by-step procedures to define your access permissions, see the following topics in the online *CA Procedures*:

- Permitting Users to Assets
- Permitting Assets to Users

## Rule Evaluation

Two decisions are made during the security evaluation process. The first decision is whether a specific access attempt is considered authorized. If the access is not authorized, a violation results and the second decision—what to do about the unauthorized access attempt—depends on the Enforcement mode in effect. The only Enforcement mode that results in access being denied is FAIL mode, which is set by the System Violation Mode.

For additional information about Enforcement modes, see Access Violations and Enforcements.

## How CAISSF Scoping Options Work

Using the CAISSF Scoping feature, you can assign the rights to administer CA NSM securable objects in a granular fashion. Because CA asset types identify specific parts of CA NSM, you can secure those parts by defining an access permission using CAISSF Scoping. Some examples of permissions using a scoping rule are as follows:

- User can administer user profiles for this node, but not for that node.
- User can use DEFINE, but not ALTER.
- User can modify other parameters in a user profile, but not the password (see Examples).

Scoping rules can be written only for CA asset types (those having the prefix CA-), but not for user-defined asset types.

Scoping can narrow an access permission to a keyword object, a data object, or a command object. Each of these has a specific CA asset type and an associated Security Management option (which must be set before scoping can be applied).

## CAISSF Scoping Options

CAISSF Scoping options let you assign access permissions to administer CA NSM securable objects. Each object, its related option, and its CA asset type (identified by the suffix) are shown in the following table:

<b>Object to Be Secured</b>	<b>Security Management Option</b>	<b>CA Asset Type Suffix</b>
Keyword	Windows: SSF Keyword Scope UNIX/Linux platforms: SSF_SCOPE_KEYWORD	KW

<b>Object to Be Secured</b>	<b>Security Management Option</b>	<b>CA Asset Type Suffix</b>
Data	Windows: SSF Data Scope UNIX/Linux platforms: SSF_SCOPE_DATA	DT
Command	Windows: SSF Command Scope UNIX/Linux platforms: no option	CM

### CAISSF Keyword Scope

Scoping of keyword asset types provides validation against Security Management policies based on the specified keywords. The use of keyword scoping provides finer granularity of access to CA asset types. For example, you can deny access to specify the UID of a user, but allow access to specify the user name.

Keyword scoping rules are specified for asset types having a suffix of KW. For instance, use CA-CALENDAR-KW to apply a keyword scoping rule to the CA-CALENDAR asset type.

### CAISSF Data Scope

Scoping of data asset types provides validation against Security Management policies based on the data specified for a keyword value. Using data scoping, access to CA asset types can be limited to the level of the value within a keyword. For example, access to update the CA-CALENDAR object can be restricted for ID=base, but allowed for other calendars.

Data scoping rules are specified for asset types having a suffix of DT. For instance, use CA-CALENDAR-DT to apply a data scoping rule to the CA-CALENDAR asset type.

When specifying the asset ID (type) for a CA data object (suffix DT), you must supply a setup character immediately preceding the operand (as the underscore is used with the node= operand in the previous example). This character is used by the rule evaluator to edit the definition, and is required to indicate to the evaluator that the next specification is a new operand. You can use one of the following characters:

- Tilde (~)
- Slash (/)
- Pipe (|)
- Underscore (\_)

**Note:** Scoping on data objects is not supported through the EmSec APIs.

### CAISSF Command Scope (Windows)

Scoping of command asset types provides validation against Security Management policies based on the specified command. Using command scoping, access to CA asset types can be limited to specific operations. For example, a user can create new user accounts, but cannot update existing accounts.

Command scoping rules are specified for asset types having a suffix of CM. For example, use CA-CALENDAR-CM to apply a command scoping rule to the CA-CALENDAR asset type. Command scoping is valid only on Windows platforms.

### How You Define Scoping Rules

You can define a scoping rule using either the Security Management GUI or the cautil command line interface when defining access permissions. For examples of using the GUI and the cautil command line, see Restricting Access Permission (Scoping) in the online *CA Procedures*.

### How You Commit Rules in WARN Mode

When you finish defining or modifying any security rules for users, user groups, access permissions, and asset groups, you must execute a commit process to put the new rules into effect. For more information, see Understanding the Commit Process. For procedures to perform the commit process, see Committing Your New Rules in WARN Mode in the online *CA Procedures*.

## Phase 4: Set Options for Production, FAIL Mode

After you have defined your security policies and carefully reviewed them, proceed with this phase of implementation to place your security policies into an active enforcement, or FAIL mode. With System Violation mode set to FAIL, and after executing a commit process, violations will be logged to the Event Console Log, access will be denied, and violation actions will be enforced.

For procedures to set options, see Setting Options for Production, FAIL Mode in the online *CA Procedures*.

### How You Commit Rules in Fail Mode

Now that you have established your security policies for FAIL mode, you must execute the commit process to put those policies into effect.

For procedures about the commit process, see Performing a Commit Using the GUI in the online *CA Procedures*. For more information about using secadmin, see the secadmin command in the online *CA Reference*.

### How You Deactivate Security Management

Under extremely serious circumstances, such as a hardware failure causing a loss of the Security Management Database and DSBs, you may find it necessary to deactivate Security Management.

**Important!** After Security Management has been deactivated, the only security in effect will be that inherent to the native operating system.

For procedures to deactivate Security Management, see Deactivating Security Management in the online *CA Procedures*.

## Security Management Reports

This section provides information about how to interpret the reports that extract security audit data from the Event Console Log.

Also included in this topic are instructions about how to use the UNIX/Linux platform whohas and whathas commands and interpret the reports they generate.

## Access Violations Written to the Event Console Log

When a user attempts to access an asset to which he does not have permission, an error message is written to the Event Console Log. The Event Console Log may include violation errors that look similar to the following message:

```
CASF_E_465 Access violation by userid to asset ( mode ) assetname from source terminal_device at node source_node for access_type access mode. (context )
```

### **CASF\_E\_465**

Specifies the general message number used for all DENY violations.

#### ***userid***

Specifies the ID of the user who caused the violation.

#### ***mode***

Specifies the user's violation mode: W=Warn, M=Monitor, F=Fail.

#### ***assetname***

Specifies the asset name of the asset involved in the violation. For WNT-FILE, UNIX-FILE, and UNIX-SETID, the asset name is a fully qualified path name.

#### ***terminal\_device***

Specifies the device the user was logged into at the time of the violation.

#### ***source\_node***

Specifies the node from which the user logged into the system.

#### ***access\_type***

Specifies the access mode, abbreviated as follows: Rd=read, Wr=write, Up=update, Sc=scratch, Fe=fetch, Cr=create, Co=control, Se=search, Ex=execute.

#### ***context***

Specifies the context of the violation. For Windows intercepted events, specifies the access type (read, write, and so on). For UNIX/Linux platforms, specifies the system call name. For CAISSF resources checks through components, the context specifies "resource."

## UNIX/Linux Reports

You can generate two reports on UNIX/Linux platforms that let you review Security Management policies.



## Whohas Report

You can use the whohas report to look at the policies that have been set for a particular asset type and asset name. To create this report, run the following command from the command line prompt:

```
whohas [asset_type] [asset_value] {user_name} {node_name}
```

The following command was used to generate the following sample report:

```
whohas CA_CALENDAR_BASE
```

```
COMPUTER ASSOCIATES INTERNATIONAL, INC. 'WHOHAS' FOR UNICENTER
USER: audit                                     PAGE 1
NODE:
ASSETNODE:
  ---- ACCESS MODES --- ----- RULE -----
NUM FILE SSF  NAME  - ORIGIN PERMISS        TEXT
-----
  rwdxc --- allfcrit    User  PERMIT
```

```
COMPUTER ASSOCIATES INTERNATIONAL, INC. 'WHOHAS' FOR UNICENTER
USER: causer1                                   PAGE 2
NODE:
ASSETNODE:
  ---- ACCESS MODES --- ----- RULE -----
NUM FILE SSF  NAME  - ORIGIN PERMISS        TEXT
-----
  r--- --- rcrit      User  PERMIT
```

Total rules: 2

Whohas run by root on Tue Jul 10, 2001 at 11:49:46

End of whohas

The USER, NODE, and ASSETNODE values identify the user, the node associated with the user, and the targeted asset node, respectively. The whohas report groups the assets by the USER, USERNODE, and ASSETNODE values. The ACCESS MODES indicate the CAISSF access modes. The access modes are abbreviated as follows: R=READ, W=WRITE, D=DELETE, X=EXECUTE, U=UPDATE, C=CREATE, S=SEARCH, N=CONTROL.

In addition to the FILE and CAISSF access type flags, the NAME field lists the internal criteria name (for diagnostic purposes) or the name of a custom jll criteria profile.

The RULE is a combination of the following:

- ORIGIN—User or user group source of this policy
- PERMISS—Permission granted by this policy
- TEXT—Asset name for this policy

## What-Has Report

You can use the what-has report to look at the policies that have been set for a particular user ID. To create this report, run the following command from the command line prompt:

```
whathas [userid] [node]
```

The following command was used to generate the following sample report:

```
whathas audit
```

```
07/10/01                WHAT-HAS REPORT                Userid: audit

Asset Type/           Access Criteria
Asset Name/Asset Node  Type    Profile  RDXUCSC  Expires  Calendar Path  Path
-----
ALENDAR                PERMIT  allfcrit  YYYYNNNN
DENY    rwxcrit  YYNNNNN
```

# Appendix A: Unicenter NSM r11.2 UNIX and Linux Support

---

This appendix describes Unicenter NSM r11.2 support for UNIX and Linux platforms, lists the components supported by UNIX and Linux, and points to the areas in this guide that apply to UNIX and Linux users.

This section contains the following topics:

[UNIX and Linux Support](#) (see page 403)

[Supported Components](#) (see page 404)

[UNIX and Linux Support Quick Reference](#) (see page 406)

## UNIX and Linux Support

Unicenter Network and Systems Management r11.2 provides support on UNIX and Linux platforms for key CA NSM manager components. This release provides an upgrade path for UNIX and Linux users with Unicenter NSM 3.1 and r11 managers installed.

Unicenter NSM r11.2 provides full manager support on UNIX and Linux for Event Management and Agent Technology and provides database abstraction for the Event Manager.

**Note:** For more information about the database abstraction, see the *MDB Overview*.

All components supported on UNIX and Linux contain all changes and updates applied to these components for the base CA NSM r11.2 product. Therefore, IPv6 addresses are supported on UNIX and Linux platforms.

Unicenter NSM r11.2 does not support UNIX and Linux for all components of the base Unicenter NSM product. For a listing of the components supported and references to the applicable areas in this guide, see the topics that follow.

## Supported Components

Unicenter NSM r11.2 supports only a subset of the components on UNIX and Linux that are included in the base Unicenter NSM product. The components supported provide upgrades for UNIX and Linux users for Event Management, Agent Technology, and other vital areas of the product.

**Note:** CA NSM r11.2 for UNIX and Linux does not support Ingres. Any UNIX and Linux information in the CA NSM documentation set pertaining to Ingres databases does not apply to CA NSM r11.2 users.

The following are the components supported on UNIX and Linux in Unicenter NSM r11.2:

### **Event Management**

Unicenter NSM r11.2 includes an upgrade of the Event Manager component on UNIX and Linux. The Event Manager supports the Calendar and uses the new free embedded database PostGreSQL. For more information in this guide about Event Management, see the section Event Management in the chapter "Administering Critical Events."

### **Distributed State Machine**

Unicenter NSM r11.2 includes the remote DSM manager component on UNIX and Linux. This component is the manager supporting Agent Technology, and it lets you remotely communicate with local DSM and WorldView Manager components in CA NSM r11.2 environments. For more information, see the section Understanding Systems Management in the chapter "Monitoring Your Enterprise."

### **Alert Management enablement**

Provides the ability to create alerts for the Alert Management subsystem. The full AMS manager is *not* supported on UNIX and Linux. Alert enablement lets you forward alerts to a remote Alert Management Server. For more information about creating alerts for Alert Management, see the section Alert Management System in the chapter "Administering Critical Events."

### **Advanced Event Correlation Agent**

Provides the ability to implement Advanced Event Correlation policy. The Advanced Event Correlation user interface is not supported. For more information about the Advanced Event Correlation Agent, see the chapter "Correlating Important Events."

### **Event Trap Processing**

Provides processing of external traps. This capability requires that you add the CCS Trap Daemon and CCS Trap Multiplexer.

**CCI**

Enables communication of certain CA NSM components with other components. For more information, see the section Common Communications Interface in the chapter "Securing CA NSM."

**DIA**

Provides cross-component communication for certain CA NSM components.

**SDK**

CA NSM r11.2 includes UNIX and Linux support for the Event Management and Agent Technologies SDKs. Utilities associated with Event and Agent managers required for administration and maintenance are also supported.

**Management Command Center**

Provides a centralized web interface for viewing and managing data collected about your monitored resources. The following MCC dependencies are also supported:

- AIS
- CAM (CA Messaging)
- EM Provider

For more information about the MCC, see the chapter "Administering Critical Events" in this guide or the *MCC Help*. For more information about CA Messaging, see the appendix "Using Ports to Transfer Data."

**High Availability Service**

Enables high-available ready components in clustered configurations. High-availability readiness is limited to Linux platforms only. Components that are high-availability ready in the base CA NSM r11.2 product are also high-availability ready on Linux, except for Event Management and the Calendar.

**Job Management Option**

Includes the manager component for Job Management Option (JMO). For more information about JMO, see the appendix "Job Management Option."

**CA NSM Security**

Provides the built-in security solution for Unicenter NSM. CA NSM security can secure your environment at the object or access level. For more information about CA NSM security, see the chapters "Securing CA NSM" and "Securing CA NSM Objects."

**Trap Manager**

Lets you manage trap databases and trap filter files. For more information about Trap Manager, see the appendix "Managing Traps Using the Trap Manager."

## UNIX and Linux Support Quick Reference

The following table lists the components supported on UNIX and Linux platforms in Unicenter NSM r11.2, where to find the component information in this guide, and any other documents that contain useful information for these components. For details about to what extent Unicenter NSM r11.2 supports these components on UNIX and Linux platforms, see Supported Components.

**Note:** CA NSM r11.2 for UNIX and Linux does not support Ingres. Any UNIX and Linux information in the CA NSM documentation set pertaining to Ingres databases does not apply to CA NSM r11.2 users.

Component	Administration Guide	Other Guides
Event Management	See the section Event Management in the chapter "Administering Critical Events"	Inside Event Management, MCC Help
Calendar	See the section Event Management in the chapter "Administering Critical Events"	Inside Event Management, MCC Help
CCI	See the section Common Communications Interface in the chapter "Securing CA NSM"	CA Reference
DIA		Implementation Guide
SDK		Programming Guide
MCC	See the chapter "Administering Critical Events"	MCC Help
High Availability Service		Inside Systems Monitoring, Implementation Guide
Database abstraction		MDB Overview
DSM	See the section Understanding Systems Management in the chapter "Monitoring Your Enterprise"	Inside Systems Monitoring
Job Management Option	See the appendix "Unicenter Job Management Option"	MCC Help

---

<b>Component</b>	<b>Administration Guide</b>	<b>Other Guides</b>
NSM Security	See the chapters "Securing CA NSM" and "Securing CA NSM Objects"	CA Procedures
Trap Manager	See the appendix "Managing Traps Using the Trap Manager"	Trap Manager Help
Alert Management enablement	See the section Alert Management System in the chapter "Administering Critical Events"	MCC Help, Inside Event Management
Advanced Event Correlation Agent	See the chapter "Correlating Important Events"	Inside Event Management

---

For information about installation and migration, see the *Implementation Guide* and *Migration Guide*. For information about the database abstraction, see the *MDB Overview*.





# Appendix B: FIPS-140-2 Encryption

---

This appendix outlines what components of CA NSM r11.2 are FIPS-140-2 compliant, what data is protected in a compliant manner, details about the eTrust PKI, FIPS configuration and migration information for each component, and other FIPS-related information.

This section contains the following topics:

[CA NSM FIPS 140-2 Compliance](#) (see page 409)

[Compliant Components](#) (see page 409)

## CA NSM FIPS 140-2 Compliance

CA NSM r11.2 does not provide full FIPS 140-2 compliance. However, several areas of the product support this level of encryption and use FIPS-140-2 compliant cryptographic libraries to protect passwords and other sensitive data.

Components in CA NSM r11.2 that support FIPS 140-2 give you the option to encrypt all sensitive data using FIPS-compliant algorithms provided by the eTrust PKI (ETPKI). The ETPKI is a SDK included with CA NSM that enables certain components to support FIPS encryption through use of its cryptographic libraries and encryption algorithms.

## Compliant Components

The following CA NSM components provide FIPS 140-2 compliant encryption:

- Systems Performance
- Active Directory Management
- Agent Technology
- Common Communications Interface (CCI)

The following component provides FIPS 140-2 encryption support in certain situations:

- Management Command Center
- Unicenter Management Portal
- Web Reporting Server

## Systems Performance

Systems Performance provides FIPS 140-2 encryption support using the ETPKI for all sensitive data. The ETPKI wraps the FIPS 140-2 validated RSA BSAFE Crypto-C Micro Edition cryptographic module. Systems Performance uses the AES encryption algorithm with a 256-bit strength key to encrypt keys and data. It also uses SHA-1 (Secure Hash Algorithm) to hash the keys to make sure they are not tampered with.

**Note:** When FIPS mode is not enabled, Systems Performance uses the PKCS #5 v2.0 algorithm to generate a password generated key.

By default, FIPS encryption is disabled. You must enable the encryption, at which time any sensitive data is re-encrypted. Data is encrypted to a Systems Performance-specific library using a Data Encryption Key (DEK), which must be distributed to all Systems Performance servers where encryption and decryption is required.

## Data Encrypted

The following Systems Performance data is encrypted when using FIPS mode:

### **SNMPv3 credentials**

Includes the SNMPv3 security name, authentication password and protocol, and privacy password and protocol used to access SNMPv3 devices and their MIBs. This information is created by Performance Configuration and decrypted by Performance Scope for real-time monitoring and the Performance Agent for historical monitoring.

**Files:** Stored by the Performance Domain Server in files named resource.ext and <machine>.cmp. Copied by the Performance Distribution Server to Performance Agents configuration file Hpaagent.cfg if it is configured to monitor an SNMP device.

### **SNMPv1/v2 credentials**

Includes SNMPv1/v2 community string information. This information is encrypted by Performance Configuration when accessing a MIB or device, and it is decrypted by Performance Scope for real-time monitoring and the Performance Agent for historical monitoring.

**Files:** Stored by the Performance Domain Server in files named resource.ext and <machine>.cmp. Copied by the Performance Distribution Server to Performance Agents configuration file Hpaagent.cfg if it is configured to monitor an SNMP device.

**Batch Reporting credentials**

Includes computer access credentials used by Performance Trend Batch Reporting to successfully generate and output reports. Batch Reporting must supply credentials to Unicenter Workload that let it interact with the desktop. These credentials are encrypted by Performance Trend Batch Reporting after being entered into the Performance Trend Batch Reporting Wizard, and they are decrypted by the Performance Trend Batch Reporting Generator when executing the Batch Reporting Profile.

**Location:** Stored by Performance Trend in the Performance Trend Batch Reporting Profile files (\*.tbp) held within the Performance Trend area of the <Logged in User> or <All Users> application data area of the computer.

**MDB credentials**

Includes MDB connection credentials used by the Performance Domain Server to publish summary performance data to the MDB. These credentials are created through the Systems Performance Installer or the Performance Domain Server configuration utility (cfgutil), and they are decrypted by the Performance Domain Server when accessing the MDB to publish summary data.

**File:** dbcred.dat in the Performance Domain Server

**Installer response file credentials**

Includes MDB connection credentials that you enter before a response file-generated installation so that the Performance Domain Server can publish Performance Data to the MDB, and Performance Reporting can use the WorldView section of the MDB for reporting. The response file is created by the Systems Performance Installer, and the MDB credentials are decrypted by the Systems Performance Installer when running an installation in response file mode.

**Location:** Stored in a response file for later use in response file-driven installations. The response file is stored in a user specified file.

**Performance Data Grid access information**

Includes the user names and domain names used to gain access to the Performance Domain Server and Performance Distribution Server data and operations. The user names and domains are created by the user on the Performance Domain Server, and they are decrypted by the Performance Domain Server and Performance Distribution Server hosts.

**File:** Maintained by the Performance Domain Server in a file named users.dat and distributed to the Performance Distribution Server so that it can also validate data requests.

### Unicenter Management Portal and Web Reporting Service credentials

Includes Unicenter Management Portal and Web Reporting Service connection details that are required for Performance Trend to publish reports to either of these tools. These credentials are encrypted by Performance Trend Batch Reporting after being entered into the Performance Trend Batch Reporting Wizard, and they are decrypted by the Performance Trend Batch Reporting Generator when publishing reports to Unicenter Management Portal or Web Reporting Service.

**Location:** Maintained by Performance Trend in Portal Profiles (\*.pop) that are held within the Performance Trend area of the <Logged in User> or <All Users> application data area of the computer.

### Data Encryption Key

In Systems Performance, sensitive data is encrypted in FIPS mode by ETPKI using a 256-bit Data Encryption Key (DEK). By default, Systems Performance uses an embedded key to perform FIPS-based encryption, but you can create a custom non-embedded key.

Custom keys are protected by an additional Key Encryption Key (KEK) embedded within Systems Performance. We recommend also securing custom keys using operating system file security. The keyfile is in the Systems Performance installation directory at the following location:

```
%CASP_PATH%\appdata\keystore
```

We recommend giving only administrators permissions to access the directory and keyfiles for Performance Manager and Performance Agent servers and read access to non-admin users on servers running UI components.

### Turn on FIPS Mode

By default, FIPS encryption is not enabled in Systems Performance. You must turn on FIPS mode manually using the CASP\_FIPS\_MODE environment variable.

To turn on FIPS mode, locate the CASP\_FIPS\_MODE environment variable and set its value to '1' on all Domain Server and Performance Trend servers.

This enables FIPS 140-2 encryption in Systems Performance. Any sensitive data is re-encrypted using the new FIPS compliant algorithm.

If necessary, you can switch back to non-FIPS mode at any time by setting the CASP\_FIPS\_MODE variable value to anything but '1' or removing the variable.

**Note:** For more information about the complete process you must perform to ensure that the switch to FIPS mode is complete and without errors, see How to Switch to FIPS Mode.

## Installation Considerations

You can perform an initial manager installation of Systems Performance in any of the following ways:

- FIPS mode turned off
- FIPS mode on with the default key
- FIPS mode on with a custom key

We recommend that you perform an initial installation with FIPS mode turned off (the default setting), in which the installer continues to encrypt all data using password-based encryption. If earlier versions of Performance Agents exist in your enterprise, the manager may not be able to configure these agents with FIPS mode enabled. Once all Performance Agents are upgraded to r11.2 levels, you can enable FIPS encryption.

**Note:** FIPS mode cannot be enabled if earlier releases of the Performance Agents are to be installed on platforms not currently supported by CA NSM r11.2. These earlier agents will be unable to decrypt the encrypted configuration information.

No special steps are required when reinstalling manager components, UI components, or Performance Agents with FIPS mode turned on or off.

For more detailed information about how the FIPS encryption option affects your installation and several deployment scenarios, see the Systems Performance documentation. For a description of the recommended installation scenario, see [How to Install Systems Performance with FIPS Mode Off](#).

## How to Install Systems Performance with FIPS Mode Off

When installing or upgrading to CA NSM r11.2, we recommend performing the Systems Performance Manager installation with FIPS mode turned off, so that the manager can configure any earlier versions of Performance Agents. Complete the following process to install Systems Performance with FIPS mode off and enable FIPS encryption after installation.

**Note:** See the Systems Performance documentation for more information, including detailed installation instructions and other deployment scenarios, such as performing a clean installation with FIPS mode enabled and performing a response file-driven installation using FIPS encryption.

1. (Optional) Create a custom key if necessary by running setup.exe with the following parameter:

```
setup.exe /genkey <File path>
```

**<File path>**

Specifies the full path of an existing directory, including the key file name, where you want to store the key file.

**Note:** By default, Systems Performance uses an embedded key to perform the FIPS-based encryption and decryption and does not require the creation of a custom key.

2. (Optional) Copy the generated custom key to the CA NSM installation image by copying the DVD image to a writeable drive and copying the generated file key to the following location:

**Windows**

```
Windows\NT\SystemPerformance\data
```

**UNIX/Linux**

```
/data/sysperf_key
```

When you place the key file on the image, the Systems Performance installer automatically copies it to the installed system.

3. Install the Manager components.
4. Install the UI components.
5. Install the Performance Agents.

**Note:** You can deploy additional Performance Agents after the initial setup process. For more information, see the Systems Performance documentation.

6. Stop all client applications and the Performance Domain Server, and [turn on FIPS mode](#) (see page 412) on the Domain Server.
7. Restart the Domain Server and any client applications.
8. Reencrypt all existing Domain Server-based data using the [CASPEncrypt utility](#) (see page 421).
9. Stop Performance Trend, and [turn on FIPS mode](#) (see page 412) on all Performance Trend servers.
10. Reencrypt existing Batch Reporting profiles on all Performance Trend servers using the [CASPEncrypt utility](#) (see page 421).
11. Redeliver all profiles to Performance Agent servers.

Performance Agents continue to run with the configurations encrypted using non-FIPS based encryption until you redeliver all profiles.

## How to Switch to FIPS Mode

If you want to switch Systems Performance to run in FIPS mode, you must ensure that all appropriate servers in your enterprise are switched and all existing data is reencrypted. Complete the following process to ensure that your switch to FIPS mode is complete and without error:

**Note:** This process applies only when you are using the default embedded encryption key. If you want to create a new key, you must do so before starting the process. For more information, see [How to Change the FIPS Encryption Key](#).

1. Close all Performance Client applications and stop the Performance Domain Server through the Windows Service Control Manager.
2. [Turn on FIPS mode](#) (see page 412) on the Domain Server.

At this point all old data is still readable but not FIPS encrypted, and new data is encrypted using FIPS.

3. Restart the Performance Domain Server and any client applications.
4. Reencrypt all existing Domain Server-based data using the [CASPEncrypt utility](#) (see page 421).

5. [Turn on FIPS mode](#) (see page 412) on all Performance Trend servers.

At this point all old data is still readable but not FIPS encrypted, and new data is encrypted using FIPS.

6. Reencrypt existing Batch Reporting profiles on all Performance Trend servers using the [CASPEncrypt utility](#) (see page 421).
7. Redeliver all profiles to Performance Agent servers.

Performance Agents continue to run with the configurations encrypted using non-FIPS based encryption until you redeliver all profiles.

## How to Change the FIPS Encryption Key

If you want to change the FIPS encryption key to a newly generated one, you must distribute the key to all appropriate servers and reencrypt all existing data using the new key. Complete the following process to ensure that your migration to the new key is complete and without error.

**Note:** If you want to switch to FIPS mode in addition to generating a new custom key, complete the first two steps of this process, and then follow the steps in [How to Switch to FIPS Mode](#). If switching to FIPS mode, you do not have to copy the previous key as in Step 2.

1. Generate a FIPS encryption key using the [CASPKKeyUtil utility](#) (see page 420) on the Domain Server and install it on the server using the `-i` parameter.
2. Distribute the FIPS key to all servers running Systems Performance, first copying the previous key to `key.old`.

**Note:** To ensure that the key is copied to the correct location on the target servers, use the `-i` parameter of the CASPKKeyUtil utility and specify the correct file to install the key to. The CASPKKeyUtil utility also automatically copies the existing key to `key.old`.

3. Close all client applications.
4. Reencrypt all Domain Server-based data using the [CASPEncrypt utility](#) (see page 421).
5. Reencrypt Batch Reporting profiles on all Performance Trend servers using the [CASPEncrypt utility](#) (see page 421).
6. Restart client applications.
7. (Optional) Delete the old key from the Performance Domain Server servers. The `-p` parameter of the CASPKKeyUtil utility removes the old key.
8. (Optional) Delete the old key from the Performance Distribution Server servers.
9. (Optional) Delete the old key from the Performance Client servers.

**Note:** Perform Steps 7-9 only if changing from one non-embedded key to another non-embedded key, and if you feel that the old key is compromised.

10. Redeliver all profiles to agent servers.

Performance Agents continue to run with the configurations encrypted with the existing key until you redeliver all profiles. Once the agent receives a configuration encrypted using the new key, it automatically deletes the previous key.



## How to Switch Off FIPS Mode

To switch from FIPS mode back to non-FIPS mode, you must do so on all appropriate servers and reencrypt all existing data. Complete the following process to ensure that your switch back to non-FIPS mode is complete and without errors.

**Note:** Non-FIPS mode is the default encryption mode, so you do not have to switch off FIPS mode unless you have previously turned it on.

1. Close all client applications, and ensure that the Performance Client applications and the Performance Domain Server are stopped.
2. Turn off FIPS mode on the Domain Server server by removing the CASP\_FIPS\_MODE variable.

At this point all new or updated data is encrypted using non-FIPS mode encryption. However, existing data remains encrypted in FIPS mode and is still accessible by all applications.

3. Restart the Performance Domain Server, and reencrypt all Domain Server based data using the [CASPEncrypt utility](#) (see page 421).
4. Turn off FIPS mode on all Performance Trend servers by stopping Performance Trend and removing the CASP\_FIPS\_MODE variable.

At this point all new or updated Batch Reporting profiles are encrypted using non-FIPS mode encryption. However, existing profiles remain encrypted in FIPS mode and are still readable by Performance Trend.

5. Reencrypt Batch Reporting profiles on all Performance Trend servers using the [CASPEncrypt utility](#) (see page 421).
6. Redeliver all profiles to Performance Agent servers.  
Performance Agents continue to run with the configurations encrypted in FIPS mode until you redeliver all profiles.
7. (Optional) Delete the old keys from the Performance Domain Server servers.
8. (Optional) Delete the old keys from the Performance Client servers (including Performance Trend servers).
9. (Optional) Delete the old keys from the Performance Agent servers.

**Note:** Perform Steps 7-9 only if you used a non-embedded key for FIPS encryption and you feel that the key is compromised. Before deleting keys, make sure that all FIPS encrypted data has been reencrypted, because FIPS encrypted data cannot be decrypted without the encryption keys. Use the -purgeall parameter in the CASPKeyUtil utility to delete all keys.

## How to Migrate from a Previous Release

When you upgrade Systems Performance from a previous release (3.1, r11, or r11.1), we recommend that you perform the upgrade in non-FIPS mode so that all Performance Agents are able to read their configurations. Following is the basic upgrade process:

1. Upgrade the manager components to r11.2.
2. Upgrade the UI components to r11.2.
3. Upgrade Performance Agents to r11.2.
4. Enable FIPS-based encryption.

The process is similar to installing Systems Performance with FIPS mode off. For details, see [How to Install Systems Performance with FIPS Mode Off](#). For more information about how to upgrade, see the [Systems Performance documentation](#).

## How to Add the Performance Agent to an Existing FIPS Enabled Enterprise

When deploying the Performance Agent to an enterprise where FIPS is enabled, we recommend that you include the current FIPS key in the master image. When the key is copied to the master image, it is automatically deployed to the agent server, allowing it to decrypt any SNMP credentials provided in its initial configuration.

Complete the following process to copy the current FIPS key to the master image and install the Performance Agent:

1. Log onto the Performance Domain Server.
2. Extract the key file by running the following command:

```
CASPKeyUtil -e <FILE>
```

**<FILE>**

Specifies the name of the file to export the active key to.

3. Copy the DVD image to a writeable drive.
4. Copy the file you exported the key to to the following location:

**Windows**

```
Windows\NT\SystemPerformance\data
```

**UNIX or Linux**

```
/data/sysperf_key or /<Platform>/data/sysperf_key
```

5. Install the Performance Agent.

The current key is automatically deployed to the Performance Agent server.

An alternative method is to copy the key to the target system after the Performance Agent has been installed. Complete the following process to copy the appropriate encryption key to the Performance Agent server after installation:

1. Ensure that the Performance Agent is installed, and install it if necessary.
2. Log onto the Performance Domain Server.
3. Extract the key file by running the following command:

```
CASPKKeyUtil -e <FILE>
```

**<FILE>**

Specifies the name of the file to export the active key to.

4. Log onto the Performance Agent server.
5. Install the key file by running the following command:

```
CASPKKeyUtil -i <FILE>
```

**<FILE>**

Specifies the name of the key file exported from the Performance Domain Server.

### How to Update the User Domain Access File in a FIPS Environment

You may need update the domain access file (users.dat) periodically to add or remove user and domain names used to access the Performance Domain Server and Performance Distribution Server data and operations. This file cannot be encrypted when you edit it. Complete the following process to un-encrypt, edit, and reencrypt the users.dat file.

1. Stop the Performance Domain Server.
2. Un-encrypt the users.dat file by running the following command:

```
CASPEncrypt -x
```

**Note:** You can only run this command on the Performance Domain Server.

3. Make the required updates to the users.dat file. Find the file at the following location:

```
%CASP_PATH%\DomainServer\data\users.dat
```

**Note:** For more information about how to update the users.dat file, see the *Inside Systems Performance* guide.

4. Re-encrypt the file after completing the updates by running the following command:

```
CASPEncrypt -y
```

5. Restart the Performance Domain Server.

## CASPKeyUtil Utility--Generate a New Key

The CASPKeyUtil command line utility lets you generate a new data encryption key if you do not want to use the provided key. With this utility, you can also view information about existing keys, delete existing keys, and specify the location of existing keys and the new key.

This utility has the following format:

```
CASPKeyUtil [  
-i, --install [FILE]  
-c, --create <FILE>  
-e, --export <FILE>  
-p, --purge  
-p, --purgeall  
-f, --info [FILE]  
-?, --help]
```

### **-i, --install [FILE]**

Generates and installs a new key into the key store. If you specify a file, this key is installed into the key store.

### **-c, --create <FILE>**

Creates a new key into the specified file, but does not install it into the key store.

### **-e, --export <FILE>**

Exports the active key to the specified file name.

### **-p, --purge**

Purges the old key in the key store.

### **-p, --purgeall**

Purges all of the keys in the key store.

### **-f, --info [FILE]**

Displays the properties for either the active or specified key.

### **-?, --help**

Displays the utility's help.

**Important!** The CASPKeyUtil utility maintains only two keys at a time - key and key.old. Therefore, you must ensure that all sensitive data has been reencrypted using the CASPEncrypt utility before you create a second custom key. When you create a second custom key, the original key is permanently deleted, and you will not be able to recover the data encrypted by that key if you have not reencrypted it.

## CASPEncrypt.exe Utility--Reencrypt Data

### Valid on Windows

The CASPEncrypt.exe utility lets you reencrypt data that has been previously encrypted with a different mode or different key. When you change the encryption mode or the encryption key, the data encrypted using the previous key or mode becomes invalid and requires reencryption. When you run this utility, Domain Server data, Performance Trend profiles, or both are reencrypted using the new key or mode, with the old key or mode being used to decrypt the existing data.

This utility has the following format:

```
CASPEncrypt [  
-d, --domain  
-t, --trend  
-a, --all  
-h, --help  
-i, --info  
-x, --usersdec  
-y, --usersenc]
```

#### **-d, --domain**

Reencrypts only Domain Server files.

#### **-t, --trend**

Reencrypts only Performance Trend files.

#### **-a, --all**

Reencrypts Domain Server and Performance Trend files.

#### **-h, --help**

Displays usage information.

#### **-i, --info**

Displays whether FIPS is enabled. Use this parameter at any time to confirm whether a switch to FIPS mode or non-FIPS mode was successful.

#### **-x, --usersdec**

Un-encrypts the user access control file (users.dat) for editing.

#### **-y, --usersenc**

Reencrypts the user access control file (users.dat) after editing.

## Active Directory Management

Active Directory Management (ADM) provides FIPS 140-2 support using the ETPKI library to encrypt user passwords. The ETPKI wraps the FIPS 140-2 validated RSA BSAFE Crypto-C Micro Edition (ME) encryption library. Passwords are encrypted using the AES encryption algorithm with a 256-bit strength encryption key.

### Data Encrypted

The following Active Directory Management data is encrypted:

#### Active Directory credentials

Includes the password required to access Active Directory forest and domain information. This information is created by the ADEM Connection Settings tool (ADEMConfig) and decrypted by the ADEM Service.

**File:** Stored by ADEMConfig in a file named forest.py and read from the file by the ADEM Service.

### Data Encryption Key

Active Directory Management encrypts passwords with a 256-bit data encryption key. The encryption key is contained in the file util\_crypt.pyc, which is part of library.zip. This file is in the installation directory at the following location:

*Install\_path\SC\CCS\ADEM\bin*

We recommend securing the file using Operating Systems file security, and giving read permissions to only ADM administrators and the local system account.

### Installation and Migration Considerations

When you perform a clean installation of ADM, it automatically uses FIPS mode encryption.

When you upgrade from a CA NSM r11.1 patch that contains ADM, the configuration utility ademininstall.exe automatically converts passwords in the forest.py file to the new encryption format. A new parameter 'filever = 2' is added to forest.py to specify the new format.

When you uninstall ADM, forest.py is removed automatically.

**Note:** For more information about installing, upgrading, or uninstalling ADM, see the ADM documentation.

## Convert Password File to FIPS Encryption

ADM automatically uses FIPS encryption in the forest.py file upon installation and upgrade. However, you may need to manually convert forest.py to use FIPS encryption if restoring an r11.1 backup file, for example.

To manually convert forest.py to use FIPS encryption, run the ADEMConfig utility.

The tool converts the file automatically on startup and notifies you of the change in a message box.

**Note:** For more information about the ADEMConfig utility, see the *Inside Active Directory Management* guide.

## Agent Technology

The Agent Technology component of CA NSM provides FIPS 140-2 encryption support using the RSA Crypto-C ME encryption library through the ETPKI. Agent Technology uses the AES CFB algorithm with a 128-bit strength encryption key, the 3DES CBC algorithm with a 64-bit strength encryption key, and the SHA-1 algorithm to encrypt configuration files, communications with other CA NSM components, and SNMPv3 communications.

## Data Encrypted

The following Agent Technology data is encrypted in a FIPS 140-2 compliant manner:

### **Configuration files**

Includes files containing sensitive data such as passwords.

### **SNMPv3 communications**

Includes encrypted content transported over a network using SNMPv3. This encryption is done using the AES CFB algorithm with a 128-bit key or the 3DES CBC algorithm with a 64-bit key.

### **Component communications**

Includes all data exchanges with other CA NSM components.

## Data Encryption Key

The data encryption keys used with Agent Technology are derived from passwords that protect the key and the sensitive data protected by the key. You can further augment key protection by creating an additional DLL to add another layer of protection to your encryption environment.

## Installation Considerations

You must turn on 'FIPS-only' mode for Agent Technology to use FIPS 140-2 compliant libraries and algorithms to protect sensitive data and communications.

## Migration Considerations

After an upgrade, all new files are encrypted using FIPS algorithms. Encrypted Agent Technology data from previous versions works with CA NSM r11.2 after an upgrade. Some of the existing data is reencrypted automatically during the initial run-time, while other data requires you to reencrypt it manually. If you want to make this data FIPS 140-2 compliant.

## Common Communications Interface

The CA Common Communications Interface (CACCI) component of CCS r11.2 provides FIPS 140-2 encryption support using the ETPKI. ETPKI provides an abstraction layer from the underlying certified encryption libraries.

CCI uses RSA's BSAFE Crypto C Micro Edition Version 2.0 for encrypting data sent over the Common Communications Interface. The encryption algorithm used is the AES Cipher Block Chaining (CBC) algorithm with a 256-bit strength key, or any other cipher suite used by the ETPKI that is FIPS compliant and compatible with CA OpenSSL TLS.

**Note:** FIPS encryption is only supported for CCI through the remote service and is unavailable through the quenetd.

## Data Encrypted

Data exchanged between CA NSM components using the Common Communications interface is encrypted using FIPS compliant libraries and algorithms. Everything sent from the local to remote hosts, including the data itself, CCI headers, and user data, is encrypted.



## Turn on FIPS Mode

For CCI to use FIPS compliant libraries to encrypt data exchanged between components, you must turn on FIPS mode.

### To turn on FIPS mode

1. Set the `CA_FIPS1402_ENABLE` environment variable value to 1.  
FIPS mode is turned on.
2. Restart the CCI remote daemon.  
CCI encrypts all communications using FIPS compliant methods.

## Installation Considerations

You must do the following to enable FIPS compliant encryption after installing CCI:

- Turn on FIPS mode.
- Set the `CAI_CCI_SECURE` environment variable's value to YES to enable SSL support.
- Restart the remote CCI daemon.

For more information about enabling SSL support, see the chapter "Securing CA NSM".

## Management Command Center

While the Management Command Center (MCC) is not fully FIPS 140-2 compliant, several types of protection and communication use or can be configured to use FIPS compliant encryption:

- Protection for login credentials used to access web-based applications has been improved through the use of a FIPS 140-2 validated encryption algorithms. RSA's BSAFE Crypto-J JCE Provider Module is used for encrypting the login credentials. MCC uses the AES Cipher Block Chaining (CBC) algorithm with a 128-bit strength key for encryption.
- You can configure Tomcat to use TLS (Transport Layer Security) protocol to protect data exchanged between the MCC and web-based applications. By default, Tomcat uses the Java Secure Socket Extension from Sun (JSSE) for implementation of TLS support. While JSSE has not received FIPS 140-2 certification, you can configure it to use FIPS 140-2 compliant encryption algorithms. The algorithm used for encryption with TLS is negotiated between the server and the client, although you can configure to restrict use to only certain algorithms. You can also configure MCC to use TLS to protect data exchanged with WorldView and Event Management through CA Messaging (CAM).
- Communications between MCC and Alert and Event Management that are handled by DIA use SSL/TLS protocol. DIA uses the Java Secure Socket Extension (JSSE) from Sun for implementation of SSL/TLS support. The algorithm used for encryption with SSL/TLS is negotiated between the server and the client. DIA does not support site configuration of the cipher suites to be enabled in JSSE.

## Data Encrypted

The following MCC data is encrypted in a FIPS 140-2 compliant manner:

### **Web application login credentials**

Includes the user names and passwords entered to access the following web-based applications: AEC Web Editor, Configuration Management (UCM), Adaptive Dashboard Services (ADS), Discovery Configuration, eHealth Report Server, Web Reporting Service (WRS), and Unicenter Service Desk. These credentials are encrypted and stored in the MCC WebApplications directory located under the user's home directory whenever the user indicates that the credentials should be remembered.

**File:** savedsettings.xml

### **Communications with web applications**

Includes data sent between the MCC and other web applications that are hosted by Tomcat. These applications include AEC Web Editor, UCM, ADS, Discovery Configuration, eHealth Report Server, WRS, and Unicenter Service Desk.

**Communications with AIS providers**

Includes data sent between the MCC and WorldView and Event Management using CA Messaging (CAM).

**Communications with DIA providers (Alerts and Console logs)**

Includes data sent between the MCC and Alert and Event Management using DIA.

**Data Encryption Key**

The 128-bit private key used for data encryption of web application login credentials is generated by MCC and stored in a keyfile located under the user's home directory, allowing for access level control using operating system level file security.

For communications between the MCC and Tomcat-hosted applications, you can configure Tomcat to use the Java Key Store (JKS) or a PKCS12 keystore. The keystore can be located anywhere on the server. You can point to the keystore using the keystoreFile attribute. Secure the keystore using operating system level file security.

DIA uses a 256-bit private key to encrypt DIA-based communications between the MCC and Alert and Event Managers. The keystore is located in the configuration directory for DIA, and you can secure the keystore using operating system level file security.

**Installation Considerations**

During installation of MCC, a private Java Runtime Engine (JRE) is installed for use by MCC and other CA NSM components. The jsafeJCEFIPS.jar file containing the RSA BSAFE Crypto-J JCE Provider Module is installed in the extensions directory of this private JRE. At run-time, the MCC triggers the JRE to load the BSAFE JCE provider module. Therefore, FIPS compliant encryption of login credentials for web-based applications requires no post-installation steps.

Encryption of communications between the MCC and web applications must be enabled manually after installation by configuring Apache Tomcat to use TLS encryption. For detailed step-by-step instructions for configuring Apache for TLS, see Configure Tomcat with SSL in the *Implementation Guide*.

When configuring Tomcat to use SSL, you can also specify to use only FIPS 140-2 compliant algorithms by setting the "ciphers" attribute of the Connector element and listing the allowed algorithms. Specify the ciphers using the JSSE cipher naming convention. For more information about configuring Tomcat to use specific ciphers, see the Apache Tomcat and Sun JSSE web sites for details.

For encryption of communications between MCC and the WorldView and Event Management components, you must configure CA Messaging (CAM) to use TLS encryption. For more information, see [Configure CAM to Use TLS Encryption](#).

You must enable encryption of DIA-based communications between MCC and Event and Alert Managers by configuring DIA to use FIPS 140-2 compliant encryption. For detailed steps on how to configure DIA encryption, see [Configure Communications for Encryption in the \*Implementation Guide\*](#).

## Migration Considerations

Login credentials encrypted using a prior release of MCC are decrypted using the previous algorithm upon the first launch of MCC and automatically reencrypted when the MCC session ends.

## Configure CAM to Use TLS Encryption

For communications between the MCC and WorldView and Event Management to be encrypted, you must configure CAM, which handles the data transport between the MCC and these components, to use TLS encryption.

For details about the command line utility required to configure CAM for TLS, see [Configure SSA to Enable CAM to use SSL/TLS](#) (see page 481).

For more information about configuring CAM, see the CA Message Queuing Service (CAM) section in the chapter "Using Ports to Transfer Data."

## Turn Off Password Caching for Event Management and WorldView Credentials

Remembering of login credentials for the WorldView and Event Management components is accomplished through CA Messaging (CAM) using password caching, which is not FIPS compliant. You may want to turn off this password caching if you are concerned about the level of security it provides. The password caching is turned on by default.

### To turn off password caching for CAM login credentials

1. Access the `ji.cfg` file.
2. Set the following parameter to a non-zero value:  
`default.SEC.bypass_cache`  
Save and close the file.  
Password caching is turned off.

## Unicenter Management Portal

While Unicenter Management Portal does not fully support FIPS 140-2 encryption, the encryption algorithm has been improved for r11.2. Unicenter MP uses the encryption libraries installed, registered, and maintained by the CleverPath Portal PEK, which has been upgraded to r4.72.

Unicenter MP supports remembering credentials with AES encryption for persistent user names and passwords using the RSA BSAFE Crypto-J JCE Provider Module version 3.5. Unicenter MP uses the AES Cipher Block Chaining (CBC) algorithm with a 128 bit strength key for encryption.

You can also configure Apache Tomcat to use TLS (Transport Layer Security) encryption to protect data exchanged between Unicenter MP and web-based applications that Tomcat hosts. By default, Tomcat uses the Java Secure Socket Extension from Sun (JSSE) for implementation of TLS support. The algorithm used for encryption with TLS is negotiated between the server and the client, although you can configure to restrict use to only certain algorithms.

## Data Encrypted

The following Unicenter MP data is encrypted using FIPS 140-2 compliant libraries:

### **Login credentials**

Includes persistent user names and passwords.

**File:** database.properties

### **Database credentials**

Includes database user names, passwords, and JDBC URLs entered in Unicenter MP to access MDB data.

**File:** database.properties

### **CA NSM credentials**

Includes Unicenter NSM login names and passwords entered in Unicenter MP to establish connections with CA NSM components.

### **Communications with web applications**

Includes data sent between Unicenter MP and other web applications that are hosted by Tomcat. These applications are AEC Web Editor, Configuration Management (UCM), Adaptive Dashboard Services (ADS), Discovery Configuration, eHealth Report Server, Web Reporting Service (WRS), and Unicenter Service Desk.

## Data Encryption Key

The 128 bit private key used for data encryption of persistent user names and passwords, database passwords, and CA NSM passwords is stored in the `INSTALL_PATH\CCS\WRS\webpages\WEB-INF\lib\ca-common-crypto-12.0.39.0.jar` resource file. You can secure this file on the local file system using operating system level file security.

For communications between Unicenter MP and Tomcat-hosted applications, you can configure Tomcat to use the Java Key Store (JKS) or a PKCS12 keystore. The keystore can be located anywhere on the server. You can point to the keystore using the `keystoreFile` attribute. Secure the keystore using operating system level file security.

## Installation Considerations

Encryption of communications between Unicenter MP and web applications must be enabled manually after installation by configuring Apache Tomcat to use TLS encryption. For detailed step-by-step instructions for configuring Apache for TLS, see *Configure Tomcat with SSL* in the *Implementation Guide*.

When configuring Tomcat to use SSL, you can also specify to use only FIPS 140-2 compliant algorithms by setting the "ciphers" attribute of the Connector element and listing the allowed algorithms. Specify the ciphers using the JSSE cipher naming convention. For more information about configuring Tomcat to use specific ciphers, see the Apache Tomcat and Sun JSSE web sites for details.

## Web Reporting Server

While Web Reporting Server (WRS) does not fully support FIPS 140-2 encryption, the encryption algorithm has been improved for r11.2. WRS uses the encryption libraries installed, registered, and maintained by the CleverPath Portal PEK, which has been upgraded to r4.72.

WRS supports remembering credentials with AES encryption for persistent user names and passwords using the RSA BSAFE Crypto-J JCE Provider Module version 3.5. WRS uses the AES Cipher Block Chaining (CBC) algorithm with a 128 bit strength key for encryption.

## Data Encrypted

The following WRS data is encrypted using FIPS 140-2 compliant libraries:

### **Login credentials**

Includes persistent user names and passwords for components that use WRS.

### Data Encryption Key

The 128 bit private key used for data encryption of persistent user names and passwords is stored in the `INSTALL_PATH\CCS\WRS\webpages\WEB-INF\lib\ca-common-crypto-12.0.39.0.jar` resource file. You can secure this file on the local file system using operating system level file security.

### Installation Considerations

WRS requires a post-installation configuration of Apache Tomcat to use TLS to encrypt Tomcat-enabled communications using FIPS-compliant libraries. For detailed instructions for configuring Tomcat to use TLS, see *Configure Tomcat with SSL* in the *Implementation Guide*.





# Appendix C: Managing Traps Using the Trap Manager

---

This section contains the following topics:

[Trap Daemon](#) (see page 433)

[Trap Filters](#) (see page 434)

[Local Versus Remote Installation](#) (see page 434)

## Trap Daemon

The Trap Manager lets you perform sophisticated trap database and trap filter file management. You can use the Trap Manager to manage trap information and translation messages stored in the Trap Database and trap filters stored in the trap filter file.

The CA Trap Daemon (CATRAPD) receives Simple Network Management Protocol (SNMP) traps on UDP port 162. These SNMP traps contain critical information about the latest status of your network environment, including the network itself and devices on that network. Since this information is received in the form of Management Information Base (MIB) variables and their numeric values, it is difficult to understand offhand. The Trap Daemon reads the MDB trap tables, which contain all trap information and translation messages, and translates SNMP traps into meaningful, easy to understand messages. These translated traps appear on the CA NSM Event Console.

For every incoming trap, the Trap Daemon also searches the trap filters file for any filters that apply. If the specified filter criteria is satisfied, the trap is dropped from further processing and does not appear on the Event Console. This can be very helpful if you are only interested in certain traps.

**Note:** By default, CATRAPD does not translate the trap information it receives. You must configure CATRAPD to use the Trap Translation Database. You can also enable or disable translation for specific traps.

## Trap Filters

The Trap Manager lets you easily manage trap filters stored in the trap filter file. You can use trap filters to filter out traps from appearing on the Event Console. For every incoming trap, the Trap Daemon searches the trap filters file for any filters that apply. If the specified filter criteria are satisfied, the trap is dropped from further processing and does not appear on the Event Console. This can be very helpful if you are only interested in certain traps. You can use the Trap Manager to view, add, edit, or delete trap filters.

## Local Versus Remote Installation

You can install the Trap Manager on the same computer as Event Management or on a different computer. With a local installation, all of the Trap Manager features are available. If you install on a different computer, the Trap Manager can connect to the database remotely however, the Trap daemon and filter management features will not function. The table that follows explains the differences between local and remote installation.

<b>Feature</b>	<b>Local Installation (Trap Daemon and Trap Manager on same Computer)</b>	<b>Remote Installation (Trap Daemon and Trap Manager on Different Computers)</b>
File - Add Vendor, MIB File, Trap, Rename, and Exit	Functional	Functional
View MIBs and Refresh	Functional	Functional
Tools Import, Find, Backup, and Restore	Functional	Functional
Trap Daemon Refresh Cache, Shutdown, and Start	Functional	Not functional because commands required to perform these tasks are executed only locally.
Help	Functional	Functional
Trap Tab (right pane)	Functional	Functional
Filter Tab (right pane)	Functional	Not functional because the filter definitions are stored locally in a flat file, not in the database.

# Appendix D: Managing Cisco Devices Using Cisco Integration

---

This section contains the following topics:

[Analyzing CISCO Integration](#) (see page 435)

[Cisco Device Recognition](#) (see page 435)

## Analyzing CISCO Integration

Cisco Integration supports all Cisco devices (routers and switches) available in the Cisco Network Management Integration Data Bundle (NMIDB). NMIDB is published by Cisco and can be downloaded. Cisco Integration uses the NMIDB file to dynamically update new Cisco devices. You can configure automatic device update or manual device update. Based on the NMIDB, Cisco Integration does the following:

- Creates class definitions for the Cisco model class.
- Updates Agent Technology policies to monitor these Cisco devices.
- Creates icon/image files to uniquely display these devices using either the Unicenter 2D Map or the Management Command Center.
- Updates Performance Management to create reports for these Cisco devices.

## Cisco Device Recognition

Cisco devices monitored by CA NSM have generic Cisco icons when they appear on the Unicenter 2D Map or on the Management Command Center. With Cisco Integration, Cisco devices appear with unique icons to make the device identifiable. These devices have default Cisco router/switch icons as their base and device model labels in the upper right corner.



# Appendix E: Replicating Objects in the WorldView Repository

---

This section contains the following topics:

- [Analyzing Repository Bridge](#) (see page 437)
- [How Repository Bridge Works](#) (see page 438)
- [Repository Bridge Architectures](#) (see page 439)
- [Repository Bridge Components](#) (see page 442)
- [Repository Bridge Supported Platforms](#) (see page 444)
- [Repository Bridge in a Distributed Organization](#) (see page 444)
- [Repository Bridge for a Restricted View of Resources](#) (see page 445)
- [Repository Bridge for Problem Notification](#) (see page 445)
- [Troubleshooting](#) (see page 445)
- [How to Create a Bridge Configuration File \(Windows Only\)](#) (see page 446)
- [Bridging Rules \(Windows\)](#) (see page 448)
- [Bridging Objects to A Repository Where a DSM is Running](#) (see page 448)
- [Start the Bridge Configuration GUI \(Windows Only\)](#) (see page 448)
- [Manage Repository Bridge Instances Using a Windows Service \(Windows Only\)](#) (see page 449)
- [Create a Configuration File \(UNIX/Linux\)](#) (see page 450)

## Analyzing Repository Bridge

The Repository Bridge component lets you replicate a subset of the managed objects in a source WorldView repository (that is, the WorldView data in the MDB) to a destination repository and maintain those objects. This subset of objects is determined by a set of user-defined rules, known as bridging policy. Once activated, a bridge instance replicates objects in the source repository that comply with the bridging policy, and then continues to monitor the source repository for the following events:

- Changes to bridged objects, updating the destination repository with those changes or removing a bridged object if the change means it no longer conforms to the bridging policy.
- Changes to non-bridged objects, bridging those objects if the change means they now conform to bridging policy.

By using a number of bridges in parallel, a single source repository can be bridged to many destination repositories or many source repositories can be bridged to a single destination repository. A many-to-one bridge configuration enables central monitoring of significant objects in a distributed enterprise environment. Classes must exist on a destination repository before bridging is done.

**Note:** To use Repository Bridging on UNIX and Linux, the destination MDB must be hosted in an Ingres database because Repository Bridge uses an Ingres client to connect to the destination MDB on the remote server.

## How Repository Bridge Works

During initialization, the Repository Bridge performs the following operations:

1. Scans the configuration and parses the bridging policy.
2. Initializes and begins logging.

Repository Bridge provides complete logging facilities with a configurable level of detail.

3. Establishes repository connections.

Establishes connections to both the source and destination repository. The Repository Bridge must obtain a connection to both databases to operate.

4. Registers for notifications.

Registers the Repository Bridge with the source repository for the receipt of object state change notifications, which are used to drive the object replication process.

5. Builds Bridging decision binary.

Builds an in-memory binary of the bridging policy. This process is analogous to the Event Manager or Security Manager building in-memory binaries of message records and actions. When notification is received, the Repository Bridge can rapidly analyze the notification to see if it should bridge any information to the destination repository.

6. Synchronizes repositories.

Scans the source repository for all objects that satisfy the bridging policy criteria, and then cross checks the destination repository to ensure the bridged objects exist with the same property values. Any previously bridged objects that no longer conform to policy are deleted from the destination repository.

After these startup procedures are complete, the Repository Bridge is driven by notifications from the source repository. When a notification is received, the object to which it relates is analyzed to determine the following:

- The object is bridged, so the notification is bridged.
- The object should be bridged as a consequence of the notification, so the object is replicated in the destination repository.
- The object is bridged, but no longer conforms to bridging policy as a consequence of the notification, so the replicated object is deleted.

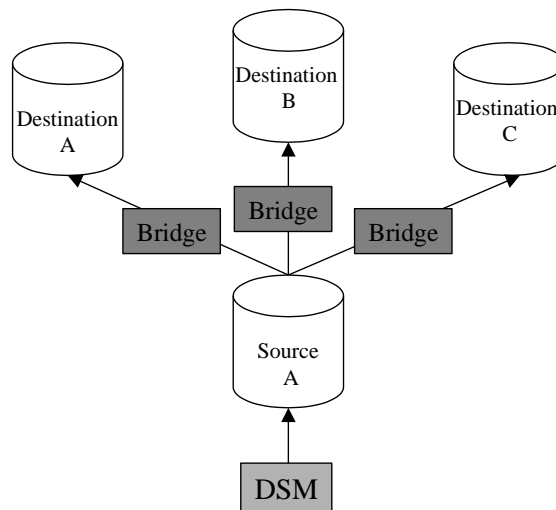
## Repository Bridge Architectures

You can implement Repository Bridge using two types of bridging architecture:

- One-to-one/one-to-many or “fanout”
- Many-to-one or “aggregation”

### Fanout Architecture

The fanout architecture consists of one source repository and one or more destination repositories. Bridge instances run between the source repository and each of the destination repositories.



Use the following guidelines when you are considering using a fanout architecture:

- The number of bridge instances running on a host affects CPU utilization on that host.

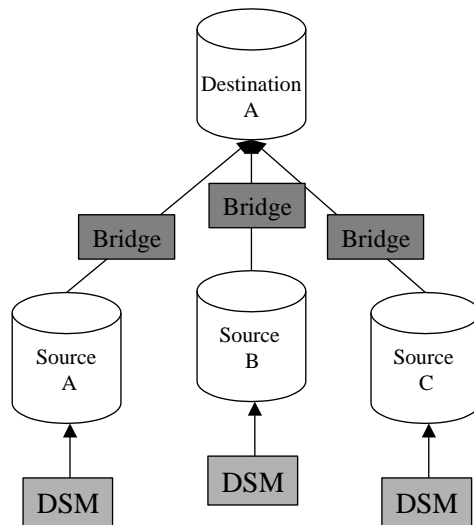
Each bridge instance independently processes notifications from the source repository. Therefore, the more activity in the repository, the more objects being bridged, and the greater the load on the host running the instances.

- The number of bridge instances associated with a source repository increases the load on the source database server.

Each destination repository in this architecture requires a separate bridge instance, which runs independently of the other instances associated with the source repository. This situation causes an increased load on the source database server as the server processes requests and queries from those instances.

## Aggregation Architecture

The aggregation architecture consists of several source repositories and one destination repository. Bridge instances run between each of the source repositories and the destination repository.





Use the following guidelines when you are considering using an aggregation architecture:

- Carefully monitor the cumulative number of objects bridged from the source repositories to the destination repository.

If several source repositories exist, the number of objects in the destination repository can quickly exceed the recommended limits. The same guidelines provided for a standard repository should be followed for a bridged repository. To avoid problems, obtain estimates for the number of bridged objects from each source repository before implementation.

- Bridging duplicate objects to a repository causes errors.

If you have a duplication of objects across source repositories (that is, objects have the same name), and those objects are bridged to the same destination repository, errors can occur.

## How to Determine Which Architecture to Use

The wrong architectural approach may have a significant impact on the performance of your bridged system, so you must consider several factors before making a decision. The following guidelines can aid your decision process:

- A bridge instance should always be located on the same machine as the source repository.

Significantly more communication occurs between the Repository Bridge and the source. If this communication occurs across a network you may encounter an increase in network loading and a delay in Bridge processing.

- The number of objects and activity in the source repository affects the performance of the Repository Bridge.

The more objects and activity in the source repository, the more information the Repository Bridge must process.


- The number of objects being bridged affects performance throughout the system.

As the number of bridged objects increases, you may have an increase in processing by the bridge, an increase in network traffic between the Repository Bridge and the destination repository, and an increase in the load on the destination database server.

## Repository Bridge Components

Repository Bridge consists of three closely related components:

- **Bridge configuration** lets you create, modify, and delete bridge configuration (.tbc) files.  
On UNIX/Linux, these files are located in `$CAIGLBL0000/wv/config/bridge`.
- **Bridge instance** is the process that actively implements the bridging policy defined in a bridge configuration (.tbc) file.
- **Bridge control** lets you initiate and stop bridge instances available on the local host.  
On UNIX/Linux, it also creates a file, `$CAIGLBL0000/wv/config/bridge/.bridgePid`, that stores the process IDs of each running bridge instance.



### Bridge Configuration

Bridge configuration lets you develop and edit bridging policy, which is maintained in a .tbc file. Although it is possible to write or edit .tbc files manually (they are flat ASCII text files), we recommend that you use the interfaces provided to ensure the accuracy and consistency of the policy.

On Windows, the Repository Bridge Configuration GUI lets you define bridging policy for a bridge instance. This interface generates the .tbc files, which are stored in the Bridge Config. Directory specified during installation.

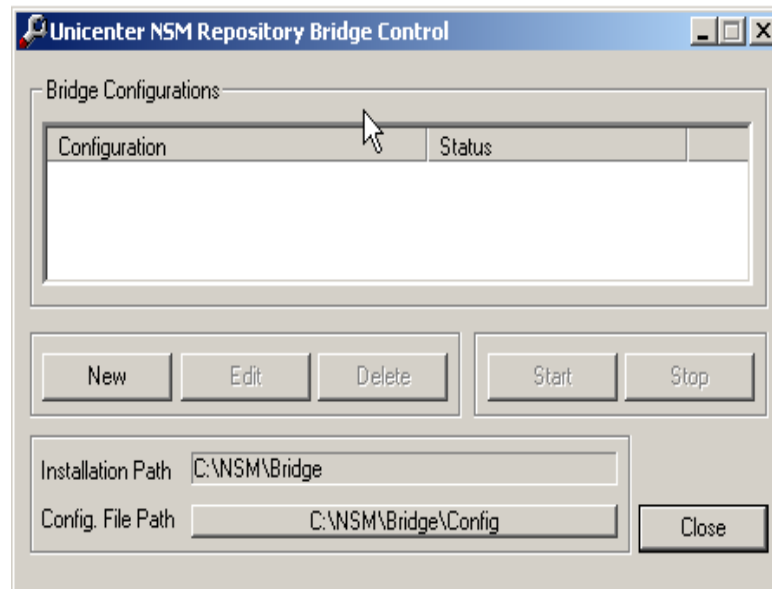
On UNIX/Linux, use the `bridgecfg` command to create bridging policy. This interface generates the .tbc files and saves them in `$CAIGLBL0000/wv/config/bridge`.

## Bridge Control

The Bridge Control provides an operator with a means of starting, stopping, and displaying the status of bridge instances available on the local host. On Windows, the Bridge Control also lets you start the Configuration GUI where you can edit or delete existing configurations.

The Bridge Control can be accessed through the Repository Bridge Control GUI or the command line. The Repository Bridge Control GUI displays information about the way the bridge was configured at installation, including the installation path and the path under which configurations are stored.

On Windows, start the Repository Bridge Control GUI by selecting Start, Programs, CA, Unicenter, NSM, WorldView, Bridge, Bridge Control.



The Repository Bridge Control also has a command line interface, `bridgectrl`, through which you can start and stop any number of configured instances, letting you write scripts or batch files to control instances on a particular host without user intervention.

On UNIX/Linux, use the `bridgectrl` command from the UNIX/Linux command line to start, stop, and display bridge instances on the local host.

The `bridgectrl` command creates a file, `$CAIGLBL0000/wv/config/bridge/.bridgePid`, that stores the process IDs of each running bridge instance.

## Bridge Instances

A bridge instance implements the bridging policy defined in a bridge configuration file. A bridge instance is a generic shell that derives its configuration from a .tbc file. Only one instantiation of a given bridge configuration can be running at any time, which means that you cannot run several instances for the same source and destination repository. However, you can run any number of *different* configuration instances on a particular host.

Depending on the logging level set in the configuration, you can monitor the status and activity of a bridge instance by inspecting the log file generated locally. In addition, you can monitor startup, shutdown, and critical error messages remotely from a designated Event Manager Message Console.

## Repository Bridge Supported Platforms

The following platforms are supported:

- Windows
- Linux

## Repository Bridge in a Distributed Organization

Your worldwide corporation wants to use CA NSM to manage your enterprise. Your organization consists of a global headquarters plus a number of regional offices. Although you want to permit a degree of local autonomy in each of the regions, you also want to have a central summary of the key business processes throughout the organization to ensure that you are aware of any emerging problems that have potential impact on the global business.

In this situation, you can deploy the Repository Bridge using an aggregation architecture. Bridging the key resources from each regional repository to the central repository facilitates the representation of global business processes, providing the operators at that level with enough information to determine the current operation state of the organization without overloading them with low level regional detail.

For example, the communications infrastructure in any organization is key to its capacity to function as a whole. The Repository Bridge could be configured to bridge the MS Exchange resources from each region into the central repository so that a Global Communications Business Process View could be created and maintained.

## Repository Bridge for a Restricted View of Resources

Your company has a central repository populated with all your customers' managed resources, from which you can do administrative tasks. A number of those customers have expressed an interest in passively monitoring the state of their own resources, to ensure your company is doing its job. For obvious reasons, your company could not simply provide a remote WorldView onto the central repository for each customer, since you cannot restrict access to specific resources in that repository. Any customer with access could browse the repository and, potentially, the resources of their competitors.

In this situation, you can deploy the Repository Bridge using a fanout architecture. By bridging the resources specific to each customer into a repository to which that customer has exclusive access, you can ensure that their requirements are met without breaching any contractual obligations elsewhere.

## Repository Bridge for Problem Notification

Your large utilities company has a number of regional offices, each managed locally by CA NSM. Your company also has a national Call Center where all problems are reported. If the problem cannot be resolved over the phone, a local contract services organization is notified which provides the necessary support at the local level.

In this situation, you can deploy the Repository Bridge using an aggregation architecture. By placing a repository at the Call Center level, and bridging all objects with a status of critical from the regional level up, operators in the Call Center can monitor and track problems throughout the organization. Furthermore, by populating the Call Center repository with only those resources showing a problem, the operators are not distracted by problem free resources, and can quickly focus their attention where required.

## Troubleshooting

Check the Repository Bridge log files for errors. If the log file does not contain any reported errors, the source of the problem may be in the way that the Repository Bridge has been configured. If errors are present, this may imply a problem with the source or destination repository, or the Repository Bridge itself.

## View Repository Bridge Log Files

You can view the log file if a bridge instance fails to start, shuts down unexpectedly, or displays any unusual behavior. Each instance has its own log file that can contain a large amount of information about the operational state of the Repository Bridge.

To view Repository Bridge log files, open any standard text editor. Log file names use the following format:

*Configuration-Name\_DDMMYYYY.LOG*

## How to Create a Bridge Configuration File (Windows Only)

To create a bridge configuration file using the Bridge Configuration GUI, follow these steps:

1. Identify the source repository from where the objects are copied.
2. Identify the destination repository to where the objects are copied.
3. (Optional) Configure targeting.

Specifying a target object helps you manage bridged objects in the destination repository. Targeting lets you specify a target object under which all bridged objects will be placed. If targeting is not specified within a configuration, bridged objects are placed in the Managed Object folder of the destination repository.

4. (Optional) Configure batch mode operation.

Although you would typically run Repository Bridge in real-time mode, to ensure updates are replicated as fast as possible, batch mode provides some additional features that are not available in a typical operation. However, the overhead associated with the additional calculations involved, particularly when applying summary state and bridging based on propagated state, may be prohibitive. Carefully consider the update period, ensuring that you permit enough time between updates for the repository synchronization to be completed.

5. Define the bridging policy by creating bridge configuration rules.

Bridging policy is defined as a set of rules. Bridge rules consist of property and values pairs. The value can be specified in terms of an explicit alphanumeric value, a mix of alphanumeric and wildcard characters, a numeric range set, or a combination. Rules are specific to a class of object. Identifying the class to which the rule relates determines the properties that can then be specified as rule criteria.

Examples of each property/value format are as follows:

<b>Format</b>	<b>Example</b>
Explicit	class_name: WindowsNT_Server, severity: 5
Wildcard	class_name: WindowsNT_Server, address: 172.24?.*.*
Range Set	class_name: WindowsNT_Server, severity: {0,2,4-6}
Combined	class_name: WindowsNT_Server, address: 172.{24-28,35}.*.*

6. Configure the logging facility, which lets you determine where logs are written and the level of logging you want to write.

Each bridge instance has its own log file that can contain a large amount of information on the operational state of the Repository Bridge. If a bridge instance fails to start, shuts down unexpectedly, or displays any unusual behavior, you can access the log file to help you determine what the problem is.

7. Configure Event Management integration.

You can send Repository Bridge events to the Event Management Console, specify the Event Management node where start up, shut down, and other messages can be sent, and specify a message prefix to enable easy identification of the instance from which the event message was initiated.

8. Configure startup options.

9. Save the configuration file.

If the configuration definition is successful, the Bridge Configuration GUI closes and the Bridge Control interface updates its list of available configurations.

## Bridging Rules (Windows)

You define your Bridging policy by creating a set of bridging rules that consist of property and value pairs. The value can be specified in terms of an explicit alphanumeric value, a mix of alphanumeric and wildcard characters, a numeric range set, or a combination. Rules are specific to a class of object. Identifying the class to which the rule relates determines the properties that can then be specified as rule criteria.

## Bridging Objects to A Repository Where a DSM is Running

When multiple MDBs are bridged together and each MDB has one or more DSMs, you must ensure that the assigned list of nodes to manage for each MDB's DSM is unique.

That is, when "bridging" one MDB to another, you must ensure that the DSM being "bridged" does not have overlapping DSMIPScoping entries. You do not want the "bridged" MDB DSM to manage the same subnet. Otherwise, when the "bridged" nodes are created in the other MDB, that MDB will begin to manage that node.

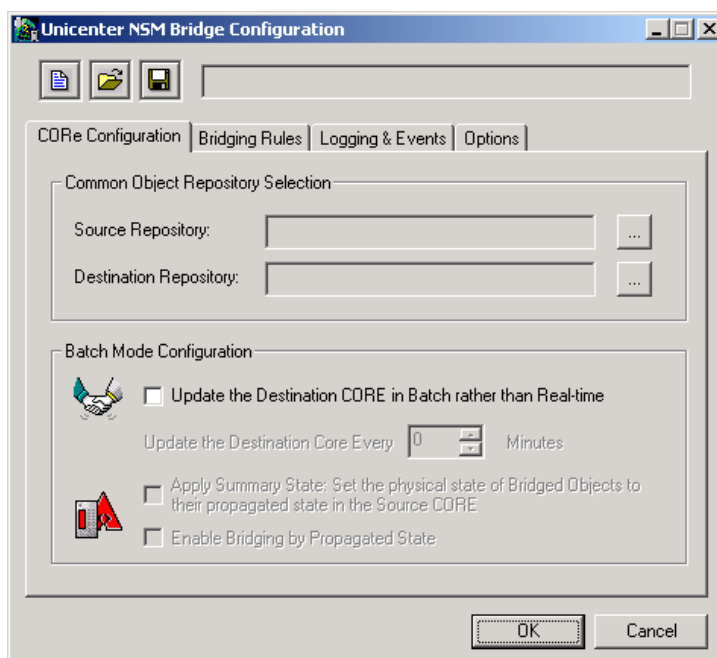
## Start the Bridge Configuration GUI (Windows Only)

The Bridge Configuration GUI lets you develop and edit bridging policy.

To start the Bridge Configuration GUI, select Start, Programs, CA, Unicenter, NSM, WorldView, Bridge, Bridge GUI.



The following window appears:



## Manage Repository Bridge Instances Using a Windows Service (Windows Only)

An optional Repository Bridge service lets you manage configured bridge instances as a Windows service. Using this service, you can automatically stop and restart configured bridge instances during a system reboot. For example, to stop the service from the command line, enter:

```
NET STOP "Unicenter Repository Bridge"
```

This command also stops all instances of the Repository Bridge currently running on the hosts that are under the control of the service.

You can use the bridge command from a command line to control bridge instances.

## Create a Configuration File (UNIX/Linux)

A Repository Bridge configuration file contains all of your bridge configuration rules. You specify which objects you want to bridge in the configuration rules. You can create rule files manually, but we recommend that you use the `bridgecfg` command to create your bridging policy. If you do create a rule file manually, you can use the `-f` parameter to append the rules in the rule file to the configuration file.

See the sample configuration file, `sample.tbc`, in the `$CAIGLBL0000/wv/config/bridge` directory.

### To create a bridge configuration file

1. Open a command line on the computer where the source repository resides.
2. Run the following command:

```
bridgecfg [-f rule_file] [-c dest=desination_repository] [Ruserid=userid]  
[Rpassword=password] [Rinstid=instance_id]
```

For more information about specifying parameters, see the online *CA Reference*.

You are prompted to choose rule options if you do not specify an existing rule file with the `-f` parameter.

The names of the source repository and the destination repository name are saved in the configuration file. The bridge configuration file is saved as *destination-name.tbc*. Configuration files are written to `$CAIGLBL0000/wv/config/bridge`.

**Note:** A field called `vspactive` is automatically created in your configuration file. It is not supported on UNIX/Linux, but it is there for compatibility with Repository Bridge on Windows.

## Rule File Parameters for UNIX/Linux

When you run the `bridgecfg` command, you are prompted to choose rule parameters if you do not specify an existing rule file with the `-f` option. See the sample rule file called `ruleTemplate` in the `$CAIGLBL0000/wv/config/bridge` directory.

Rule file parameters include the following:

**active**

Specifies that the rule is active.

**rulelabel**

Specifies the label of the rule.

**inclevelsup**

Includes this many levels upward. Specifies the level of parent inclusions. This option determines the number of parent branches bridged when a compliant object is bridged.

**inclevelsdwn**

Includes this many levels downward. Specifies the level of child inclusions. This option determines the number of child branches bridged when a compliant object is bridged. The bridging of child inclusions ensures that any status changes are correctly propagated for objects in the destination repository.

**class\_name**

Specifies the name of the class in which to bridge. Rules are specific to a class of object.

**address**

Specifies the address of the object in which to bridge.

**label**

Specifies the label of the object in which to bridge.

**name**

Specifies the name of the object in which to bridge.



# Appendix F: Support for DMI, MOM, and SCOM

---

This section contains the following topics:

[Desktop Management Interface \(DMI\)](#) (see page 453)

[Unicenter Management for Microsoft Operations Manager](#) (see page 457)

[Integration with Microsoft System Center Operations Manager \(SCOM\)](#) (see page 461)

## Desktop Management Interface (DMI)

The rapid change and growth in the computer industry has produced myriad hardware and software products that users and network administrators need to understand, manage, and inventory. In 1992, the Distributed Management Task Force, Incorporated (DMTF), began developing a standard framework for managing and tracking components installed on a PC. The DMTF established the Desktop Management Interface (DMI) specification, which was developed for users to determine the hardware and software installed on a PC.

The DMI consists of the following components:

### **Service Provider**

Consists of a set of programs that collects information from products, manages the MIF database, and passes information to management applications. Under Windows, the Service Provider starts running at system startup.

### **Management Interface**

Handles communications between the Service Provider and management applications, and allows a management application to query and edit the MIF database. The MI lets a management application view and edit the MIF database. Unicenter Support for DMI is a management interface.

### **Component Interface**

Allows a Component Instrumentation (CI) program that may come with a component to provide real-time component information. The CI handles communications between components and the Service Provider. The CI communicates with components that supply Component Instrumentation programs, which provide real-time access to the values for component attributes in the MIF database.

### **MIF Database**

Contains information about hardware and software installed on a system. Each product that adheres to the DMI specification is shipped with a MIF. Upon installation of the component, information in the MIF file is stored in the MIF database. A MIF file may contain static information about a component, or it may describe how component data can be obtained through the component instrumentation.

## **DMI Service Provider**

The DMI Specification requires that a DMI Service Provider be installed on the managed system. At the time of this writing, more and more desktop computer OEMs are shipping a DMI Service Provider as a pre-installed component on new systems. In addition, major hardware vendors and OEMs such as Intel and Microsoft provide desktop management kits on their respective Web sites. Many of these kits include a DMI Service Provider and are free to the OEM's customers.

Check with your hardware vendor, your motherboard manufacturer, and your software vendor to see if a DMI kit is available for your machine. Follow their instructions for installing the kit; install the WorldView DMI Manager and Agent, and then verify that the DMI Service Provider was installed correctly by starting the DMI Browser.

## Unicenter Support for Desktop Management Interface (DMI)

Unicenter Support for DMI is a desktop management tool that conforms to the Desktop Management Interface (DMI). Support for DMI extends the Desktop Management Interface (DMI) by providing a powerful software tool to help network administrators track installed hardware and software. Support for DMI strictly adheres to the Desktop Management Interface (DMI) specification to provide local machine management. In addition, DMI specification provides remote management of desktop PCs across a network.

Unicenter Support for DMI incorporates the latest technology to implement the DMI specification on your PC and has features that do the following:

- Display current data from various resources.
- Provide inquiry capability for local or remote hardware and software.
- Allow modification of DMI data by the user (unless the data is designated read-only).

## Install the DMI Manager and DMI Agent

To use the Unicenter Support for DMI functionality, you must install the DMI Agent and the DMI Manager. You can install these components from the CA NSM installation DVD. Install the DMI Manager on a computer that will be used to monitor remote computers; install the DMI Agent on each computer that you want to monitor remotely.

### **To install the DMI Agent and the DMI Manager on Windows platforms**

1. Insert the installation DVD into the DVD drive.

The installation program starts automatically, and the Unicenter Product Explorer window appears. The option for Installation Wizard for CA NSM is preselected.

2. Click Install.

The CA NSM Installation Wizard window appears.

3. Select "Install any or all CA NSM components" and click Next.

The Component Selection window appears.

4. Select WorldView DMI Manager and WorldView DMI Agent.

Click Next and continue with the installation. For more information about installation, see the *Implementation Guide*.

## Set SNMP Destinations in the CA DMI Agent

If installed on a remote PC, the CA DMI Agent converts DMI events into SNMP traps and can send them to multiple CA NSM event consoles.

### To define SNMP trap destinations using the DMI Browser

1. Choose Start, Programs, CA, Unicenter, NSM, WorldView, DMI, DMI Browser.

The DMI Browser appears.

In the left pane, expand the CA DMI Agent, and expand SNMP Trap Destinations.

The right pane displays the SNMP Trap Destinations group. It contains lines for sixteen SNMP trap destinations. The first destination is the destination you entered at installation.

2. Double-click the trap destination address under TCP/IP address.

The DMI-Browser - Edit attribute window appears.

3. Edit the trap destination address by changing the IP address in the Value field, and click OK.

The trap destination is updated.

4. To enable the trap destination, double-click the Disable attribute for the destination.

The DMI-Browser - Edit attribute window appears.

5. Use the Value field drop-down list to change the value to Enable, and click OK.

The trap destination is enabled. Repeat Steps 3-6 for each trap destination you want to change and enable.



## Unicenter Management for Microsoft Operations Manager

Unicenter Management for MOM (MOM Management) integrates CA NSM with Microsoft Operations Manager (MOM).

Microsoft Operations Manager delivers operations management by providing event management, proactive monitoring and alerting, reporting, and trend analysis. It helps administrators monitor and manage the events and performance of Windows 2000 or 2003 server systems. MOM is similar to CA NSM Event Management.

The integration between MOM and CA NSM provides a central location for performing management functions. You can see the status of MOM servers and managed PCs using WorldView, and you can change the status from there. You can also view open MOM alerts and CA NSM events in one location, the Event Console.

**Note:** CA NSM provides integration kits to both of Microsoft's management applications, Microsoft Operations Manager (MOM) and System Center Operations Manager 2007 (SCOM). Although the integrations to MOM and SCOM can coexist on the same management server, each one integrates only with its Microsoft counterpart.

### MOM Terminology

The following terms define industry standards related to Microsoft Operations Manager (MOM). Also included are definitions that are specific to MOM.

#### **Web-Based Enterprise Management (WBEM)**

Web-Based Enterprise Management (WBEM) is a standard set of management tools and Internet technologies. The Distributed Management Task Force (DMTF) has developed standards for WBEM that include the Common Information Model (CIM) for databases, xml/CIM for coding, and CIM Operations over HTTP for transporting information.

#### **Windows Management Instrumentation (WMI)**

Windows Management Instrumentation (WMI) is a Microsoft infrastructure that supports the CIM model and Microsoft-specific extensions of CIM. It offers query-based information retrieval and event notification.

#### **MOM Entity**

A MOM entity is a MOM Server or a MOM Managed PC.

### **MOM Server**

A MOM Server is a computer that has access to the MOM database.

### **MOM Managed PC**

A MOM Managed PC is a computer with a MOM agent running on it. The MOM agent monitors the computer and reports problems to a MOM Server.

### **MOM Administrator Console**

The MOM Administrator Console is the GUI where MOM is configured. It also provides the central monitoring point in MOM.

## **How MOM Management Works**

Unicenter Management for MOM (MOM Management) discovers MOM servers and managed PCs so that you can see the status of MOM entities in CA NSM WorldView. MOM Management also updates MOM object status and sends MOM alerts to the CA NSM Event Log. The process happens like this:

- WorldView Auto Discovery runs on your MOM network to populate your MDB with WBEM-enabled computers.

**Note:** MOM entities can be classified only when they are already in the Unicenter Management Database.

- MOM Discovery communicates with MOM and classifies entities according to their roles as MOM Servers or MOM Managed PCs. A MOM object is created in Unispace for each role that a computer plays.

To keep MOM information current, run MOM Discovery at regular intervals.

**Note:** The discovery process adds WorldView classes and icons for the objects it creates. After discovery, restart the RMI server (`rmi_server`) so that the icons are displayed correctly.

- MOM Management organizes MOM entities into Business Process Views called MOM Management Views. One view groups entities according to their roles as MOM Servers or MOM Managed PCs. Other views group MOM entities based on their severity.

MOM entities are visible in the left-pane Topology view of the Management Command Center and on the classic 2D map. From there you can view actual MOM alerts, filter out ones you are not interested in, and acknowledge or resolve alerts.

- MOM Management interfaces with MOM, collects unresolved MOM alerts, and sends them to the Event Log. The alerts are converted into the format of Event messages. You can use Event message records and actions to identify important MOM alerts and act upon them.

**Note:** Resolved MOM alerts are not collected.

- MOM Management updates MOM alerts. Use the MOM Management GUI or the momaleralter command in CA NSM to acknowledge an alert, assign someone to fix the situation that caused an alert, and indicate the progress toward resolving the situation.

**Note:** On Windows, the node running the CA NSM integration with MOM must be in the local Administrators group on the node where MOM is running.

- The error that caused the MOM alert is corrected.
- MOM Management notifies MOM that the alert is resolved. Use the MOM Management GUI or the momaleralter command.

## MOM Alerts as Event Messages

Unicenter Management for MOM converts MOM alerts to Event messages. The actual MOM alerts are not displayed verbatim in the Event log and on the Console.

**Note:** You can see the actual MOM alerts from the Topology view of Management Command Center and the classic 2D map.

The following table shows what information from MOM alerts is put into Event messages:

MOM Alert	Event Message
MOM server that generated the alert	Node
Process that gathers MOM alerts	User
MOM Managed PC where the event that caused the MOM alert occurred	Station
CAMM prefix	Message ID
MOM alert description field. The MOM alert URL is appended to the message, if possible.	Message Text
GUID (Globally Unique Identifier)	User Data
MOM alert source field	Category

The following table shows how MOM severity is converted to Event severity:

MOM Alert Severity	Event Message Severity
10 (Success)	S (Success)

<b>MOM Alert Severity</b>	<b>Event Message Severity</b>
20 (Information)	I (Information)
30 (Warning)	W (Warning)
40 (Error)	E (Error)
50 (Critical Error)	F (Fatal)
60 (Security Breach)	E (Error)
70 (Unavailable)	I (Information)

**Note:** Event Management does not communicate directly with MOM.

### Status of MOM Entities in WorldView

MOM alerts contain a resolution state value, which indicates what is being done to resolve the situation that caused the alert. Unicenter Management for MOM uses the resolution states and the severity to determine the status of servers and managed PCs in WorldView. The color of objects indicates whether the status is normal, warning, critical, or down.

The following table shows how the default MOM resolution states are equated to WorldView status:

<b>MOM Resolution State</b>	<b>WorldView Status</b>																
0 (New)	Depends on severity: <table border="1"> <thead> <tr> <th><b>WorldView</b></th> <th><b>MOM</b></th> </tr> </thead> <tbody> <tr> <td>Normal</td> <td>Success</td> </tr> <tr> <td>Normal</td> <td>Information</td> </tr> <tr> <td>Warning</td> <td>Warning</td> </tr> <tr> <td>Critical</td> <td>Error</td> </tr> <tr> <td>Critical</td> <td>Critical Error</td> </tr> <tr> <td>Critical</td> <td>Security Breach</td> </tr> <tr> <td>Down</td> <td>Unavailable</td> </tr> </tbody> </table>	<b>WorldView</b>	<b>MOM</b>	Normal	Success	Normal	Information	Warning	Warning	Critical	Error	Critical	Critical Error	Critical	Security Breach	Down	Unavailable
<b>WorldView</b>	<b>MOM</b>																
Normal	Success																
Normal	Information																
Warning	Warning																
Critical	Error																
Critical	Critical Error																
Critical	Security Breach																
Down	Unavailable																
85 (Acknowledged)	Warning																
170 (Level 1: Assigned to help desk or local support)	Warning																
180 (Level 2: Assigned to subject matter expert)	Warning																

<b>MOM Resolution State</b>	<b>WorldView Status</b>
190 (Level 3: Requires scheduled maintenance)	Warning
200 (Level 4: Assigned to external group or vendor)	Warning
255 (Resolved)	Normal

## Using MOM Management

The online Management Command Center help contains the following procedures for using MOM Management:

- Run MOM Discovery
- View Alerts for a MOM Server
- Filter Alerts in the Unicenter Alert Container
- Add, Remove, or Rearrange Columns in the Unicenter Alert Container
- Refresh the Unicenter Alert Container Automatically
- Acknowledge or Resolve MOM Alerts Using the GUI
- Acknowledge or Resolve MOM Alerts Using the momalertalter Command

## Integration with Microsoft System Center Operations Manager (SCOM)

CA NSM integrates with Microsoft's System Center Operations Manager 2007 (SCOM).

SCOM delivers operations management by providing event management, proactive monitoring and alerting, reporting, health state, and trend analysis. It helps administrators monitor and manage the events and performance of Windows systems.

The integration between SCOM and CA NSM provides a central location for performing management functions. You can use WorldView to see the status of SCOM servers and managed PCs addition and to view and update SCOM alerts. When you update an alert using CA NSM, the change is also reflected on the Microsoft side. You can also view SCOM alerts as CA NSM events on the Event Console.

**Note:** CA NSM provides integration kits to both of Microsoft's management applications, Microsoft Operations Manager (MOM) and System Center Operations Manager 2007 (SCOM). Although the integrations to MOM and SCOM can coexist on the same management server, each one integrates only with its Microsoft counterpart.

## Minimum Software Requirements

Before installing the CA NSM SCOM Integration, your system must meet minimum software requirements.

### Computer where the Integration Is Installed

The computer where the SCOM Integration is installed must have, at minimum, the following software:

- The Microsoft System Center Operations Manager Console Client, which provides the files necessary to integrate with SCOM
- A CA NSM Event Agent
- A CA NSM WorldView Client

### Elsewhere in the Domain

The domain where the SCOM Integration is installed must meet the following minimum requirements:

- An instance of System Center Operations Manager must be installed and running.
- A CA NSM Event Manager must be present because the SCOM Integration creates CA NSM events from SCOM alerts. The Event Manager may be on the same computer as the SCOM Integration.
- A CA NSM WorldView Manager must be present because the SCOM Integration creates objects in the WorldView Repository. The Integration reflects the status of those objects, and SCOM alerts are created based on that status. The WorldView Manager may be on the same computer as the SCOM Integration.
- An instance of the Management Command Center must be present because it is the user interface that shows SCOM objects in the WorldView Topology, and alerts in the SCOM Alert Viewer. The Unicenter MCC may be on the same computer as the SCOM Integration.

## SCOM Terminology

The following terms define industry standards related to Microsoft System Center Operations Manager (SCOM). Also included are definitions that are specific to SCOM.

### **Windows Management Instrumentation (WMI)**

Windows Management Instrumentation (WMI) is a Microsoft infrastructure that supports the CIM model and Microsoft-specific extensions of CIM. It offers query-based information retrieval and event notification.

### **SCOM Entity**

A SCOM entity is any device that SCOM manages regardless of the method used to do it.

### **SCOM Management Server**

A SCOM Management Server is a computer that has access to the SCOM database.

### **SCOM RMS**

A SCOM Root Management Server (RMS) is a computer that runs the SDK service needed for the integration to communicate. It is a SCOM Management Server as well.

### **SCOM Agent Managed PC**

A SCOM Agent Managed PC is a computer with a SCOM agent running on it. The SCOM agent monitors the computer and reports problems to a SCOM Server.

### **SCOM Agentless Managed PC**

A SCOM Agentless Managed PC is a remotely managed computer that has no health service installed.

### **SCOM Gateway Server**

A SCOM Gateway Server is a computer that provides a trust relationship between two managed domains.

### **SCOM Operations Console**

The SCOM Operations Console is the GUI where SCOM is configured. It also provides the central monitoring point in SCOM.

## How the SCOM Integration Works

The CA NSM Integration with System Center Operations Manager discovers SCOM servers and managed PCs so that you can see the status of SCOM entities in CA NSM WorldView. The integration also updates the SCOM health status of these entities in WorldView, and sends SCOM alerts to the CA NSM Event Log. The process happens like this:

- WorldView Auto Discovery runs on your SCOM network to populate your MDB.  
**Note:** You must know the name of and have Ops admin privileges on the Root Management Server (RMS) of your SCOM domain. In addition, the CA NSM Manager must be in the same domain in order for authentication and discovery to take place.
- SCOM Discovery communicates with SCOM and classifies entities according to their roles as SCOM management servers, gateway servers, and agent or agent-less managed PCs. A SCOM object is created for each role that a computer plays.

To keep SCOM information current, run SCOM Discovery at regular intervals.

**Note:** The discovery process adds WorldView classes and icons for the objects it creates. After discovery, restart the RMI server (rmi\_server) so that the icons are displayed correctly.

- The SCOM Integration organizes SCOM entities into Business Process Views called Ops Management Views. One view groups entities according to their roles. Other views group SCOM entities based on their severity.

SCOM entities are visible in the left-pane Topology view of the Management Command Center. From there you can view actual SCOM alerts, filter out ones you are not interested in, and acknowledge or resolve alerts.

- The SCOM Integration interfaces with Operations Manager, collects SCOM alerts, and sends them to the Event Log. The alerts are converted into the format of Event messages. You can use Event message records and actions to identify important SCOM alerts and act upon them.
- The SCOM Integration updates SCOM alerts. Use the SCOM GUI in the Unicenter MCC to acknowledge an alert, assign someone to fix the situation that caused an alert, and indicate the progress toward resolving the situation.
- When an alert is updated through the Unicenter MCC, it is also updated in Operations Manager.

**Note:** In order for the update on the Microsoft side to be successful, the user logging into SCOM through the alert viewer must be in the SCOM Administrators group on the node where Operations Manager is running.



## SCOM Alerts as Event Messages

The CA NSM SCOM Integration converts SCOM alerts to Event messages. The actual SCOM alerts are not displayed verbatim on the Console.

**Note:** You can see the actual SCOM alerts in the SCOM Alerts Viewer of the Management Command Center.

The following table shows what information from SCOM alerts is put into Event messages:

SCOM Alert	Event Message
SCOM server that generated the alert	Node
Process that gathers SCOM alerts	User
SCOM object where the event that caused the alert occurred	Station
CAOPS prefix	Message ID
SCOM alert description field	Message Text
GUID (Globally Unique Identifier)	User Data
SCOM object display name	Category

The following table shows how SCOM alert severity is converted to Event severity:

SCOM Alert Severity	Event Message Severity
Success	S (Success)
Information	I (Information)
Warning	W (Warning)
Uninitialized	U (Unknown)
Error	F (Fatal)

**Note:** Event Management does not communicate directly with SCOM.

## Status of SCOM Entities in WorldView

SCOM contains a health state value for monitored objects that indicates the object's status as defined by Operations Manager policies. The CA NSM SCOM integration synchronizes the health state of computers it discovered during SCOM discovery with the WorldView status of those objects in the repository.

The following table shows how Operation's manager health states map to WorldView status:

<b>SCOM Health State</b>	<b>WorldView Status</b>
Success	Normal
Warning	Warning
Error	Critical
Uninitialized	Unknown

## SCOMMsgconfig Utility

The SCOMMsgconfig utility lets you select the SCOM alert fields to be included in the corresponding CA NSM EM message from the CA Unicenter System Center Operation Manager integration. You can select any SCOM alert field in the EM message and in any order.

## Configure SCOMMsgconfig Utility

By using the SCOMMsgconfig utility, you can select any SCOM alert field that you want in the EM message.

To configure SCOM utility, do the following:

1. Go to the CCS\WVEM\OpsMgr folder
2. Run the SCOMMsgconfig.exe command from the command prompt.  
The Message setup for NSM SCOM Integration window opens.
3. From the SCOM Alert Fields drop-down list, select the fields to be included inside the message to EM. You can select multiple fields. Selected fields are sent to EM in the order selected.
4. Click Add to add the SCOM alerts to the "Fields to be included inside EM message" text box.
5. Select the OpsSeverName check box to send the SCOM server name inside the message.
6. Select the Enable check box to enable all SCOM configuration you have set up in the above steps.  
**Important!** You must select the Enable check box to enable all SCOM configuration otherwise the changes will only be saved and not be enabled.
7. Click OK and restart CA Unicenter SCOM service to enable the new SCOM configuration.

In CA NSM r11.2 or CA NSM r11.2 SP1, the Node field contained SCOM server machine and the Workstation field contained the SCOM alert originator name. However, in CA NSM r11.2 SP2, the Node and the Workstation field contains the SCOM alert originator name.



# Appendix G: Scanning the Systems for Viruses

---

This section contains the following topics:

[Virus Scan](#) (see page 469)

[Downloading Virus Signature Updates](#) (see page 469)

[Deleting Old Scan Logs](#) (see page 470)

## Virus Scan

The Event Management suite of functions includes Unicenter Virus Scan. This utility, available only on Windows, automatically scans the local drive daily at midnight. Parameters can be set to customize the type of scan, the drive to be scanned, and the actions to be taken upon virus detection.

CA NSM provides predefined message policies that automatically launch the Virus Scan utility at midnight. The installation procedure defines a Midnight\* message record and associated action to run Virus Scan every day at midnight.

**Note:** To prevent Virus Scan from running automatically, simply delete the Midnight Virus Scan message record (Message ID: Midnight\*) from the Message Records container. You can also delete the message record's only associated message action from the Message Action Summary container. Deleting either record prevents the automatic execution of Virus Scan.

You can also run Virus Scan on demand by launching the `inocmd32.exe` command with your desired parameters. Parameters can be used to specify the type of scan performed and the action to be taken upon virus detection. Upon virus detection, Virus Scan sends messages to the CA NSM Console Log and the Event Log.

For more information on the `inocmd32.exe` command, including a list of valid parameters, see the online *CA Reference*.

## Downloading Virus Signature Updates

See the instructions for updating your virus signature in the online *CA Procedures* under the topic, *Downloading Virus Signature Updates*. The procedure includes a link to the appropriate CA support web page.

## Deleting Old Scan Logs

Virus Scan maintains its Scan Logs indefinitely in the \NSM\db directory on the hard drive. If you need to reclaim hard disk space, you can manually delete the old logs.

# Appendix H: Using Ports to Transfer Data

---

This section contains the following topics:

[Utilizing and Configuring Ports](#) (see page 471)

[Required Open Ports](#) (see page 472)

[Optional Ports](#) (see page 473)

[Configure the DIA Communications Port](#) (see page 474)

[CA Message Queuing Service \(CAM\)](#) (see page 476)

## Utilizing and Configuring Ports

CA NSM employs the following primary data transport mechanisms to transfer data between CA NSM enabled nodes:

- Distributed Intelligent Architecture (DIA)
- CA Message Queuing Service (CAM)

Other legacy communications mechanisms may also be supported by CA NSM for the purposes of backwards compatibility.

To support these communication mechanisms, certain ports in a firewall must be open.

**Note:** CA is committed to reducing the number of ports that are required to use CA NSM. For this reason, we have identified the ports that are required and the ports that are optional. CA Technology Services can help you design the best port solution for your enterprise.

## Required Open Ports

This table provides a list of the ports that are required to be open in a firewall to support basic communication among CA NSM components.

<b>Component</b>	<b>Default Port</b>	<b>Port Type</b>	<b>Install checks if default port is in use?</b>	<b>Install prompts for different port if conflict exists?</b>	<b>Comments</b>
CA Common Communications Interface (CAICCI)	1721	TCP	Y	Y	CA IANA registered CAICCI listener port (CCI Remote Service).
CAM	4104	UDP	Y	Y	Used for System Performance, Continuous Discovery.
DIA	5635	TCP	N	N	Used for MCC to Manager communication
DIA	5636	TCP	N	N	Used for MCC to Manager communication
CA Common Communications Interface (CAICCI)	7001	TCP	Y	Y	Required for NSM communications that use CCI between Windows servers.
Apache Tomcat	9090	TCP	Y	N	Services incoming requests by component applications that expose functionality through Tomcat, such as Unicenter Web Reporting Server.
Apache Tomcat	9005	TCP	Y	N	Used when Tomcat waits for a shutdown command.
DIA	11501	TCP	N	N	Only requires firewall to open outbound
DIA	11502	TCP	N	N	Only requires firewall to open outbound
DIA	11503	TCP	N	N	Only requires firewall to open outbound
DIA	11504	TCP	N	N	Used by DIA for manager to manager communication
Ingres Remote Client	19016	TCP	Y	Y	Used to communicate with MDB server if you are using an Ingres



Component	Default Port	Port Type	Install checks if default port is in use?	Install prompts for different port if conflict exists?	Comments
					database. Port number is bound to Ingres instance name. Default code is EI, but can be changed at installation.

## Optional Ports

This table provides a list of the optional ports that are required to be open in a firewall to support only certain low-level features or compatibility with a previous version. These ports are grouped by their specific component in the Ports by Component section.

Component	Default Port	Port Type	Install checks if default port is in use?	Install prompts for different port if conflict exists?	Comments
Non-CA agents, Enterprise Management	162	UDP	N	N	Native SNMP traps sent to DSM policy. Required to receive traps from non-CA agents. Additionally, Enterprise Management <i>requires</i> this port if Trapmux is used to support SNMP V3.
CAM	4105	TCP	Y	Y	Used for System Performance, Continuous Discovery, CAM. This port needs to be opened in a firewall only if the CAM communications method is set to TCP. Default mode of operation among computers is UDP.
Enterprise Management	6161	UDP	N	Y	When trapmux is active, used as the catrapd command port.
Enterprise Management	6163	UDP	N	Y	If Trapmux is being used (SNMP v3 support is activated), catrapd opens this port on which to listen
Agent Technology	6665	UDP	N	N	Used only if DIA is not installed

Component	Default Port	Port Type	Install checks if default port is in use?	Install prompts for different port if conflict exists?	Comments
Mobile Services	8888	TCP	N	N	Used for communications between a CA NSM manager and Pocket PC devices. Disabled by default. If Pocket PC connectivity is required, Mobile Services must be activated and this port opened.
CA Common Communications Interface (CAICCI)	7000	TCP	N	N	Default for AP/OPS interface component. Configurable at site. In field since 2000. Is not required if all nodes are at r11 or higher.
Agent Technology	7774	TCP	N	N	Required when using Agent Technology processes using the -@ option.

## Configure the DIA Communications Port

DIA requires the use of four distinctive ports for communication. By default, these ports are 11501, 11502, 11503 and 11504. You can change them in DIA configuration files.

**Note:** You must complete the following procedure on each machine where DIA is installed.

### To change the DIA ports

1. Open the dna.cfg file in the folder *InstallPath\CA\SharedComponents\CCS\DIA\dia\dna\config* using a text editor of your choice.
2. Set RMI\_REGISTRY\_PORT and RMI\_DATA\_PORT to the port numbers you want.  
The two port numbers must be different.
3. Save the dna.cfg file.
4. Open the ukb.cfg file in the folder *InstallPath\CA\SharedComponents\CCS\DIA\dia\ukb\config* using a text editor of your choice.
5. Set RMI\_DNA\_REGISTRY\_PORT to the same port number that you specified in RMI\_REGISTRY\_PORT in dna.cfg at Step 2.

6. Set RMI\_REGISTRY\_PORT and RMI\_DATA\_PORT to the port numbers you want.

The two port numbers must be different from each other and should be different from those you set for RMI\_REGISTRY\_PORT and RMI\_DATA\_PORT in the dna.cfg file.

7. Save the ukb.cfg file.
8. Modify the SRV record and set the port number field to the same value as you have for RMI\_REGISTRY\_PORT in the ukb.cfg file.  
See the topic Configure Unicenter Domain Name Services earlier in this appendix. If you do not have an SRV record in the domain, skip this step.
9. Stop both of the following services or daemons and restart them to apply the changes:
  - CA DIA 1.3 DNA
  - CA DIA 1.3 Knowledge Base

The following example assumes you want to set customized ports to 16001, 16002, 16003, and 16004:

1. In the dna.cfg file, set the following:
  - RMI\_REGISTRY\_PORT = 16001
  - RMI\_DATA\_PORT = 16002
2. In the ukb.cfg file, set the following
  - RMI\_DNA\_REGISTRY\_PORT = 16001
  - RMI\_REGISTRY\_PORT = 16003
  - RMI\_DATA\_PORT = 16004
3. In the SRV record of DNS, make the following change:  
Set "port number field" to 16003.

## CA Message Queuing Service (CAM)

The CA Message Queuing Service (CAM) is one of the two principle data transport mechanisms used by CA NSM. CAM provides connection-less application-to-application messaging with reliable delivery. In addition to the standard IP protocols, CAM provides the following features:

- Store and forward
- Protocol independence
- Routing
- Auditing
- Delivery/non-delivery notification
- User-defined quality of service
- Unified name/address mapping
- Scalability and performance
- Low resource server/network utilization
- Application and server "ping" availability detection
- Support for very large messages
- Comprehensive diagnostics and statistics

CAM combines the lightweight benefits of UDP with the reliable delivery of TCP. A CAM server process runs on each host supporting CAM. CA's applications that use CAM communicate with the CAM local server, which then forwards messages to other CAM servers or to other CAM client applications on the same computer. For more information about CAM, see the CAM product documentation.

CAFT is a simple file transfer protocol (similar to FTP) that uses CAM for its data transport.

### Supported Transport Layer Protocols

CAM supports the following Transport Layer Protocols:

- UDP on port 4104--This is the default for intra-host communications. The port number is configurable. TCP (4105) can optionally be configured instead of UDP for increased performance on high-loss/unreliable networks.
- TCP on port 4105--This is the default for inter-host (application to CAM server) communications.
- SPX on port 4905

## Components That Use CAM/CAFT

The following CA NSM components make use of CAM and in certain cases, CAFT, as their principle messaging/data transport mechanism.

<b>Component</b>	<b>Subcomponent</b>	<b>Windows Executable</b>	<b>UNIX/Linux Executable</b>
Continuous Discovery	Continuous Discovery Manager	dscvmgrservice.exe	CaDiscMgrService
	Continuous Discovery Agent	dscvagtsservice.exe	CaDiscAgentService
Systems Performance	Performance Scope	perfscope.exe	N/A
	Performance Trend	perftrend.exe	N/A
	Performance Web Reporting	java.exe	java
	Performance Configuration	egc30n.exe	N/A
		discover.exe	N/A
		hpaprofile.exe	N/A
	Performance Configuration (one-click)	capmwvc.exe	N/A
	Performance Data Grid	pdectl.exe	pdectl
		pdgstat.exe	pdgstat
		capmpde.exe	capmpde
		pdesumgen.exe	pdesumgen
	Performance Domain Server	configserver.exe	configserver
	Performance Distribution Server	profileserver.exe	profileserver
	Performance Agent	prfagent.exe	prfAgent
		hpaagent.exe	hpaAgent
hpacbman.exe		hpacbman	
hpacbcol.exe		hpacbcol	
Performance Utilities	cubespider.exe	cubespider	
	rtpmon.exe	rtpmon	
	cfgutil.exe	cfgutil	
	pdtodm_m.exe	pdtodm_m	
	pdtodb_u.exe	pdtodb_u	

Component	Subcomponent	Windows Executable	UNIX/Linux Executable
		pdtoxml.exe	pdtoxml

## CAM/CAFT Configuration Files

The configuration files for CAM and CAFT are as follows:

### **cam.cfg**

Main CAM configuration file. Configures most aspects of CAM's operation, including basic configuration settings, routing rules, logging and tracing control, port and transport layer control, and so forth.

### **camclient.cfg**

Contains a list of registered processes that can be automatically started by CAM when a message arrives for them.

### **caftenv.cfg**

CAFT environment variables (for example, used to store locations set at install time of Systems Performance Cube Store directories).

## CAM/CAFT Binaries

The following list of CAM binaries includes the principle CAM components, as well as utilities and configuration tools.

Windows	UNIX/Linux	Description
cam.exe	cam	CAM server
camabort.exe	camabort	Stops the CAM server (forcefully)
camben.exe	camben	Benchmarks a communications link
camclose.exe	camclose	Stops the CAM server cleanly (informs clients first).
camconfig.exe	camconfig	Changes the CAM configuration and routing.
camping.exe	camping	Similar to ICMP echo request (ping), but can check availability of client applications as well as hosts.
camq.exe	camq	Lists and manipulates queues in the CAM server.
camsave.exe	camsave	Saves the CAM server's configuration in the same format as cam.cfg.
camstat.exe	camstat	Displays detailed status information for a CAM server.
camswitch.exe	camswitch	Forces a log file switch.

## How to Encrypt the MCC Data Transport (CAM) for AIS Providers

For security reasons, you may want to encrypt the information going over the network between the MCC and the AIS Providers (WorldView, DSM, etc.). These providers use the AIS subsystem, which in turn uses CAM. Complete the following process to encrypt CAM for AIS providers:

1. Install the CA Secure Socket Adapter (SSA).
2. Reconfigure CAM to use the newly installed SSA component.

### Install the CA Secure Socket Adapter

You must install the CA Secure Socket Adapter (SSA) so that you can use it to encrypt CAM.

#### To install the CA Secure Socket Adapter

1. Navigate to the following location on the CA NSM installation DVD:  
Windows\NT\%LANGPACK%\%ALL%\SSA
2. Run the following executable:  
`CASockAdapterSetupWin32.exe`
3. Follow the prompts to install SSA.

**Note:** SSA 2.0 is not currently supported on Solaris on Intel. All other manager/client configurations are supported.

### Integrating with the Secure Socket Adapter

The Secure Socket Adapter provides a means of adding extra services to TCP connections without requiring changes to the programs that make these connections. This description applies to SSA 2.0 or later but is also applicable to the earlier version of SSA, known as the Dylan Socket Adapter. For the Dylan Socket Adapter, the `csamconfigedit` command described below is known as `configedit`.

CAM makes use of SSA to enhance the security of its messaging. When using SSA, TCP paths can be secured using SSA's ability to provide SSL/TLS encryption.

**Note:** Communications within a machine are not normally encrypted, and the CAM API library does not integrate with SSA.

On UNIX platforms, the CAM server detects SSA's presence at startup and makes use of the SSA library to interface with the underlying communications layer. On Windows, we provide a version of the CAM server code that has been adapted to use SSA. In both cases, the adapted CAM performs no differently unless SSA is configured to adapt the port that CAM uses for TCP.

## Configure CAM for SSA

SSA adapts TCP connections, but does not adapt UDP messages. To account for the lack of UDP support, if all CAM communications are to use SSL/TLS, you can disable UDP, making TCP the default protocol. To disable UDP for CAM, include the following line in the CAM configuration file, `cam.cfg`:

```
*CONFIG
udp_port=0
```

Alternatively, you can configure specific paths that you want to adapt by defining them in the `*PATHS` section. However, this approach could be cumbersome on large networks and difficult to maintain when machine addresses are determined by DHCP.

A more usable option is to select the machines on which you always want to use SSL for CAM communications and configure them as follows in `cam.cfg`:

```
*CONFIG
udp_port=-1
```

With this setting, all remote paths created by CAM are TCP paths (and can then be adapted, using SSA to use SSL). Also, if other machines attempt to establish a UDP path, they are rejected and switched to TCP. However, the part of the network where security is not required can continue to use UDP.

**Note:** With this configuration, one unencrypted UDP message is sent and rejected for connections that are switched to TCP.



## Activate SSA-enabled CAM on Windows

On Windows, you must activate the SSA-enabled version of CAM.

### To activate SSA-enabled CAM on Windows

1. Open a command prompt of a user with administrator privileges.

**Note:** On Windows Vista and later Windows versions, you must explicitly request administrator privileges when opening the command prompt, even if you are an administrator.

2. Enter the following command:

```
camdsa install
```

The command associates the SSA-enabled version of the CAM server with the CAM service.

3. Restart the CAM service.

CAM is SSA-enabled.

You can reverse this process by running the following command and restarting the CAM service:

```
cam install
```

## Configure SSA to Enable CAM to use SSL/TLS

You must configure the SSA-enabled CAM to use SSL/TLS by requesting the support and enabling use of the SSA connection broker.

To configure SSA-enabled CAM to use SSL/TLS, run the following command at a command prompt:

```
csamconfigedit port=4105 EnableSSL=True EnablePmux=True PmuxLegacyPortListen=True  
PmuxLegacyPortBindAddress=127.0.0.1
```

This command requests SSL/TLS encryption on the CAM TCP port and enables use of the SSA connection broker (port multiplexer) for connections using that port. The port multiplexer must be used to allow support for non-encrypted TCP connections, as it enables the SSA software to differentiate between the two. Legacy connections are also allowed on the port but are restricted to within-machine connections by binding to 127.0.0.1, the IPv4 localhost address.

In an IPv6 environment, you may need to replace the final parameter value of 127.0.0.1 with localhost (or 127.0.0.1:::1 if localhost cannot be resolved to one or both of these addresses).

**Note:** On some machines, localhost may not resolve to any addresses.

On AIX, in rare circumstances, CAM may not be able to accept local connection when using CAM 1.12 or later if you set a bind address. If you experience this issue (one symptom is that the camf process is running but utilities such as camstat claim that it is not), remove the PmuxLegacyPortBindAddress parameter from the initial definition or set it back to its default value.

If you want to accept unadapted connections from remote machines, you can omit the PmuxLegacyPortBindAddress parameter, but you also have to define an appropriate OutboundHostList to ensure that outward connections to these machines are not adapted. This operation may prove complex in practice, and the most viable policy is to encrypt all connections. You could use UDP for non-encrypted connections, but this would require you to explicitly configure (in CAM) all encrypted connections. SSA 2.1 will improve flexibility in this area.

# Appendix I: Integrating with CA Spectrum Service Assurance

---

This section contains the following topics:

[CA NSM Connector Import](#) (see page 483)

## CA NSM Connector Import

The following information provides details about the CA NSM connector and what it imports. The connector retrieves WorldView and DSM objects.

### WorldView Import

#### Automatic CI and service synchronization

Provides this support.

#### Object updates

Receives WorldView notifications for add, update, and delete actions of managed objects and inclusions.

#### Events and Alarms

Receives WorldView notification of status updates. Status text specifies the reason for an alarm.

#### Object types and classes

- BPV—All objects of the Worldview class BusinessView and its subclasses are imported as services.
- Managed Objects—All WorldView managed objects are imported as CIs.
- DSM granular objects—All DSM granular objects used in WorldView BPVs are imported if you use the SAMP functionality in CA NSM to store these objects in the WorldView repository.

#### Alarm types

All WorldView status updates are sent as state change alarms.

WorldView has no true alarms, only object severity and status text properties. The connector creates an alarm for any object with non-normal severity. Because each object can have only one severity, only one alarm per object can exist at any time. If a state changes, a new alarm replaces the previous one, or clears the alarm if the state changes to Normal.

## DSM Import

### Connection to Silo

Uses the DSM ObjectFactory to connect to ORB.

### Automatic CI and service synchronization

Provides this support for DSM objects. Services are not imported from DSM.

### Object updates

Registers for DSM callbacks for additions, deletions, and updates to DSM objects.

### Events and Alarms

The Object Factory calls the connector's event handler. Text provided as *reason* in the ObjectFactory callback is imported as alarm text.

Starting with CA NSM r11.2 SP2, the DSM connector imports more detailed information for each event.

### Object types and classes

DSM objects of following types are imported as CIs. All standard classes for these objects are supported.

- Host level objects
- Agent objects
- Leaf node objects (metrics)

**Note:** No DSM objects are imported as services.

### Alarm types

All standard DSM alarm types are supported.

**Note:** The metrics and resources being monitored by DSM agents that appear in the WorldView topology under the agent objects (DSM granular objects) do not exist in the WorldView repository by default. The WorldView connector only imports objects that exist in the WorldView repository, so these DSM granular objects are not automatically included in any imported BPV service definition. However, you can insert the DSM granular objects into the WorldView repository using SAMP. With this feature enabled, the WorldView connector can automatically import services containing DSM objects.

# Appendix J: Integrating with CA Spectrum

---

This section contains the following topics:

[CA Spectrum-NSM Integration Kit](#) (see page 485)

[CA Spectrum Infrastructure Manager and CA NSM Integration Guide](#) (see page 485)

## CA Spectrum-NSM Integration Kit

The CA Spectrum integration with CA NSMs is achieved using the CA Spectrum-NSM Integration Kit. After you install the CA Spectrum-NSM Integration Kit, you can do the following from CA NSM administrative interfaces:

- Monitor states of CA Spectrum objects
- View and manage CA Spectrum device model alarms
- View CA Spectrum device model alarms as events in the Event Console
- Access the CA Spectrum OneClick Console from the MCC or 2D Map

The CA Spectrum-NSM Integration Kit is included on the product media for both the CA NSM and CA Spectrum applications.

**Note:** For more information about NSM Integration Kit, see the *CA Spectrum Infrastructure Manager and CA NSM Integration Guide*.

## CA Spectrum Infrastructure Manager and CA NSM Integration Guide

The information that was previously in this appendix is now in the *CA Spectrum Infrastructure Manager and CA NSM Integration Guide*. This guide is distributed with both CA NSM and CA Spectrum. See the *Integration Guide* for detailed information about the integration between CA NSM and CA Spectrum.



# Appendix K: Integrating with CA Virtual Performance Management 11.7 VC AIM

---

This section contains the following topics:

[Introduction to CA Virtual Performance Management](#) (see page 487)

[CA SystemEDGE Agent](#) (see page 488)

[Integration with CA Virtual Performance Management](#) (see page 490)

[Discover VPM Resources](#) (see page 491)

[Enable AIMS in VPM integration](#) (see page 496)

## Introduction to CA Virtual Performance Management

CA Virtual Performance Management (CA VPM) is a policy-based product that automatically monitors, reconfigures, and provisions physical and virtual resources to dynamically meet the load demands of complex service-oriented data centers. CA Virtual Performance Management is built on a Service Oriented Architecture (SOA) and continuously analyzes your data center to help ensure that your servers are optimally provisioned to perform required tasks. You can manage your data center and obtain detailed information about each managed computer using the web-based CA VPM user interface.

The adoption of virtual monitoring technologies has dramatically changed the Enterprise IT Management (EITM) marketplace. Virtualization is no longer viewed as a specialized, niche solution. As a result, customers are expecting EITM products to provide out-of-the-box support for their virtualized environments. CA Virtual Performance Management represents the evolution of the Unicenter ASM product from an add-on, optional product specific to CA NSM into a common component and stand-alone product that enables other CA EITM products to support virtualization environments.

The integration with CA NSM provides associations between managed objects in the CA NSM database and the CA Virtual Performance Management AOM database.

## CA SystemEDGE Agent

CA Virtual Performance Management uses the CA SystemEDGE agent and its application insight module (AIM) plugins for monitoring and managing a broad range of physical systems and virtual resources. CA NSM lets you manage the CA SystemEDGE agent and the AIMs that are distributed with CA Virtual Performance Management. For more information, see the guide *Inside Systems Monitoring* in the CA NSM documentation set and the CA Virtual Performance Management documentation.

Some CA VPM features are not available with version 4.3 of the CA SystemEDGE agent, and with version 5.0 running in legacy mode. Use version 5.0 of the agent in its regular operating mode to enable full CA VPM functionality.

**Note:** For more information about the CA SystemEDGE agent, see the *CA SystemEDGE User Guide*.

### Logical Partition (LPAR) AIM

Logical partitions (LPARs) divide the resources of a server into subsets and make a server run as if it were multiple independent servers. You can install software on an LPAR, which runs as an independent server with the resources allocated to it.

The CA IBM LPAR AIM is a plug-in to the CA SystemEDGE agent that lets you manage the infrastructure of your P5 and P6 environment. When you integrate the LPAR AIM with CA NSM, you can discover the LPAR AIM, view its monitored data, and configure how it monitors.

### Service Response Monitor (SRM) AIM

CA Service Response Monitor (CA SRM; previously CA eHealth Service Availability) is an AIM that provides response time monitoring for common network services. You can create response time tests for services such as DNS, HTTP, SNMP, and others to ensure that your network services are available and responding within defined thresholds.

CA SRM supports an object state model by using the CA SystemEDGE agent object model to enable state monitoring of objects represented by a response time test. You enter instance name and class parameters for each test, and the CA SystemEDGE agent uses these and other defined values to create an entry in the Self Monitor table that can track the severity of the object represented by the test (a URL, a server, and so on).



The CA SRM AIM Agent View (abrowser) lets you view AIM summary information and view, create, and manage response time tests. The Agent View supports CA SRM r3.0, which must be running under CA SystemEDGE r5.0. The CA SRM cannot exist or run without CA SystemEDGE.

**Note:** For more information about how to use the CA SRM AIM, see the *Service Response Monitor User Guide* and SRM AIM Agent View Help.

## VMware vCenter (VC) AIM

VMware vCenter Server runs as a service in Microsoft Windows operating systems and provides the central point of control for configuring, provisioning, and managing virtualized IT environments. The CA SystemEDGE AIM for VMware vCenter Server runs on this VMware vCenter Server computer and communicates with VMware vCenter Server through web-services and with the vCenter Server PMM on the Manager.

## Xen AIM

XenServer is a server virtualization platform that offers near bare-metal virtualization performance for virtualized server and client operating systems. XenServer uses the Xen Hypervisor to virtualize each server on which it is installed, enabling each to host multiple virtual machines simultaneously with guaranteed performance. XenServer allows you to combine multiple Xen-enabled servers into a powerful Resource Pool, using industry-standard shared storage architectures and leveraging resource clustering technology created by XenSource. In doing so, XenServer extends the basic single-server notion of virtualization to enable seamless virtualization of multiple servers as a Resource Pool, whose storage, memory, CPU, and networking resources can be dynamically controlled to deliver optimal performance, increased resiliency and availability, and maximum utilization of data center resources.

XEN AIM agent is developed to monitor and configure the resources of the Citrix Xen Server machine. The Xen AIM will reside on a Windows machine and gather information from Xen hosts using RPC-XML protocol.

## Zones AIM

A Sun Solaris Zone defines a virtualized operating system platform (called a zone) that provides an isolated, secure environment in which to run applications. This allows allocation of resources among applications and services, and helps ensure that processes do not affect other zones. Solaris manages each zone as one entity. A *container* is a zone that also uses the operating system's resource management. The Solaris Zones PMM provides health monitoring, management, and provisioning of Solaris Zones environments.

The Zone application insight module (AIM) is a plugin to the CA SystemEDGE agent that lets you manage the infrastructure of your Sun Solaris systems environment. When you integrate the Sun Solaris Zone AIM with CA NSM, you can discover the Sun Solaris Zone AIM, view its monitored data, and configure how it monitors.

**Note:** For more information about enabling and configuring the Sun Solaris Zone AIM in CA Virtual Performance Management (CA VPM), see the CA VPM documentation.

## Integration with CA Virtual Performance Management

We can monitor and configure VPM (Sun Solaris Zones, Citrix XenServer, VMware vCenter Server, and IBM LPAR) AIM agent from NSM manager which can reside on same or remote windows server using Abrowser and DSM policy.

VPM Abrowser support is based on the CA NSM existing framework and Unicenter TNG Technology. Agent Browser utility is a generic application that is used to build GUI applications to display the contents of SNMP-compliant agents. This utility reads simple text files to build a graphical interface that can be tailored to the individual needs of a particular agent.

VPM DSM policy and WVC is based on the CA NSM DSM framework and responsible to perform platform level discovery and WV object state management according to polling and traps received from AIM.

VPM Performance sponsor module is responsible to collect data from VPM proxy service component.

## Discover VPM Resources

The CA NSM discovery process creates and classifies new objects in the Management Database (MDB). After VPM objects (Sun Solaris Zones, Citrix XenServer, VMware vCenter Server, and IBM LPAR) are discovered, the VPM Integration for CA NSM performs a second-level discovery process to find any VPM objects present on the nodes. Your agents must be installed and running properly for the discovery process to succeed. The resulting objects are then linked to the appropriate Business Process Views in WorldView.

To discover VPM resources, enter the following command at the command prompt:

```
dscvrbe -7 hostname -v 9
```

### **hostname**

Specifies the host name of the VPM server that the VPM AIM is managing.

Objects discovered as a part of the VPM Integration with CA NSM are represented by the Business Process View icon labeled VPM. When you drill down, icons for Solaris Zones, Citrix XenServer, VMware vCenter, and IBM LPAR environments appear, depending on the type of environments that are discovered.

**Note:** For more information about the CA NSM discovery process, see the *Administration Guide*. For more information about discovery command options, see the *CA Reference Guide*. Both documents are available with the CA NSM documentation set.

## IBM LPAR Object Discovered

The discovery process generates a comprehensive view of the IBM LPAR topology that you can view in the MCC. The following IBM LPAR objects are discovered:

- LPAR servers and LPARs. You can set the states of these objects and perform other management operations.
- Slots, profiles, and virtual disks, virtual Ethernet adaptors, virtual SCSI drives, and virtual serial drives.

## Start the LPAR AIM Agent View

The LPAR AIM Agent View lets you configure how the LPAR AIM manages all IBM LPAR resources, such as CPU and memory settings for LPAR servers, LPARs, and so on.

### To start the LPAR AIM Agent View from the MCC

1. Open the Management Command Center.
2. Navigate to the IBM LPAR server in the Topology view.
3. Find the IBM LPAR Aim Agent object in the topology, right-click the object, and select Actions, View Agent.

The LPAR AIM Agent View appears.

To start the LPAR AIM Agent View from the command line, open a command prompt and enter the following command:

```
abrowser -c browser.calparaim -h lparaimhost
```

### **lparaimhost**

Specifies the host name of the server on which the LPAR AIM is installed.

## Sun Zones Objects Discovered

The discovery process generates a comprehensive view of the Solaris Zones topology that you can view in the MCC. The product discovers the following Solaris Zones objects:

- Zone servers and Zones. You can set the states of these objects and perform other management operations.
- Projects, Processor Sets, Host Disks, Physical NIC, and Virtual NIC.

## Start the Zones AIM Agent View

The Zones AIM Agent View lets you configure how the Zones AIM manages all Solaris Zones resources, such as CPU and memory settings for Zones servers, Zones, and so on.

### To start the Zones AIM Agent View

1. Open the Management Command Center.
2. Navigate to the Zones server in the Topology view.
3. Find the Zones Aim Agent object in the topology, right-click the object, and select Actions, View Agent.

To start the Zones AIM Agent View from the command line, open a command prompt and enter the following command:

```
abrowser -c browser.cazoneaim -h zoneaimhost
```

### **zoneaimhost**

Specifies the host name of the server on which the Zones AIM is installed.

## Citrix XenServer Objects Discovered

Objects discovered as a part of the Citrix XenServer Integration with CA NSM are represented by the Business Process View icon labeled VPM. When you drill down, icons for XenServer environments appear, depending on the type of environments that are discovered.

CA NSM discovers the following XenServer resources:

- Physical CPU, memory, and network interfaces of a XenServer
- Virtual CPU, memory, and network interfaces of VMs
- Resource Pools
- Storage Repositories
- Virtual Disks
- Templates

### Notes:

- If you use Citrix XenServer resource pools, CA NSM can only discover the pool master. Since a XenServer resource pool is represented by the pool master only, the other pool members are not visible to the network. For more information about resource pools, see the Citrix XenServer documentation.
- For more information about the CA NSM discovery process, see the *Administration Guide*. For more information about discovery command options, see the *CA Reference Guide*. Both documents are available with the CA NSM documentation set. For more information about Citrix XenServer Management, see the *CA Virtual Performance Management Implementation Guide*.

## Start the Citrix XenServer AIM View

The CA Citrix XenServer AIM View lets you configure how the XenServer AIM manages all Citrix XenServer resources, such as CPU and memory settings for VMs, XenServers, resource pools, and so on.

### To start the Citrix XenServer AIM View from the MCC

1. Open the Management Command Center.
2. Navigate to the Citrix XenServer in the Topology view.
3. Find the Citrix XenServer AIM object in the topology, right-click the object, and select Actions, View Agent.

The Citrix XenServer AIM View appears.

To start the Citrix XenServer AIM View from the command line, open a command prompt and enter the following command:

```
abrowser -c browser.cacxenaïm -h xenaïmhost
```

#### **xenaïmhost**

Specifies the host name of the server on which the XenServer AIM is installed.

## VMware Objects Discovered

The discovery process generates a comprehensive view of the VMware vCenter Server topology that you can view in the MCC. The following VMware vCenter Server objects are discovered:

- vCenter servers, ESX servers, resource pools, datacenters, clusters, VM instances, and other virtual entities. You can set the states of these objects and perform other management operations. When these objects are classified, the AIM correlates the relationships between the ESX servers and VM instances.
- Disks and networks for all discovered VM instances. You can set the states of these objects and configure the VC AIM to monitor or ignore these devices.

## Start the VC AIM Agent View

The VC AIM Agent View lets you configure how the VC AIM manages all vCenter Server resources, such as CPU and memory settings for VMs, ESX servers, resource pools, and so on.

### To start the VC AIM Agent View from the MCC

1. Open the Management Command Center.
2. Navigate to the vCenter Server in the Topology view.
3. Find the VMware Virtual Center Aim Agent object in the topology, right-click the object, and select Actions, View Agent.

The VC AIM Agent View appears.

To start the VC AIM Agent View from the command line, open a command prompt and enter the following command:

```
abrowser -c browser.cavmvcaim -h vcaimhost
```

#### **vcaimhost**

Specifies the host name of the server on which the VC AIM is installed.

## Enable AIMs in VPM integration

By default all AIMs in the VPM integration are disabled. If you want to enable them, then perform the following steps:

1. Open MCC.  
MCC window opens.
2. Navigate to Tools, DSM Configuration, DSM Agent Class Scope and select the AIMs you want to enable.  
Selected AIMs are enabled.

**Note:** You can enable AIMs by using dsmwiz also.



# Appendix L: Job Management Option

---

This section contains the following topics:

- [How CA NSM Job Management Option Works](#) (see page 497)
- [How to Specify Where to Perform Work](#) (see page 500)
- [How to Identify Resource Requirements for Workload Balancing](#) (see page 501)
- [How to Schedule Work by Dates](#) (see page 502)
- [How to Form Groups of Related Tasks \(Jobsets\)](#) (see page 504)
- [How to Identify Work to Perform](#) (see page 507)
- [How to Schedule Work by Special Events](#) (see page 513)
- [How to Test Your CA NSM JM Option Policy Definitions](#) (see page 517)
- [Autoscan](#) (see page 518)
- [Workload Processing](#) (see page 520)
- [Maintenance Considerations](#) (see page 521)
- [Agent/Server Configurations](#) (see page 524)
- [Cross-Platform Scheduling](#) (see page 526)
- [Windows Configuration Environment Variables](#) (see page 530)
- [UNIX/Linux Configuration Environment Variables](#) (see page 532)
- [Environment Variables for Jobs and Actions](#) (see page 533)
- [Monitor Workload Status](#) (see page 534)

## How CA NSM Job Management Option Works

The CA NSM Job Management Option automates submission and tracking of batch processes. It facilitates the definition and enforcement of complex interrelationships among units of work (including predecessor controls, manual tasks, and cause-and-effect scheduling policies). This chapter explains how to implement your workload policy through the CA NSM Job Management Option.

Today's production control administrators need flexibility to control their view of the job flow. Sometimes they need to see the big picture; at other times they must focus on a specific job or group of jobs. Jobflow provides the tools to meet those needs. Jobflow is a feature of the CA NSM Job Management Option that expands your view of the workload to include multiple job relationships.

CA NSM Job Management Option is made up of two major subsystems—the *CA NSM Job Management Option job server* and the *Unicenter Universal Job Management Agent*. Typically, sites have one job server and multiple job agents. The job agents run on all the computers where the job server component will submit and track work.

## CA NSM Job Management Option Job Server

The job server is responsible for scheduling work, sending job submission requests to the job agent, processing user commands, and job tracking. It consists of two functional areas:

- The *Monitor*, which identifies jobs that should run and determines when and where they should be submitted
- The *job tracker*, which collects job event data from all job agents where jobs have been submitted and uses that data to update job status and resource availability information

## Unicenter Universal Job Management Agent

The Unicenter Universal Job Management Agent runs jobs on behalf of the job server and sends workload-related information to the central job tracker. The Unicenter Universal Job Management Agent processes should be running on every computer where the CA NSM JM Option submits and tracks work (including the server where the job server resides).

The job agent consists of two functional areas:

- The *remote submission agent* executes the submit file specified by the Monitor.
- The *remote tracking agent* collects job event data and sends it to the job tracker.

The following section describes how the CA NSM JM Option processes act together to automatically schedule, submit, and track a unit of work. The Monitor determines when it is time to submit a job and sends a submission request to the appropriate agent. The job server and job agent together accomplish the tasks of scheduling, submitting, and tracking your jobs as follows:

1. The remote submission agent receives the submission request, which includes data items like submit file name, user ID information, and parameters required by the submit file.
2. The remote submission agent performs the submission. When the job starts, the submission agent records the event in the CA NSM JM Option checkpoint file.
3. Shortly thereafter, the tracking agent reads and forwards the event data to the job tracker.
4. The job tracker marks the job as started.
5. The same flow of events occurs when the job terminates.

## CA NSM JM Option Profiles

You define your CA NSM JM Option policy through profiles that specify where, when, what, and how work is to be performed. It is through these profiles (described in the following list) that CA NSM JM Option manages your workload.

### **Calendar profile**

Defines the days on which jobs can run.

### **Station profile**

Identifies the location *where* work is to be performed. A station can be a CPU where jobs are processed or a physical location where manual tasks, such as distributing reports, are performed.

### **Station group profile**

Defines a group of stations where work can be performed so that you can schedule work to multiple stations or the best available station in a group.

### **Job profile**

Specifies *what* work is to be performed. There are various types of jobs; for example, a CPU job specifies a command or program that runs as a background process.

### **Jobset profile**

Generally shows *how* the work is to be performed. A jobset groups logically related jobs into a set and defines scheduling criteria for the group as a whole. It provides default values for all jobs grouped in the jobset.

### **Resource profile**

A resource profile specifies the hardware, software, or supporting resources that must be available before a job can be submitted.

### **Predecessor profile**

A predecessor profile controls the sequence or order of work being processed. A jobset or job predecessor specifies a list of jobs or jobsets that must complete before a job can be run.

A predecessor profile is accessed from a notebook page on the Job Detail and Jobset Detail windows.

## CA NSM JM Option Variables

In addition to predecessors, the CA NSM JM Option lets you specify the following variables:

- *Early start time* indicates the earliest time a jobset can be marked as started or a job can be submitted for processing. (Think of this as the “don’t start before” time.)
- *Must start time* indicates the latest time that work can be started without being considered late. The work can start anytime prior to the indicated time, but not before the early start time.
- *Must complete time* indicates the time by which work must be completed; otherwise it is considered late. (Think of this as the “must be finished by” time.)
- *Maximum time* indicates the maximum length of time that work can run.
- Advanced calendar override criteria let you set conditions for holiday and non-workday processing.

## Types of Job Scheduling

The use of calendars to schedule jobs is called *predictive scheduling*. Jobs are scheduled and automatically placed in the tracking file on the days you specify, using calendars.

The use of triggers and message actions to schedule jobs is called *event-based scheduling*. Event-based scheduling is event driven--a message action starts when a trigger is tripped.

Both approaches have inherent strengths and weaknesses, and most enterprises use a combination of the two, using predictive scheduling for the bulk of their workload and event-based scheduling for exception event processing.

## How to Specify Where to Perform Work

The first step in implementing a CA NSM JM Option policy is to identify the locations (the *where*) and operations (the *what*) used in running your current work.

You can identify “where work gets done” through the definition of station profiles and station group profiles, which allow you to assign logical names to the locations where tasks are performed. Think of each station as the place where one or more steps in a procedure are performed.

A typical production job involves more than just executing a program. Often, job setup and post-processing requirements must also be performed to ensure the correct processing of jobs. Recognizing this, the CA NSM JM Option provides three categories of Job Management stations:

**PRECPU**

Specifies the location where a manual task, such as loading a printer with special forms, must be performed prior to running jobs on the CPU.

**CPU**

Specifies the computer where a job actually runs.

**POSTCPU**

Specifies the location where a manual task, such as verifying or distributing printed reports, must be performed after the CPU completes its work.

By providing the same level of attention to critical non-CPU tasks that you do to CPU-based processes, the CA NSM JM Option helps you ensure that jobs are set up correctly and have appropriate manual checks and balances.

For procedures to specify where to perform work, see the following topics in the online *CA Procedures*:

- Defining Station Profiles
- Defining Station Group Profiles

## How to Identify Resource Requirements for Workload Balancing

**Note:** If you did not set workload balancing capabilities for the CA NSM JM Option, go to *How to Schedule Work by Dates*. Identifying resource requirements is necessary only if you will be using advanced workload balancing.

The CA NSM JM Option maintains a list of available resources and the special resource requirements of any given job to automatically balance an active workload. This feature is typically referred to as *workload balancing*. Workload balancing enables the CA NSM JM Option to avoid bottlenecks and resource contention problems on a single computer or across a network of computers.

Tape drives, for example, are serial-use devices. It would be inappropriate to simultaneously submit two jobs, each requiring a tape drive for processing, on a computer that has only one tape drive. Between these two unrelated jobs, a special resource requirement would exist that requires coordination (in this case, serial processing).

To derive maximum benefit from workload balancing, you must identify those resources that have special usage or access coordination requirements and define them to the CA NSM JM Option as resource profiles.

Changes made to resource profiles affect all jobsets and jobs that reference the resources, including work currently scheduled for processing (entries in the *tracking file*). Work currently scheduled for processing can be monitored or updated using Jobset Status and Job Status.

For procedures to identify resource requirements, see Defining Resource Profiles in the online *CA Procedures*.

## How to Schedule Work by Dates

The next step in implementing a CA NSM JM Option policy is to define a default calendar and any other calendars necessary for scheduling jobs. The default calendar, typically named BASE (base on UNIX/Linux platforms), is associated with any jobset for which no calendar has been explicitly specified. Job profiles, however, are handled differently. If you do not specify a calendar name in the job profile, the value "Def" appears in the *Calendar* field to indicate that the job's calendar setting, by default, inherits the jobset's calendar.

**Note:** A job is eligible for processing only when the jobset to which it belongs is also eligible for processing. Jobs can be excluded from running on certain days that the jobset is selected by using a different calendar for the job that has those days turned off.

On Windows, you can assign a name other than BASE to the default calendar on the Job Management Option Options page in the EM Classic Configuration Settings notebook.

On UNIX/Linux, assign a name other than base to the default calendar by setting the CA NSM JM Option configuration option environment variable CAISCHD0009. For information about setting this variable, see Important Configuration Environment Variables.

By associating a jobset or job with a calendar, you define the days on which a jobset or job is eligible to run. The CA NSM JM Option does not use any time settings defined in a calendar profile. The times that jobs run are determined by the settings in the jobset and job profiles you define.

**Note:** You may want to establish a naming convention for your Job calendars that clearly distinguishes them from the calendars used by other CA NSM functions.

## Expanded Calendar Processing

As an alternative to using a calendar to let jobsets, jobs, or triggers run on specific days, you can use expanded calendar processing to define conditions under which jobsets, jobs, or triggers run.

Expanded calendar processing lets you set conditions for holiday and non-workday processing. By setting the holiday action and non-workday action switches, you can tell a job what to do on these days; you would not need to maintain special calendars to handle the situation.

Using expanded calendar processing also can reduce the number of calendars you have to maintain. For example, when you want one job to run on Mondays and another job to run on Tuesdays, you have the choice of defining two calendars or setting expanded calendar processing criteria for a single holidays calendar.

Use the Profile notebook page on the Jobset, Job, or Trigger Detail windows to extend your calendar processing criteria. The fields on the Profile notebook page are as follows:

### **Holiday calendar**

The name of the holiday calendar to use when testing any of the criteria keywords. An entry in this field overrides any other calendar processing that may be set on the Jobset or Job Main Info notebook page.

### **Holiday action**

The action to take if the selected day falls on a holiday.

### **Non-workday**

The action to take if the selected day falls on a non-workday (Saturday or Sunday).

### **Adjustment**

The adjustment in days (positive or negative) to make when the other criteria do not meet expectations.

### **Conditions**

The selection criteria keywords to use when processing this calendar. This field allows Boolean logic to create compound conditions.

For a complete description of these fields and conditions, see Job/Jobset/Trigger Profile - Detail in the online help.

For procedures to schedule work by dates, see Defining Calendars in the online *CA Procedures*.

## How to Form Groups of Related Tasks (Jobsets)

Once you define at least one calendar and a station profile, you can define a *jobset*. A jobset is a collection of one or more logically related jobs that are typically part of the same production run.

The jobset profile establishes the relationships, criteria, and default attributes of the jobs that you later define as members of this jobset.

**Note:** Every job must belong to a jobset.

### Jobset Resources

Use of jobset resources is optional. If you do not want to use workload balancing with jobsets, go to Jobset Predecessors.

Jobset resource profiles specify the minimum resource requirements for each job in the jobset.

Because jobsets do not allocate or require resources, a resource identified for the jobset does not have to be available for the jobset to be declared eligible for processing and selection to the tracking file. (Jobsets never have a status level of WRSRC—Waiting for Resources—and are never included in workload balancing.)

Resource requirements specified at the jobset level define the minimum resource requirements that must be available before any individual jobs of the jobset are considered eligible for submission.

### Resource Amount, Weight, and Usage

The jobset resource profile references a previously defined CA NSM resource profile that describes resources required by the jobs in the jobset. Amount describes how much of this resource is needed. Weight indicates the relative importance of a resource. Usage indicates how this resource will be used by the jobs in the jobset and can be specified as one of the following:

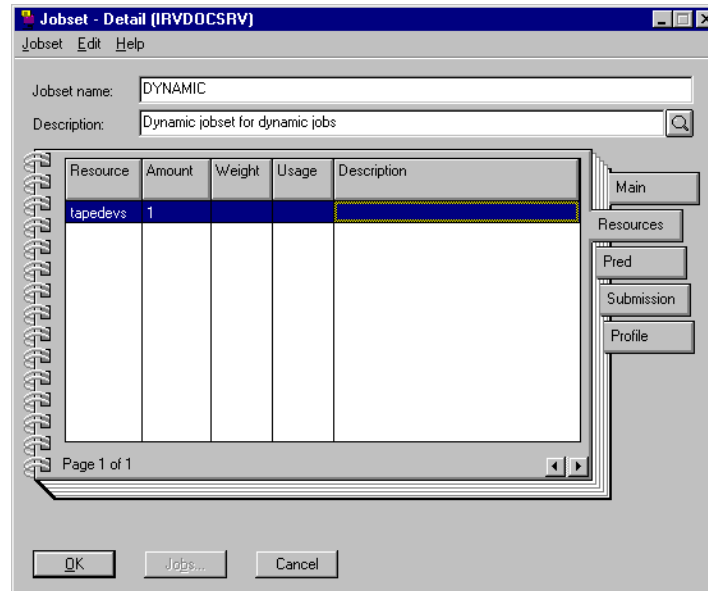
- PRIVATE—each job needs exclusive access to a resource.
- SHARED—multiple jobs share a resource.
- COREQ—jobs can access the resource only when another job has PRIVATE access.

**Note:** COREQ control typically is used to direct jobs that have a particular service to the CPU where the system is active. Consider, for example, a job that has a CPU dependency of *Any available CPU* and requires a co-requisite resource named db. This job can only run on a CPU where another job, which has PRIVATE use of the resource db, is currently active.



An example of resource allocation follows. Assume that an existing resource profile named `tapedevs` specifies that four tape drives exist on station `europa`.

The following Jobset - Detail Resources window indicates that each job in the jobset requires one of the four tape drives defined as available resources by the previously defined resource profile named `tapedevs`, and each job needs exclusive access to the tape drives.



**Note:** The CA NSM JM Option does not verify the physical existence or actual availability of tape drives or any other device. These resource definitions are logical rather than physical.

## Jobset Predecessors

Use of jobset predecessors is optional. If you do not want to use predecessor relationships at the jobset level, go to [How to Identify Work to Perform](#).

Jobset predecessors are used to specify the list of jobs, jobsets, or trigger profiles that must complete successfully before this jobset can be started. A predecessor requirement will be marked satisfied if any of the following conditions are true:

- The predecessor completed successfully (COMPL).
- The predecessor aborted and shows a status code of ABORT but has an *Abend action* of CONTINUE in its profile.
- The predecessor is specified as dynamic and is missing from (that is, does not exist in) the current workload tracking file.

## Nonexistent Predecessor Evaluation

If a jobset has a dynamic predecessor requirement (the default) that specifies a job, jobset, or trigger that is not in the tracking file (not currently scheduled to run) at the time the predecessor requirement is evaluated, that predecessor requirement is automatically satisfied through an implicit posting. In effect, it is ignored. If the predecessor requirement type is static, no implicit posting takes place. The jobset remains in "wait on predecessor" state (WPRED) until the predecessor requirement is satisfied.

The reason the CA NSM JM Option performs this implicit posting with dynamic predecessor requirements is simple. To have a dependency on a job that does not exist is illogical, because any job with such a predecessor dependency would never run. As such, if the CA NSM JM Option detects such a condition—a predecessor requirement referencing a job not scheduled to run (that is, does not exist in the current workload)—it ignores the illogical predecessor requirement and continues to evaluate eligibility.

You can use static predecessor requirements in those instances where a jobset must not run unless its predecessor job completes first. A jobset can have both static and dynamic predecessor requirements, but cannot reference the same predecessor requirements as both dynamic and static.

## Canceled Predecessors

The only predecessor requirements honored are those that represent predecessor jobs that are scheduled to run in the current workload (also referred to as being "in the tracking file"). Because of this rule, the cancellation of any predecessor job (which removes that job from the tracking file) results in that predecessor requirement being ignored. The cancellation of a predecessor job that was referenced as dynamic effectively satisfies the predecessor requirement, allowing any successor jobs to run. If the predecessor requirement was referenced as static, any successor jobs remain in a "wait on predecessor" state (WPRED).

While many enterprises find this behavior intuitive and useful for dynamic predecessor requirements, others do not. To support each enterprise's preference, the CA NSM JM Option components provide an option that lets you change the default effect of the cancel command so that successors (those jobs that define the canceled job as a predecessor) are not automatically posted. For information about specifying this option using the Post on Cancel and CAISCHD0014 environment variables, see Configuration Environment Variables.

For procedures to form groups of related tasks, see Defining Jobsets in the online *CA Procedures*.

## How to Identify Work to Perform

A job is a unit of work performed. A job can be as follows:

- An executable file used to process an application
- A manual task to be performed before or after processing an application

Before reading about how to define jobs, review the following sections that describe how the CA NSM JM Option evaluates the policies you define to determine when jobs are eligible to run.

### Jobset Membership

Every job must belong to a jobset, and only after a jobset is marked as started will Job Management evaluate the jobs in the jobset.

**Note:** DYNAMIC jobs are a special exception to this rule and are discussed separately in Demand a DYNAMIC Job.

### Early Start Time

Early start time indicates the earliest time a job may start. For CPU jobs, this is the earliest time they would be submitted to the operating system. Once the job's jobset has started, the CA NSM JM Option checks the job's early start time. If you specify an early start time for the job, the job remains in WSTART status (waiting for the job's early start time) until that time arrives. When the time arrives, or if the job does not have an early start time specified, the job advances from WSTART to WPRED status (waiting for predecessor requirements to be satisfied).

## External Predecessors

The EXTERNAL job type setting in a job profile lets you define a job to the CA NSM JM Option that will be submitted by an external job manager. The job is actually submitted by another job management system such as Unicenter<sup>®</sup> CA-7<sup>®</sup> Job Management (Unicenter CA-7), Unicenter<sup>®</sup> CA-Scheduler<sup>®</sup> Job Management (Unicenter CA-Scheduler), Unicenter<sup>®</sup> CA-Jobtrac<sup>®</sup> Job Management (Unicenter CA-Jobtrac), and so forth. The CA NSM JM Option treats this job as a non-CPU job with a class of EXTERNAL (similar to PRECPU and POSTCPU).

The job's presence on the tracking file causes JOBINIT and JOBTTERM triggers to be generated internally and the job is automatically tracked by the client where the job actually runs. The job server tracks the job, but does not submit it. Since this is a JOB base entry, predecessor definitions and calendar selection can reference it.

You specify the job type at the Main - Info notebook page of the Job - Detail window.

## Job Resources

Job resource requirements do not override jobset resource requirements. Rather, resource requirements defined at the job level are added to any resource requirements that may already be defined at the jobset level. Jobset and job resource requirements are therefore cumulative.

## Job Submission

The CA NSM JM Option submits CPU jobs based on user-defined processing requirements. You define what to submit and any special considerations through a submission profile that assigns these submission characteristics:

- User ID (the user for whom the job is submitted)
- Name of the file or program to submit
- Optional parameter values to pass to the submitted file or program
- Password of the user for whom the job is submitted
- Domain of the user for whom the job is submitted (Windows only)

## Job Predecessors

Jobs, jobsets, and triggers can all be defined as predecessors for a job. All predecessors defined for the job must be satisfied before a job can become a candidate for submission.

**Important!** The jobset starts only after all of the jobset's predecessor requirements have been met, and only after the jobset starts will the CA NSM JM Option evaluate the predecessor requirements for the individual jobs that are members of that jobset.

On UNIX/Linux, you can also specify advanced system conditions that must be met before a job can run on a specific node.

For example, if a job must not run when a particular user is logged on, you can define a system condition criterion that specifies this. The CA NSM JM Option holds the job until that user is no longer on the system and then releases it.

SYSCON objects for administering system condition requirements are available using `cautil` command syntax and are described in `JOBSYSCON`, `JOBSETSYSCON`, and `TJOBSYSCON` in the online *CA Reference*.

## Password Validation for Job Submission (UNIX/Linux)

By default, job submission requires a password and a valid user ID for a job to be executed.

You can change this rule at any time after installation by executing the `cawrksec` utility located in the `$CAIGLBL0000\sche\bin` directory. The utility allows only the `uid 0` user to maintain the file and preserve the file permissions. The file can also be maintained using a UNIX/Linux text editor. For more information about using the `cawrksec` utility, see the online *CA Reference*.

The ExtNodeL.sch configuration file is located in the \$CAIGLBL0000\sche\config directory. You can use this file to maintain policies that specify how password validation is to be performed based on the submitting node and user ID. The file must be owned by root, and only a uid of 0 may have write access to it. An individual entry in the file has the following format:

```
-n=nodename, user-id, flag
```

where:

**nodename**

Specifies the node from which the job is initiated; it can contain a trailing generic mask character.

**user-id**

Specifies a user ID to whom the rule applies; it can contain a trailing generic mask character.

**flag**

Specifies D for disable (perform no password authorizations), E for enable (unless the proper password is supplied, the job will not run), or W for warn (check the password; if invalid, run the job but issue a warning message).

## Examples

The following rule is the default rule in effect if you elected to enable password checking during installation. The rule states that for all nodes and all users password validation is to occur.

```
-n=*,*,E
```

The following rule is the default rule in effect if you elected to disable password checking during installation. The rule states that for all nodes and all users password validation is bypassed.

```
-n=*,*,D
```

The following combination of rules only enforces a password validation on user root and allows all other users to bypass password validation.

```
-n=*,*,D
```

```
-n=*,root,E
```

The following combination of rules allows all users to bypass password validation unless the request comes from the node mars. In that case, password validation is enforced for all users. The last entry sets a warning type password validation for user root if it comes from a node other than mars.

```
-n=*,*,D
-n=mars,*,E
-n=*,root,W
```

Job Management scans the entire configuration file for a best match and uses that rule. It uses the node field as a high level qualifier when searching for a best match. For example, if the following entries are the only two entries in the file, any request coming from the node mars uses the enforce rule. The user root only uses the warning rule if the request comes from a node other than mars.

```
-n=mars,*,E
-n=*,root,W
```

## Cyclic Job Submission

The CA NSM JM Option supports *cyclic job* submission, which allows a job to be submitted multiple times during a particular time period. For example, if you have a job that needs to run every hour between 8:00 a.m. and midnight, you can define the job as *cyclic* and specify the *frequency* (how often it should run) and the number of *iterations* (how many times it should run). Cyclic job submission is much easier than defining 17 jobs, one for every hour between 8:00 a.m. and midnight. It also simplifies maintenance because you need only change a value once to alter the cycles, frequency, or number of iterations.

**Note:** You should define cyclic jobs as ineligible for backlogging to ensure that the job is purged from the tracking file at the next new-day autoscan.

## Cyclic Job Special Considerations

Processing cyclic jobs has special considerations of which you should be aware if you plan to take advantage of this powerful facility. The following discussion provides detailed explanations of these special considerations.

When a cyclic job is demanded into the tracking file after its early start time has passed, the CA NSM JM Option adjusts for this situation by doing the following:

- Adding one instance of the job using the normal early start time job, so that the job starts immediately
- Submitting *x* instances of the job to start at the first eligible time interval after the current time

For example, if the job has an early start time of 1:00 p.m. and a frequency of 60 minutes, and the job is demanded at 3:15 p.m., the CA NSM JM Option generates entry one to start at 1:00 p.m. When the job enters the current workload (by entering the tracking file), the CA NSM JM Option recognizes that it is late (it should have started at 1:00 p.m.) and starts it immediately. The next submitted instance of this job has a start time of 4:00 p.m. with subsequent entries to start at 5:00 p.m., 6:00 p.m., and so forth.

By default, when a cyclic job is brought in late (demanded), the CA NSM JM Option skips the jobs that are too late to run and only runs the number of iterations left, based on the calculation of when the jobs should have run. If you set the Windows configuration setting "cycle count precedence" or the CAISCHD0540 environment variable on UNIX/Linux to Y, the cycle count takes precedence over the remaining count based on current time. Thus, the defined number of cyclic jobs is scheduled up to the time of the new-day autoscan.

Using the default, consider, for example, a cyclic job that has a frequency of 30, a cycle count of 16, and an early start time that defaults to the new-day autoscan time of 1:00 a.m. The job is typically brought in at 1:00 a.m. and runs 16 times: at 1:00 a.m., 1:30 a.m., 2:00 a.m., and so forth until the last (sixteenth) run at 8:30 a.m. However, if the job is demanded in at 4:05 a.m., the occurrences up to that time are skipped and the job is run at 4:05 a.m., 4:30 a.m., 5:00 a.m., 5:30 a.m., and so forth. At 8:30 a.m., the job runs for the last time, for a total of 10 runs. The occurrences that would have run (at 1:00 a.m., 1:30 a.m., 2:00 a.m., 2:30 a.m., 3:00 a.m., 3:30 a.m.), are skipped. The CA NSM JM Option does not attempt to "catch up" if it means running the jobs after the last calculated start time of 8:30 a.m.

If you enable cycle count precedence using the above example, the jobs would run at 4:05 a.m., 5:00 a.m., 5:30 a.m., and so forth, until the last (16th) ran at 12:30 p.m. All counts up to the time of the new-day autoscan are scheduled.

If you elect to run with the autoscan process disabled, cyclic jobs are submitted only when you execute a manual autoscan. The jobs are not automatically submitted each day. For an explanation of the autoscan process, see Autoscan.

You cannot define a cyclic job as a predecessor to itself. In other words, when a cyclic job is submitted, the job begins execution when its early start time is reached. It does not wait for any preceding instances of itself to complete.

You can define a cyclic job as a predecessor to another job or jobset. All occurrences of the cyclic job are treated as predecessors and must complete before the successor job or jobset is eligible to run.



The predecessor criteria for cyclic jobs include a qualifier option that allows cyclic jobs to run in one of two ways. For example, assume you have two cyclic jobs, CycJobA and CycJobB, where both have an iteration number of 3, frequency is set to 60, and CycJobA is the predecessor of CycJobB. When both jobs are demanded into the tracking file, the file appears as follows:

```
CycJobA  QUAL=xx01  
CycJobA  QUAL=xx02  
CycJobA  QUAL=xx03  
CycJobB  QUAL=xx01  
CycJobB  QUAL=xx02  
CycJobB  QUAL=xx03
```

If you set the qualifier *Use Qualifier Predecessor Criteria for cyclic job?* to N, CycJobB runs after all CycJobA iterations have run.

If you set the qualifier to Y, CycJobB QUAL=xx01 runs after CycJobA QUAL=xx01 completes, CycJobB QUAL=xx02 runs after CycJobA QUAL=xx02 completes, and so forth.

If you plan to use cyclic job submission for Windows platforms, review the environment settings for the CA NSM JM Option in the Configuration Settings window (Options tab) for *Max active jobs*, *Max resources*, *Max predecessors*, and *Use Qualifier Predecessor Criteria for cyclic job?* and set the appropriate values to accommodate the additional jobs in the daily workload.

If you plan to use cyclic job submission for UNIX/Linux platforms, review the values set for the environment variables CAISCHD0025, CAISCHD0026, CAISCHD0027, and CAISCHD0040 and set appropriate values to accommodate the additional jobs in the daily workload.

For procedures to identify work to perform, see *Defining Jobs* in the online *CA Procedures*.

## How to Schedule Work by Special Events

Occasionally situations or events may arise that require a quick, consistent, and correct response to ensure that operations continue to run smoothly. The CA NSM JM Option provides facilities to identify these events and specify automatic responses.

A *trigger profile* specifies a special event for which the CA NSM JM Option should be prepared. When an event occurs that matches a trigger profile, the CA NSM JM Option automatically “trips” that trigger and sends a message to the Event Console where a message action is invoked.

For example, assume you download an updated database to your system on a weekly basis and you want a number of reports to run as soon as the file transfer is complete. To automate the sequence, you can define a File Close (DCLOSEU) trigger profile and a message action profile. When the file closes, the trigger is detected, the appropriate message is issued, and the message action profile demands a jobset containing all of the necessary report jobs.

The CA NSM JM Option lets you define triggers for the following events:

- Job initiation
- Job termination
- The caevent command
- File close (when a file, opened for write or newly created, is closed, a File close event is created)
- File unlink (deletion)
- IPL event

Once a trigger “trips,” its associated message action profile executes. For a description of message action processing, see *Trap Important Event Messages and Assign Actions* in the “Administer Critical Events” chapter.

Triggers can be defined to be associated with a calendar so that although a triggering event may occur at any time, the trigger is only tripped when its associated calendar determines it is appropriate. When a trigger profile has no associated calendar, it is always in effect; it is scheduled every day and never marked complete.

## Use caevent

For most situations you can choose a trigger type of File close, Job termination, or Job initiation. However, there may be times when the triggering event you want to define is not a system-initiated even and a logical event may be more suitable. The trigger type of CA Event lets you create these user-initiated logical events.

Using the executable caevent (provided with CA NSM), you can generate a logical event by executing the caevent command and supplying two command line arguments--the first is the logical event name, and the second optional parameter is an event status code. By running the caevent executable in scripts or using it interactively, the system alerts the CA NSM JM Option to events as they take place.

In addition, since the event name and status code are user-specified, they can describe any trigger event. For example, the following caevent command sends an event of type caevent, with an event name of bkupnow (a previously defined trigger profile), and a status code of 10 to the CA NSM JM Option:

```
caevent bkupnow 10
```

If a defined trigger profile matches these criteria, the trigger trips and the associated message action profile is processed. The message action profile may demand a jobset into the workload automatically or execute a cautil backup command. Message action profiles are flexible enough to meet your needs. For additional information about the caevent command, see the online *CA Reference*.

### Triggers as Predecessors (UNIX/Linux)

Often, events not related to a job need to be tied directly to a job or jobset, such as creation or deletion of a file, completion of a job that ran on another server, or occurrence of an IPL event. These events can be detected by the trigger mechanism in the CA NSM JM Option and can be used to satisfy a job dependency.

For instance, you may not want a payroll job to start until the data file has transferred from another system. You can define a trigger to watch for the creation of the data file and define the payroll job with a predecessor to wait for the trigger.

In another scenario, CA NSM JM Option server SERVA runs a job on system SYSA. CA NSM JM Option server SERVB cannot run a job until the job submitted by SERVA completes. Define a trigger on SERVB to watch for the completion (JOBTERMU) of the job submitted by SERVA, and define this trigger as a predecessor to the job on SERVB.

**Note:** For additional information about keywords that let you use trigger profiles as predecessors, see the Control Statements for JOBSETPRED and JOBPRED objects in the online *CA Reference*.

For procedures to schedule using triggers, see Defining Trigger Profiles in the online *CA Procedures*.

## Run a Job on Demand

Many installations' jobs are not performed on a predictable schedule or in response to any particular event but are run "on demand."

You can use the Jobset Demand option on the Jobset list container window to tell the CA NSM JM Option to demand a jobset that is defined in the Job Management Database into the current workload. Demanding a jobset lets you bring jobsets and jobs into the current day's workload (into the tracking file) that would not be automatically selected for processing.

When you demand a jobset into the tracking file, the individual jobs in the jobset are also brought into the tracking file if both of the following conditions are true:

- The job profile specifies Autoselect=YES.
- The calendar profile for the job indicates that today is a valid execution day.

Both of these conditions must be satisfied to bring jobs into the tracking file when you demand the jobset of which they are members. Otherwise, you must specifically demand the job into the tracking file using the Job Demand option on the Jobs list container window.

## Demand a DYNAMIC Job

You can demand a job into the current workload that is not defined in the job Management Database by designating the job as DYNAMIC. This demand job, while not in the job Management Database, can have the following CA NSM JM Option capabilities:

- Its messages go to the Event Console.
- History information is written to the CA NSM JM Option history database.
- The running process appears on the Jobflow status GUI.
- The job can be tracked and canceled by the CA NSM JM Option.
- You can specify basic parameters such as start time and effective date.

The next autoscan removes all references to the job.

For procedures to run a job on demand, see Demanding a Job into the Workload in the online *CA Procedures*.

## How to Test Your CA NSM JM Option Policy Definitions

After you have defined your workload policy to the CA NSM JM Option, we recommend that you run the Simulation Report to “see what happens” before committing the policy to your daily production cycle.

The Simulation Report invokes a simulated autoscan process, producing a set of detailed reports that identify:

- Jobs that would be selected, executed late, or carried over to the next day (backlogged).
- Date and time that each job would be processed.
- Location where each job would be processed.
- Resources that would be required.
- Amount of utilization for each resource.

The information obtained from the Simulation Report will help you to evaluate your implementation and more readily identify any adjustments that may be required to the workload policies you have defined.

Run the following command from a command line prompt to create a simulation report:

```
wmmodel
```

The `wmmodel` executable lets you play “what if” scenarios, override job dependencies, and change the duration of a job execution.

**Note:** To run a CA NSM JM Option report on UNIX/Linux platforms, you must be logged in as an authorized user such as `root`.

The following command also runs the Simulation report.

```
schdsimu BOTH
```

Report output goes to `stdout`, which can be directed to any printer or file, using standard redirection operators.

For additional information about CA NSM JM Option reports, see the online *CA Reference*.

## How to Run Additional CA NSM JM Option Reports

The CA NSM JM Option shell scripts and programs shown in the following list generate the standard CA NSM JM Option reports.

On Windows platforms, the report generating programs are located in the %CAIGLBL0000%\bin directory.

On UNIX/Linux platforms, the report generating shell scripts are located in the \$CAIGLBL0000/sche/scripts directory.

<b>Windows</b>	<b>UNIX/Linux</b>	<b>Creates Report</b>
schdchk.cmd	schdcheck	Check Report
schdfore.cmd	schdfore	Forecast Report
schdhist.cmd	schdhist	History Report
schdpxr.cmd	schdpxr	Predecessor/ Successor Cross-Reference Report
schdsimu.cmd	schdsimu	Simulation Report
wmmodel.exe	wmmodel	Simulation Report

## Autoscan

The selection, submission, tracking, and cleanup of jobs begin with the autoscan procedure. The CA NSM JM Option treats a workday as a 24-hour time period that begins at new-day time. The default setting for the new-day autoscan is midnight for Windows platforms and 01:00 a.m. for UNIX/Linux platforms. At this time, all the jobsets and jobs that qualify are added to the new-day's workload (added to the tracking file). A periodic autoscan occurs at specific intervals to determine if any changes or additions should be made to the current day's workload.

New-day autoscan performs the following tasks:

- Cleans up the tracking file by purging finished work from the previous day's workload. Any unfinished work that is backlog eligible is carried over to the new day's workload.
- Scans the job Management Database to select the new day's work. Autoscan first selects all jobsets that qualify for processing the current day (based on their calendar, expiration, and effective date criteria). From these jobsets, autoscan similarly selects the individual jobs that qualify. Once these jobsets and their jobs are in the current day's workload (in the tracking file), they are ready for processing, review, and any last minute changes.

---

## How a Job or Jobset Qualifies for Selection During Autoscan

To determine whether a jobset (and job) qualifies for automatic inclusion in today's workload, autoscan examines several operand values in the following sequence:

1. **Auto selection**—Is the jobset or job eligible for automatic selection? Typically, jobs and jobsets that run on a regular schedule use calendars and have the auto selection operand set to YES. If the auto selection operand is set to NO, the jobset or job is ignored and no further evaluation is performed.
2. **Effective date and expiration date**—Has the effective date in the jobset or job profile been reached? Is the expiration date in the jobset or job profile still in the future?
3. **Skip count**—Has the workload administrator indicated that this jobset or job should be skipped? The workload administrator would specify a positive integer for the *Skip* box of the jobset or job to be skipped. When the autoscan encounters a jobset or job with a non-zero skip count, it automatically ignores that jobset or job and decreases the skip count. When the skip count is again zero, either because the workload administrator explicitly set it to zero or because the jobset or job has been "skipped" the necessary number of times, the jobset or job again becomes eligible for automatic selection.
4. **Calendars**—Does the calendar named by the calendar operand in the jobset or job profile indicate that today is a workday? Jobsets and the jobs subordinate to them, whose calendars indicate that today is a workday, are automatically brought into the tracking file during autoscan. The CA NSM JM Option uses only the date information provided by a calendar and ignores any time information. The CA NSM JM Option gets the early start time from the job or jobset profile.

## Cleanup and Backlogging

What happens to jobs that are not completed when the next new-day autoscan occurs? The answer to this question depends on whether the job is eligible for backlogging.

Jobs that are unable to start on the day they are selected and have Backlog=YES in their profile (or default to Backlog=YES through their Jobset profile) are backlogged. When an unfinished job is backlogged, it remains in the tracking file as part of the workload for the following day and is automatically rescheduled for processing.

Jobs that are unable to run on the day they are selected and which have Backlog=NO in their profile (or default to Backlog=NO through their jobset profile) are automatically canceled and purged from the tracking file as part of the new-day autoscan.

Jobs that are running when the new day's autoscan occurs and have Backlog=RUNNING in their profile are backlogged. When a running job is backlogged, it remains in the tracking file as part of the workload for the following day and is automatically rescheduled for processing. If the job is not running during the autoscan, the job is removed from the tracking file.

**Note:** The time of the new-day autoscan should occur after the expected completion of all scheduled jobs. This makes the Backlog option more meaningful. If the new-day autoscan runs before the day's work is completed, all unfinished jobs are purged or backlogged.

## Workload Processing

The processing of the current day's workload (work in the tracking file) is status-driven. The CA NSM JM Option constantly monitors and tracks jobs, moving jobs from one status level to the next only after specific conditions or requirements have been satisfied.

Jobsets and jobs advance through various status levels during workload processing. The autoscan process brings qualifying jobsets and jobs into the tracking file with an initial status of LOADNG. After all the new work is loaded, the CA NSM JM Option marks it WSTART (Waiting to Start). Autoscan is then complete.

After bringing all new qualifying jobsets and jobs into the current day's workload (into the tracking file), the CA NSM JM Option processes the workload by reviewing:

1. Early start time
2. Predecessors
3. Priority
4. Resource usage and availability

Jobsets are processed first. Jobset criteria are evaluated prior to criteria for jobs in the jobset. First, the CA NSM JM Option checks that a jobset's early start time has arrived; second, it checks to see if all predecessor requirements for the jobset have been satisfied.

When the jobset is in START status, the CA NSM JM Option looks at similar criteria for jobs in the jobset. A jobset must be marked as started (START) before its jobs can be evaluated. For example, when a job is in WSTART (Waiting to Start) status, the CA NSM JM Option evaluates its early start time. If the early start time has been reached, the job is placed in WPRED (Waiting for Predecessors) status. Otherwise, the job stays in WSTART until the specified early start time is reached.



The CA NSM JM Option prioritizes jobs based on resource requirements (if any), early start time, and priority. Although you can fine-tune the sequencing of your workload using early start time and priority, you typically use predecessors as the primary means to establish the sequence of workload jobset and job processing.

When all the jobs in the jobset are complete, the jobset is marked complete (COMPL). A jobset is also marked complete if none of its jobs meet the criteria for selection, or if there are no jobs in the jobset.

To summarize, the CA NSM JM Option first looks at the jobset to determine that early start time has arrived and predecessor requirements are satisfied. The CA NSM JM Option then looks at the jobs in the jobset to determine that early start time has arrived, predecessor requirements are satisfied, and that resources, must complete time, and priority can be satisfied.

## Maintenance Considerations

The following topics describe how to maintain the CA NSM JM Option.

- Job Management Logs
- Tracking File
- Undefined Calendars During Autoscan
- Unload Job Management Database Definitions to a Text File
- Purge Old History Records (UNIX/Linux)
- How to Submit Jobs on Behalf of Another User

### Job Management Logs (UNIX/Linux)

Whenever you start the CA NSM JM Option, all but the last ten versions of the audit trail log files are automatically purged by an implicit execution of the cleanlog shell script. The cleanlog script is located in the `$CAIGLBL0000/sche/scripts` directory.

You can run the cleanlog script whenever you want to purge old audit trail log files by entering the following command:

```
$CAIGLBL0000/sche/scripts/cleanlog
```

The `$CAIGLBL0000/sche/log/$host` directory contains the following job management log files:

**mtrsuf**

Identifies the current suffix number of the `schdxxx.nnnn` files; descriptions of these files follow. The version number can be a maximum of 9999.

**ca7xxx.log and ca7xxx.out**

Represent the stdout and stderr for each of the CA NSM JM Option daemons. There is only one version of each of these files; `ca7xxx.out` provides an audit trail of all display command processing that flows through the CA NSM JM Option. The files are rebuilt each time the CA NSM JM Option is stopped and restarted.

**schdlcs.nnnn**

Provides an audit trail of all update command processing that flows through the CA NSM JM Option, including trigger commands. A new version of this file is created each time the CA NSM JM Option is cycled and at midnight every day; `nnnn` is identified by the `mtrsuf` file.

**schdtrk.nnnn**

Provides an audit trail of all tracking activity by the CA NSM JM Option. A new version of this file is created each time the CA NSM JM Option is cycled and at midnight every day; `nnnn` is identified by the `mtrsuf` file.

**schdmtr.nnnn**

Provides an audit trail of all submission and selection activity by the CA NSM JM Option. A new version of this file is created each time the CA NSM JM Option is cycled and at midnight every day; `nnnn` is identified by the `mtrsuf` file.

## Tracking File

Work remains in the tracking file until it is cleaned out during the new-day autoscan, or until you manually cancel and purge it from the current day's workload.

Because any failed jobs (and their jobsets) that are eligible to be backlogged would remain in the tracking file after the new-day autoscan, you may want to check the tracking file on a periodic basis and clean out (cancel and purge) any obsolete entries that were not automatically removed. These backlogged jobs were not eligible to be purged as part of the new-day autoscan.

## Undefined Calendars During Autoscan

To prevent incorrect scheduling of jobs, the CA NSM JM Option automatically suspends job submission (using an implicit `cautil stop autosub` command) when it detects calendar validation or reference errors. Typically, this is caused by a job or jobset reference to an undefined calendar profile.

Upon issuing this implicit `cautil stop autosub` command, the CA NSM JM Option issues a message to alert the Event Console, and, in turn, the workload administrator, that a calendar validation problem was detected during autoscan. The message provides an opportunity to identify and correct errors that may exist before any jobs are run.

After completing the necessary maintenance tasks to correct the condition, you can instruct the CA NSM JM Option real-time components to rescan the CA NSM JM Option databases for eligible jobs and jobsets by running the command:

```
cautil start autoscan
```

To resume automatic job submission, run the following command:

```
cautil start autosub
```

## Purge Old History Records (UNIX/Linux)

To delete old history records that are no longer associated with a job or jobset record in the CA NSM JM Option database, run the History Report using the `CLEAN` parameter:

```
$CAIGLBL0000/sche/scripts/schdhist CLEAN
```

**Note:** You must specify the `CLEAN` parameter in uppercase characters.

The `CLEAN` parameter deletes the associated history entries for any job or jobset record that no longer exists in the job Management Database (has been deleted). The History Report lists the history entries that have been deleted from the database.

## Unload the CA NSM JM Option Database Definitions to a Text File

You can use the `cauexpr` command to unload the CA NSM JM Option definitions to a file called `cauexpr.txt` in `cautil` type format that can then be used by `cautil`. You can upload the file if necessary, take the definitions to another computer, or use `cautil` to input them.

For a description of the syntax for the `cauexpr` command, see the online *CA Reference*.

## How to Submit Jobs on Behalf of Another User

The CA NSM JM Option lets you submit jobs on behalf of another user. To enable this policy, you must set the configuration setting "Submit on behalf of another user" on the Job Management Option Options page. The user ID used to submit the job must also be defined to Windows with the authority "logon as batch job."

For procedures to submit jobs on behalf of another user, see Submitting a Process as Another User in the online *CA Procedures*.

## Agent/Server Configurations

The CA NSM JM Option provides a flexible architecture that lets you distribute work to multiple computers where the necessary resources are available for a job to be processed.

The CA NSM JM Option has two primary components that are used in combination to provide both local and remote scheduling. These two components are as follows:

- The CA NSM JM Option server, which provides support for the database and has the capacity to build schedules for execution.
- The Unicenter Universal Job Management Agent, which processes and tracks the execution of work on behalf of the job server.

By combining these two elements, the CA NSM JM Option can provide different combinations of local and remote scheduling.

Distributed remote scheduling can be performed between a single full-function job management server and a single agent, or many agents, or among multiple job management servers installed in the same network. Each job management server can have many clients, and each job management client can be designated as a client of many servers.


The CA NSM JM Option is not restricted to having an agent service a single job management server. Regardless of platform--UNIX/Linux, Windows, or z/OS or OS/390--a server can have its work tracked for it by the Unicenter Universal Job Management Agent.

The following sections describe how you can apply this architecture to defining CA NSM JM Option policies across the following:

- A single server computer
- Multiple servers in an agent/server domain

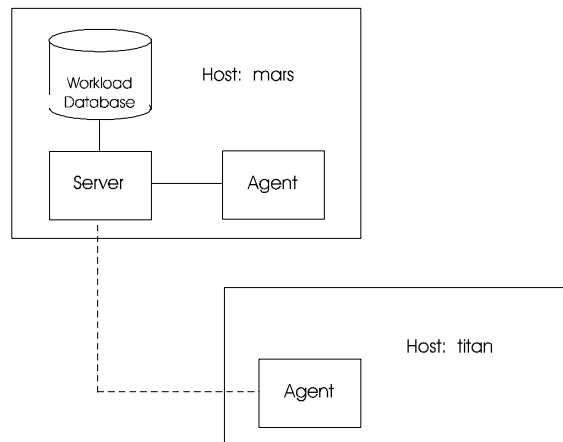
## Single Server

The CA NSM JM Option on a single server uses a local server and a local agent to define, service, and track all scheduling activity. The following diagram shows you this configuration for a single node called Mars:



## Multiple Hosts in an Agent/Server Relationship

The CA NSM JM Option in agent/server mode uses a local server, a local agent, and a remote agent to define, service, and track all scheduling activity across multiple hosts. The following diagram shows two host computers in an agent/server relationship. Mars, the server computer where the job Management Database is located, has a local agent and a remote agent. The local agent processes work to be done on mars, and the remote agent processes work sent by mars to the remote node titan.

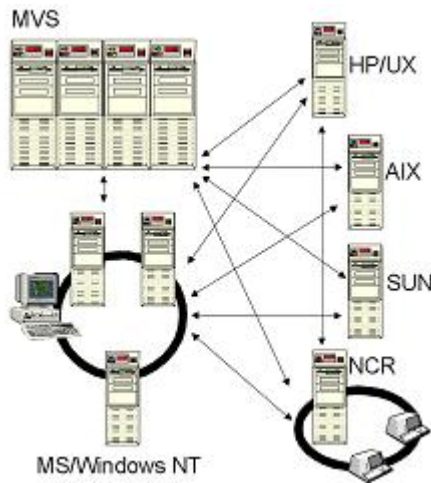


## Cross-Platform Scheduling

Scheduling, as defined in the past, was the ability to schedule units of work (jobs) within a particular host (for example, z/OS or OS/390, Windows, UNIX/Linux, AS/400). This scheduling activity started with simple definitions:

- Time dependencies. (Start Job A at 10:00 AM.)
- Unit-of-work dependencies. (Only start Job A after Job B has completed successfully.)

Today many data centers are working with more complex production workloads that cover a wider variety of processing environments—including OS/390, UNIX/Linux, Windows, OpenVMS, and other platforms communicating with each other. Such an environment requires cross-platform scheduling.



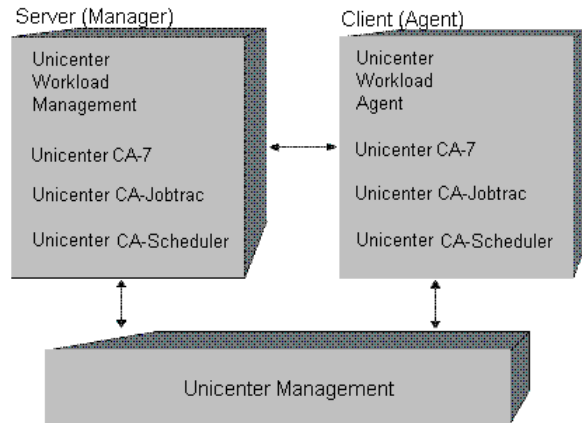
Cross-platform scheduling provides integrated enterprise control. It defines and sets dependencies. It submits and tracks the status of units of work (jobs or events) not only on the traditional platforms (z/OS or OS/390, UNIX/Linux, Windows, AS/400) but on a variety of other platforms.

Actions result in the immediate and automatic notification of status changes to the unit of work (executing, completed successfully, error) and the ability to perform additional automatic actions. These status changes can trigger event-dependent processing not only on the local platform but on other systems/resources throughout the environment to maximize enterprise processing efficiency.

CA provides distributed scheduling capabilities that enable integration of its z/OS or OS/390 (formerly MVS) products—Unicenter CA-7, Unicenter CA-Jobtrac, and Unicenter CA-Scheduler—with the CA NSM JM Option.

## Job Management Managers and Agents

Cross-platform scheduling is implemented using a manager/agent model.



In this architecture, the manager performs the following functions:

- Maintains job definitions and relationships
- Evaluates job execution and job completion information
- Uses a database for workload definitions
- Interfaces with agents to initiate jobs and collect status information

The Unicenter Universal Job Management Agent is a small set of programs that execute on each target computer where jobs will be processed. The agent performs the following functions:

- Receives job requests from one or more managers and initiates the requested program, script, JCL or other unit of work
- Collects status information about job execution and file creations
- Sends the status information to the requesting workload manager for evaluation

Many environments choose to use a centralized repository for defining jobs and monitoring their workload. The CA NSM JM Option provides a Unicenter Universal Job Management Agent so you can initiate and track jobs on a computer without maintaining a workload database on that computer. The Unicenter Universal Job Management Agents can process a request from a job management manager (the CA NSM JM Option on another computer or one of our OS/390 scheduling products) by initiating the requested process and returning tracking information about that process to the requesting job management manager. Any job management manager can submit job requests to any Unicenter Universal Job Management Agent.

## Configuring Job Management Managers and Agents

All managers and agents that participate in cross-platform scheduling require a valid CAICCI connection.

**Note:** Review the CAICCI considerations listed in the CA Common Communications Interface overview in the online *CA Reference*. Careful consideration of these issues will help to ensure that your job management scheduling tasks work across a variety of platforms.

The job management manager requires a station definition for the job agent node. The job agent does not require any additional configuration because it simply processes requests on behalf of the job management manager.

By default, any job manager can submit jobs to the agent. Even though user ID and password validation can be enforced, you may not want to allow all job managers to schedule work on various agents. Using the configuration settings option "authorized remote manager nodes" on Windows or the CAISCHD0530 environment variable on UNIX/Linux, you can configure a comma-delimited list of authorized remote manager nodes to restrict requests to certain job manager nodes.

## Implementation

Cross-platform scheduling can include several different platforms and several different products, as indicated in the following examples:


- CA NSM JM Option servers on Windows can submit work to Unicenter Universal Job Management Agents on UNIX/Linux.
- CA NSM JM Option servers on UNIX/Linux can submit work to Unicenter Universal Job Management Agents on Windows.
- Unicenter CA-7, Unicenter CA-Scheduler, and Unicenter CA-Jobtrac can submit work to Unicenter Universal Job Management Agents on Windows and UNIX/Linux.
- CA NSM JM Option servers on Windows and UNIX/Linux can submit work to Unicenter CA-7, Unicenter CA-Scheduler, and Unicenter CA-Jobtrac.

The implementation of cross-platform scheduling can be tailored to the needs of each site. Before installing the software components, you must determine where you want to implement managers and agents.




## Centralized Implementation

You can implement the CA NSM JM Option in a centralized fashion so that there is a central repository on one server. You can use this central repository to schedule and monitor work on any other server. In the following scenario, the CA NSM JM Option server is installed on one server, and the Unicenter Universal Job Management Agents are installed on the other systems.



## Decentralized Implementation

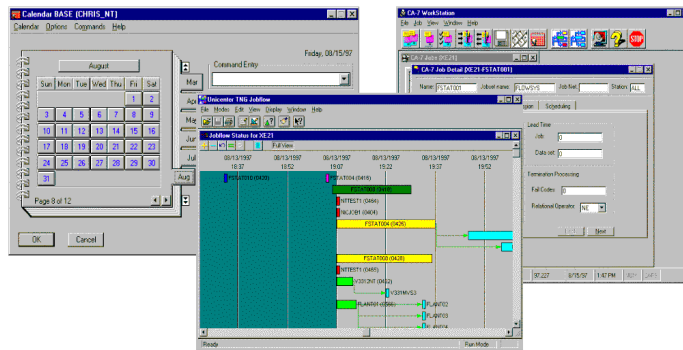
You can implement the CA NSM JM Option in a decentralized fashion so that there are multiple job management managers. Each of these can manage jobs on their own server and request processing on other servers running Unicenter Universal Job Management Agents (and, optionally, the CA NSM JM Option server).



## WorldView Implementation

Regardless of whether you use a centralized or decentralized method of implementation, an important element of any implementation is the ability to access and monitor all of the enterprise's workload from a single point. To accomplish this, use Unicenter WorldView and the scheduling WorkStations for the z/OS or OS/390 scheduling products:

- Unicenter CA-Scheduler WorkStation
- Unicenter CA-Jobtrac WorkStation
- Unicenter CA-7 WorkStation



## Recommended Approach

For procedures to implement cross-platform scheduling, see Implementing Cross-Platform Scheduling in the online *CA Procedures*.

## Scheduling Workstation Usage

Scheduling workstations provide a graphical tool for defining jobs and monitoring current workload processing on the OS/390 scheduling products. The scheduling workstations, when used with the CA NSM JM Option, provide the ability to administer and monitor all of your batch processing through WorldView from a Windows workstation.

## Windows Configuration Environment Variables

The environment variables in the following list are referenced in other various sections of this chapter. For a summary of all Windows CA NSM JM Option environment variables, see the online *CA Reference*.

### Default calendar

Specifies the name of the default CA NSM JM Option calendar. The default name is BASE. Set in the Configuration Settings dialog by clicking the Options tab and then the Job Management Options tab.

### **Autoscan hour and Interval between Autoscan**

Sets the time of day for the new-day autoscan and the interval between autoscan for the current workday. Set to different values to change the default time of the new-day autoscan (default is 0 for 00:00 or the default autoscan interval (default is every 3 hours). Set in the Configuration Settings dialog by clicking the Options tab and then the Job Management Options tab.

### **Post on Cancel?**

Controls the effect of the cancel command on successor jobs. The default setting is Y. When set to Y (yes), the CA NSM JM Option posts on cancel, which allows successor jobs to run because the predecessor requirement is satisfied by the predecessor having been removed (by the cancel command) from the current day's workload.

If the variable is set to N (no), successor jobs dependent on the prior completion of a canceled job are not posted. Because the predecessor requirement is not satisfied, successor jobs remain in a WPRED (waiting for predecessors) status.

If you change this variable, you must stop and restart the CA NSM JM Option using the following commands:

```
unicntrl stop sch
```

```
unicntrl start sch
```

**Note:** Running the unicntrl stop sch command stops job submission and disables display of the status tracking information for jobsets and jobs. These processes resume when you run the unicntrl start sch command. There is no loss of tracking data for jobs that complete or abnormally terminate while the CA NSM JM Option is stopped. The Job Event Logger continues to run and track completion of scheduled processes even while the CA NSM JM Option is stopped.

### **Cycle count precedence?**

Controls how the number of cyclic job iterations are calculated when cyclic jobs are demanded after their early start time.

### **Authorized remote manager nodes**

Specifies a comma-delimited list of authorized remote manager nodes allowed to submit jobs. The default (blank) is to allow all job managers to submit jobs to the job agent.

## UNIX/Linux Configuration Environment Variables

The environment variables in the following list are referenced in other sections of this chapter. For a summary of all UNIX/Linux CA NSM JM Option environment variables, see the online *CA Reference*.

### **CAISCHD0009**

Specifies the name of the default CA NSM JM Option calendar. The default name is base. Set this variable in the file `$CAIGLBL0000/sche/scripts/envset` file. For C shell users, set in the `$CAIGLBL0000/sche/scripts/envusrctsh` file.

### **CAISCHD0011 and CAISCHD0012**

Sets the time of day for the new-day autoscan and the interval between autoscans for the current workday. You can set these environment variables to different values if you want to change the default time of the new-day autoscan (default is 1 for 01:00 or the default autoscan interval (default is every 3 hours). Valid values for CAISCHD0012 are 0, 1, 2, 3, 4, 6, 8 and 12.

### **CAISCHD0014**

Controls the effect of the cancel command on successor jobs. The default setting is Y. When set to Y (yes), the CA NSM JM Option posts on cancel, which allows successor jobs to run because the predecessor requirement is satisfied by the predecessor having been removed (by the cancel command) from the current day's workload.

If the variable is set to N (no), successor jobs dependent on the prior completion of a canceled job are not posted. Because the predecessor requirement is not satisfied, successor jobs remain in a WPRED (waiting for predecessors) status.

Set this variable in the `$CAIGLBL0000/sche/scripts/envset` file by adding the following statements to `$CAIGLBL0000/sche/scripts/envset`:

```
CAISCHD0014=No
export CAISCHD0014
```

Then shut down and restart the CA NSM JM Option using these commands:

```
unishutdown sche
unistart sche
```

**Note:** Running the `unicntrl stop sche` or `unishutdown sche` command stops job submission and disables display of the status tracking information for jobsets and jobs. These processes resume when you run the `unicntrl start sche` or `unistart sche` command. There is no loss of tracking data for jobs that complete or abnormally terminate while the CA NSM JM Option is stopped. The Job Event Logger continues to run and track completion of scheduled processes, even while the CA NSM JM Option is stopped.

**CAISCHD0530**

Specifies a comma-delimited list of authorized remote manager nodes allowed to submit jobs. The default (blank) is to allow all job managers to submit jobs to the job agent.

**CAISCHD0540**

Controls how the number of cyclic job iterations are calculated when cyclic jobs are demanded after their early start time.

## Environment Variables for Jobs and Actions

When the CA NSM JM Option submits a job for execution as a Windows process, several environmental variables are set with information you may find useful. The programs or .BAT files executed as part of a job's submission can reference these variables, but the variables cannot be altered.

**JOBSET**

Specifies the name of the CA NSM JM Option jobset.

**JOBNAME**

Specifies the name of the job within the identified jobset.

**JOBJNO**

Specifies the job number associated with the identified jobset and job.

**JOBSTATION**

Specifies the CA NSM JM Option station to which the job was submitted.

**JOBQUAL**

Specifies the CA NSM JM Option job qualifier.

**JOBNODE**

Specifies the computer name (node ID) associated with the job.

**JOBNUMBER**

Specifies the unique job identifier assigned by the CA NSM JM Option for internal use.

**JOBSUBFILE**

Specifies the command string the CA NSM JM Option submitted to start this job.

**JOBUSER**

Specifies the user ID for which the job was submitted.

The following sample BAT file references settings within the following submitted job definition:

```
define jobset id=Trees  
  
define job id=(Trees,0ak,01) station=AMERICA autose1=yes  
  
define jobparm id=(Trees,0ak,01)subfile=0ak.bat
```

The OAK.BAT file contains the following:

```
@echo Off  
  
set ID=%JOBSTATION%: Job(%JOBSET%,%JOBNAME%,%JOBJNO%) Qual(%JOBQUAL%)  
  
cawto %ID% Phase 1 has ended  
  
cawto %ID% Phase 2 has ended
```

If job Oak is selected or demanded, the following messages are sent to the Event Console Log:

```
AMERICA: Job(Trees,0ak,01) Qual(3001) Phase 1 has ended  
  
AMERICA: Job(Trees,0ak,01) Qual(3001) Phase 2 has ended
```

## Monitor Workload Status

At completion of the new-day autoscan, the entire day's workload resides in the tracking file. (For a detailed explanation of the autoscan process, see Autoscan in this chapter.) For the remainder of the day, the CA NSM JM Option schedules, submits, and tracks the progress of that workload.

The CA NSM JM Option does much more than just submit jobs on time. It actually *monitors* or "tracks" the status of the workload and updates control information as jobsets and jobs progress from one status to another.

Using the tracking facilities of the CA NSM JM Option, you can monitor job status and, as necessary, modify the day's workload (modify the contents of the tracking file). Any changes you make in the Job and Jobset Tracking windows are temporary and affect only the contents of the tracking file; they do not affect the Job Management policy stored as job or jobset profiles in the Job Management Database.

**Note:** Although job and jobset resource requirements appear in the Tracking windows, this information is read-only and cannot be modified through the tracking facilities.

If you discover that you need to make changes that affect CA NSM JM Option policies, as opposed to minor changes in the current day's workload, see the procedures in *Modifying Job Management Policies* in the online *CA Procedures*.

For procedures to monitor workload status, see the following topics in the online *CA Procedures*:

- Displaying Jobset Status
- Changing the Status of Non-CPU Jobs
- Tracking the Time Spent on Manual Tasks
- Changing the Current Workload

## Jobflow Tracking on Windows

**Note:** The Jobflow GUI is available on Windows platforms only, but can be used to display Windows, UNIX/Linux, and z/OS or OS/390 CA NSM JM Option policy definitions and status.

To access the tracking facility of the CA NSM JM Option, click the Jobflow Forecast icon or the Jobflow Status icon from the CA NSM JM Option GUI.

By reviewing the sections that follow, you will learn to do the following tasks:

- Start the forecast and status monitor components and select jobs to view.
- Interpret the jobflow displays.
- Customize the number of successor/triggered jobs to appear in the displays.
- View job dependencies.
- Print jobflow displays and save them as files.

Jobflow provides an online, graphical representation of an enterprise's workload. You can view relationships between jobs that are defined in the scheduling database and monitor the status of jobs in real time. The graphical interface lets you tailor your view of the workload. You can limit the number of jobs to view by specifying job names, time span, and number of successor levels. You can display the results as either a Gantt or Pert chart. You can zoom in on a single job's successors and triggers, or zoom out for a larger perspective.

Jobflow provides two types of views:

- Jobflow forecast views
- Jobflow status views

*Jobflow forecast views* display job relationships (that is, predecessor-successor or job-triggers) for selected jobs within a given time frame. You can view dependencies (requirements) for any job in the flowchart.

*Jobflow status views* display real-time status information for selected jobs within a given time frame. Color coding lets you identify job status and pinpoint trouble spots. For example, abended jobs are red and late jobs are yellow. Status information is updated at regular intervals, so your view of the workload remains current.

Once you display a jobflow view, you can expand the scope of the display to include additional levels of successors (or triggered jobs), and to include job dependencies.

For procedures to view jobflows, see the following topics in the online *CA Procedures*:

- Starting the Jobflow Forecast Function
- Starting the Jobflow Status Function

The following sections present an overview of the jobflow forecast view and jobflow status view.

## Jobflow Forecast View

The jobflow forecast view lets you view relationships between jobs that are defined in the CA NSM JM Option database.

The Jobflow window contains the Jobflow Forecast window, which presents the jobflow forecast view of a single job; this job, the starting *job*, is the one you specify in your selection criteria.

The jobflow forecast presents each job as a bar extending across the appropriate number of time increments representing the job's start and end times. *Links* (or arrows) connect jobs that are related to the starting job. The *jobflow* consists of the starting job and the jobs that are successors to the starting job.

The jobflow forecast view can be either a Gantt chart or Pert chart. Gantt charts display the relationships between jobs across a time span; Pert charts display job relationships without regard to time. For information about Pert charts, see Changing the Chart Type: Gantt or Pert in the online help.

Jobs with names followed by plus signs can be expanded to show additional trigger (or successor) levels.



## Viewing Dependencies in the Jobflow Forecast Window

The Forecast window displays job predecessor-successor connections, but does not show job dependencies. To see these relationships and their status, use the Dependency View.

The Dependency View window lets you see two levels of dependency for a specific job. You can see the following:

- The job's predecessor and all of the job's dependencies (such as resources)
- All jobs that are successors of the specified job

In a single view you can see all of the jobs that affect the target job's execution, as well as all of the jobs that depend on the target job for their execution.

Further, you can display multiple dependency views.

For procedures to view dependencies, see *Viewing Dependencies in the online CA Procedures*.

## Viewing the Jobflow Status

The jobflow status view displays the real-time status of jobs within a given time frame.

The Jobflow window contains the Jobflow Status window, which presents several jobflows (a job and its successors). Unlike the jobflow forecast view, the jobflow status view can contain multiple jobflows, depending on your selection criteria.

Jobs are represented by bars extending across the appropriate number of time lines reflecting the starting and ending times of the jobs; the bars are color coded to indicate their status. Jobs that are currently executing span the vertical line representing the current time. Jobs that appear in the Future area are jobs that have not yet executed. Jobs in the Past area are jobs that have completed execution.

Job names appear inside or next to the job symbols.

When you first display the Jobflow Status window, the active jobs become the starting point of the flow. As the status is updated, a job remains in the display until the following three conditions are met:

- The job completes.
- All jobs associated with the job complete.
- The end time of the job precedes the start time of the display.

The default color codes for job status are as follows:

**Light blue**

Identifies jobs that are supposed to be processed in the near future based on what is defined in the database.

**Dark green**

Identifies active jobs.

**Medium green**

Identifies ready queue jobs.

**Light green**

Identifies request queue jobs.

**Dark blue**

Identifies completed jobs.

**Yellow**

Identifies late jobs.

**Red**

Identifies abnormal jobs.

**Pink**

Identifies canceled jobs.

**Note:** To change the default color scheme, see *Customize the Environment*.

## Displaying Multiple Views

You can display multiple jobflow forecast and jobflow status views simultaneously.

For procedures to display multiple views, see *Displaying Multiple Views* in the online *CA Procedures*.

## Customize the Environment

You can customize the jobflow display as well as the appearance of objects in the Jobflow Forecast and Jobflow Status windows.

## Operational Modes

Jobflow has two operational modes:

- Run mode
- Design mode

*Run mode* is the basic operating mode. When the system is in run mode, you can perform all of the Jobflow functions except changing the appearance of objects in the jobflow and saving jobflow forecast and status files.

*Design mode* lets you open the Tools window where you can change the appearance of the Jobflow environment and save jobflow forecast and status files.

The default mode at startup is run mode.

For procedures to customize your Jobflow environment, see the following topics in the online *CA Procedures*:

- Adjusting Job Name Labels
- Changing the Appearance of Objects
- Changing the Chart Type: Gantt or Pert
- Changing the Display Color
- Changing the Display Fonts
- Changing the Display Size
- Choosing Run Mode or Design Mode
- Expanding and Collapsing Jobflow Levels

## Scrolling Through the Forecast and Status Displays

You can scroll left and right to change the time frame displayed, or scroll up and down to display multiple jobflows (a job and its successors/triggered jobs) by using the navigation buttons or commands described in Navigating Through the Jobflow in the online *CA Procedures*.

You can tailor your view of the jobflow by expanding or collapsing the levels of successor jobs that appear. For example, you may want to see several levels of detail for a specific job but not for others. Job names that are followed by a plus sign designate jobs that can be expanded.

### Refreshing the Display

The display is automatically refreshed at regular intervals to update the Jobflow Status window. (You specified a refresh rate for the Jobflow Status window in the Status Selection Criteria dialog when you selected the workload to monitor.) You can control the updating of status information by refreshing the display at any time, stopping the refresh process, or changing the refresh rate.

### Printing a Jobflow Forecast

You can print the *current view* of the jobflow (that is, the contents of the active Jobflow Forecast window), or you can print the entire jobflow forecast. You can customize your output by specifying fonts, colors, titles, legends, and other elements.

The typical jobflow forecast spans multiple pages, both horizontally and vertically, and the output must be assembled. Jobflow generates a two-part number on each page of the printed output indicating the location of a particular segment on the assembled, printed flowchart. The digits preceding the decimal indicate the horizontal position (the row) of the segment; the digits following the decimal indicate the vertical position (the column) of the segment. The following diagram illustrates the numbering convention:

1.1	1.2	1.3
2.1	2.2	2.3
3.1	3.2	3.3

The sections in the remainder of this chapter provide more detail about the following topics:

- Customizing the output, specifying margins, titles, and reference information using the Page Setup dialog.
- Previewing the output before printing it.
- Adjusting the number of pages and the scale.

## Customizing and Previewing Your Printed Output

You can customize the printed output using the Page Setup dialog. The Page Setup dialog lets you customize the margins of the printed jobflow and to specify titles and reference information that appear on the printed output.

Settings for page margins, network margins, and borders apply to every segment of the printed output. Titles and legends can appear on every segment, or only on the segments that span the bottom of the printout, depending on the options you specify in the Chart Print Options dialog (for more information, see *Adjusting the Number of Pages and the Scale*).

You can preview the printed output using the Print Preview dialog. Select File, Print Preview to view a segment of the jobflow as it will appear in the printed output.

The Print Preview dialog shows all jobs that fall within the time frame specified in the Jobflow Forecast window, including jobs that do not appear on the screen (that is, more jobs may be active during a time period than can fit on a single screen). The view includes document titles, document reference information, page title, and legend, when specified.

## Adjusting the Number of Pages and the Scale

Typically, you will use Jobflow's automatic settings, which determine the number of horizontal and vertical pages required to most effectively accommodate the objects in the jobflow. You can override the automatic settings with your own numbers for horizontal and vertical pages. The numbers you choose will affect the appearance of the printed output.

Before printing, adjust the scale by zooming in or out. The number of pages required to print the jobflow depends on the time span you specified (the longer the time span covered, the greater the number of pages) and the scale of the objects in the jobflow display (the larger the magnification, the greater the number of pages).

**Note:** If you turned on the page title option in the Page Setup dialog, the page title (that is, the title bar text) appears at the bottom of every page segment, regardless of how you set the *Title on bottom pages only* option.

## Opening and Saving Jobflow Forecast Files

After you tailor your view of a jobflow in the Jobflow Forecast window (by expanding and collapsing the levels that appear and by modifying the scheduling database), you can "take a snapshot" of the jobflow by saving it as a jobflow forecast file with an extension of .GBF.

**Note:** This option is available only in Unicenter Classic.

When you open a flowchart forecast, the displayed workload is not connected to the underlying scheduling database; thus, you cannot expand and contract trigger levels or display dependencies. The advantage of working with a jobflow forecast file is that you can open the file quickly, because Jobflow does not have to build a selection list of jobs in the database. Working with a file, rather than a “dynamic” flowchart is useful if your schedule does not change frequently.

For procedures to open and save a Jobflow Forecast file, see the following topics in the online *CA Procedures*:

- Opening a Jobflow Forecast or Status File
- Saving a Jobflow Forecast or Status

# Index

---

## 2

### 2D Map

- background maps • 153
- billboards • 153
- custom views • 154
- favorite views • 155
- how navigation works • 156
- overview • 151, 152

## A

### access permissions

- asset perspective • 394
- user perspective • 394

### action menus in Alert Management System • 317

### active directory agent • 238

### Active Directory Enterprise Manager (ADEM) • 237

### Active Directory Management FIPS encryption about • 422

- converting password file to FIPS encryption • 423

- data encrypted • 422

- data encryption key • 422

- installation and migration • 422

### administrator ID for Unicenter NSM • 42

### administrator password, changing • 47

### Advanced Event Correlation

- about • 339, 340

- Boolean logic rules • 356

- Boolean rule pipeline items • 355

- correlation rules • 340, 351

- creating rules • 342

- credentials • 359

- deploying policy • 350

- event definitions • 340, 341

- global constants • 357

- impact analysis • 348

- implementing • 349

- individual root events • 354, 356

- introduction • 30

- pipeline items • 351

- regular expressions • 359

- starting the Integrated Development Environment(IDE) • 342, 343

- starting the web policy editor • 344

- template rules • 359

- timing parameters • 356

- user-defined tokens • 356, 357

- using tokens • 356, 357

### Advanced Event Correlation (AEC) • 304

### agent discovery • 131

### Agent Technology FIPS encryption

- about • 423

- data encrypted • 423

- data encryption key • 423

- installation considerations • 424

- migration considerations • 424

### Agent Technology overview • 221

### agents

- Agent View • 250

- auto watchers and available lists • 228

- call-back mechanism • 229

- cluster awareness • 230

- Distributed Services Bus • 247

- Distributed State Machine • 248

- generic resource monitoring • 231

- in Unicenter MP • 203

- maximum and minimum metrics • 232

- overloading thresholds • 233

- periodic configuration • 234

- Poll method • 234

- remote monitoring • 222

- resource monitoring • 227

- scoreboards and dashboards • 203, 204, 205

- SNMPv3 support • 235

- status deltas • 234

- watchers • 236

### Alert Management System • 29

- about alerts • 312

- action menus • 317

- AEC policies for alerts • 319

- alert console • 211

- alert scoreboard • 211

- classes • 315

- consolidate alerts • 317

- display attributes • 317

- escalate alerts • 317

- how AMS works • 313

- impact • 317

- in Unicenter MP • 209

---

- priority • 317
- queues • 316
- Service Desk integration • 321
- urgency • 317
- user actions • 212, 317
- user data • 317
- viewing alerts in MCC • 320

#### architecture

- remote monitoring • 223
- systems management • 244

asset types • 290

association browser • 169

## B

backlog scheduling • 519

Base calendar • 502, 532

billboards • 153

Boolean logic

- using in AEC rules • 356

bridge configuration • 442

bridge configuration files, creating • 446, 450

bridge control • 442, 443

bridge instance • 442, 444, 449

bridgecfg command • 442

bridgecntrl command • 443

bridging policy

- creating • 442

- implementing • 442

- overview • 437

browsers

- Agent View • 250

- DSM View • 251

- Event Browser • 251

- MIB Browser • 252

- Node View • 253

- Remote Ping • 253

- Repository Monitor • 254

- Service Control Manager • 254

- SNMP Administrator • 254

Business Process View Management

- about • 25, 180

- business process objects • 180

- integration with Event Management • 183

- rules • 180, 181, 182

Business Process Views

- dynamic • 159

- Dynamic Containment Service • 160

- overview • 157

- scoreboards • 199

- types • 158

## C

CA File Transfer (CAFT)

- binaries • 478

- components using CAFT • 477

- configuration files • 478

- definition • 476

- environmental variables • 478

CA Message Queuing Service (CAM)

- binaries • 478

- components using CAM • 477

- configuration files • 478

- configuring to use TLS encryption • 428

- optional ports • 473

- overview • 476

- required ports • 472

- transport layer protocols • 476

CA Secure Sockets Facility (CCISF)

- about • 81

- configuring • 84

- enabling • 82

- OpenSSL • 82

Ca7xxx log • 521

CA-CONVIEW asset type • 290

CA-CONVIEW-ACCESS asset type • 290

caevent command • 513, 514

CAICCI

- about • 81

- data encrypted • 424

- FIPS encryption • 424

- functions • 89, 90, 91

- installing with FIPS • 425

- required ports • 472

- turning on FIPS mode • 425

- user customization • 95

CAISCHD0009 environment variable • 502, 532

CAISCHD0011 environment variable • 532

CAISCHD0012 environment variable • 532

CAISCHD0014 environment variable • 506

calendars

- default • 502, 532

- expanded processing • 503

- profiles • 499

- scheduling by • 500

- undefined • 523

cancel command • 532

capagecl command • 308

catrap command • 298



---

- catrapd daemon • 292
- cauexpr command • 523
- cawrksec utility • 509
- ccicrypt utility • 87
- child update rule • 182
- CICS agent • 238
- Cisco device recognition • 435
- Cisco Integration • 435
- class editor • 166
- classes in Alert Management System • 315
- Classic Discovery
  - determine device names • 137
  - discover a single network • 136
  - effect of timeout values • 103
  - how subnets are used • 134
  - methods • 130
  - multi-homed device support • 102
  - preparing for • 135
- classifyrule.xml file • 106, 112
- CLEAN parameter • 523
- cleanlog shell script • 521
- CleverPath Portal • 190
- command messaging, implementing • 310
- Commands
  - bridgecfg • 442
  - bridgecntrl • 443
  - caevent • 513, 514
  - cancel • 532
  - cauexpr • 523
  - cautil • 392
  - cawto • 283
  - dscvrbe • 101
  - modp • 50, 51
  - schdchk • 518
  - schdfore • 518
  - schdhist • 518
  - schdsimu • 518
  - secadmin • 381
  - unicntrl • 530
  - unishutdown • 532
  - unistart • 530, 532
  - whathas • 402
  - whohas • 400
  - wmmodel • 518
- Common Discovery
  - common discovery gui • 141
  - configuring discovery agent • 145, 146
  - configuring discovery server • 142, 143, 144
  - discovery agent • 139
  - discovery request client • 140
  - discovery server • 139
  - discovery web client • 140
  - import service • 147
  - IPv6 import tool • 175
  - overview • 138, 139
- communication protocol security
  - about • 79
  - encryption levels • 79
- configuration and diagnostics • 329
- Configuration Manager
  - about • 26, 264
  - base profile • 265, 266
  - configuration bundle • 270, 271
  - delivery schedule • 269, 270
  - differential profile • 267, 268
  - file package • 268, 269
  - reports • 273
  - resource model groups • 264
- configuration, agent/server
  - configuration, agent/server, multiple node • 525
  - configuration, agent/server, single node • 525
- connect remotely to another MDB using WorldView Classic GUI • 53
- console log • 292
- console views • 290
- Continuous Discovery
  - and CAFT • 477
  - and CAM • 477
  - behind firewalls • 127
  - create unmanaged devices • 118
  - default configuration • 117
  - Discovery agents • 115, 123
  - Discovery Manager • 115, 121
  - exclude classes • 119
  - how it works • 116, 117
  - overview • 115
  - set up SNMP community strings • 120
- Correlation
  - creating rules • 342
- correlation rules, Advanced Event Correlation • 340
- CPU station • 500
- cross-platform scheduling
  - description • 526
  - implementation • 528
  - manager/agent model • 527

---

---

- manager/agent model configuration • 528
- cross-platform scheduling implementation
  - centralized • 529
  - decentralized • 529
  - overview • 528
  - worldview • 530
- cyclic job submission • 511

## D

- daily cubes • 367
- dashboards • 191
- Data Scoping
  - about • 62
  - activating • 75
  - deactivating • 76
  - implementing • 78
  - in the 2D map • 75
  - rule editor • 77
  - rule evaluation • 71, 73
  - rule inheritance • 65
  - rule performance issues • 68
  - rules • 63, 64
  - security • 69
  - user IDs for evaluation on Windows (Ingres)
    - 70
  - user IDs for evaluation on Windows (Microsoft SQL Server) • 70
- data transport mechanisms • 471
- Define Logical Repository utility and VNODEs • 43
  - connect remotely to another MDB • 53
  - using • 55
- demand scheduling • 516
- description of • 22
- Desktop Management Interface (DMI)
  - components • 453
  - DMI agent • 455
  - DMI browser • 454
  - DMI manager • 455
  - overview • 453
  - service provider • 453, 454
  - set trap destinations in the DMI agent • 456
  - Unicenter support for DMI • 455
- destination repositories, multiple • 439
- destination repository, single • 440
- DIA communications port, configuring • 474
- DIA, required ports • 472
- Discovery
  - about • 99

- classic interface • 25
- classification engine • 102
- classification rules • 106
- combining continuous and classic • 101
- Common • 138, 139
- configuration files • 104, 106, 112
- Continuous • 115
- create users with administrator privileges (Ingres) • 50
- create users with administrator privileges (SQL Server) • 49
- create users without administrator privileges (Ingres) • 51
- create users without administrator privileges (SQL Server) • 50
- events reported to the Event Console • 123
- IPv6 devices • 138, 147
- IPv6 import tool • 175
- object creation rules • 104
- subnet filters • 103
- timestamp • 102
- types of methods • 104
- display attributes in Alert Management System • 317
- Distributed Intelligence Architecture (DIA) • 20, 80
- DSM (Distributed State Machine)
  - configuration • 258
  - discover resources • 249
  - interfaces • 262
  - manage object properties • 251
  - monitoring resources • 221
  - overview • 248
- Dynamic Business Process Views • 159
- Dynamic Containment Service (DCS) • 160

## E

- early start time • 500, 507, 520
- eHealth integration
  - about • 34, 217
  - how it works in Unicenter MP • 218
- encrypting • 81
- encryption levels • 79
- encryption See also CAICCI • 81
- encryption utility for use with CCISF • 87
- encryption, FIPS compliance • 409
- enterprise cubes • 367
- enum\_cert\_info function • 92
- environment variables • 287, 310

---

ETPKI • 409  
eTrust Access Control  
    integrating with • 59, 60  
    not migrated • 61, 62  
event agent  
    configuring • 279  
    implementing • 278  
Event agent • 277  
event console • 208, 289  
Event correlation • 289, 304  
Event Management  
    about • 28, 275  
    Discovery events sent to the Event Console •  
        123  
    filters • 209  
    in Unicenter MP • 206  
    integration with Business Process View  
        Management • 183  
    maintenance considerations • 469  
    scoreboard • 207  
    testing policy • 283  
    virus scan utility • 469  
events • 276  
    actions • 208  
    impact • 184  
    notification • 183

## F

fanout repository architecture • 439  
file close event • 513  
file privileges, changing • 48  
FIPS 140-2 compliance  
    about • 409  
    Active Directory Management • 422  
    Agent Technology • 423  
    CCI • 424  
    compliant components • 409  
    data encrypted • 410, 422, 423, 424, 426,  
        429, 430  
    data encryption key • 412, 422, 423, 427,  
        430, 431  
    installation considerations • 413, 422, 424,  
        425, 427, 430, 431  
    Management Command Center • 426  
    Systems Performance • 410  
    Unicenter Management Portal • 429  
    Web Reporting Server • 430

## H

Historical Performance Agent • 366, 367  
hpaAgent • 366, 367

## I

impact events • 184  
impact in Alert Management System • 317  
Industry Standard MIBs • 300  
Ingres  
    remote connections to the MDB • 43  
    user groups • 40  
    users • 41  
Intel Active Management • 34  
International Standards Organization (ISO) •  
    300  
Internet Assigned Numbers Authority (IANA) •  
    300  
IPv6 discovery • 138, 139, 147, 175  
IPX Discovery • 132

## J

Job Management Agent  
    agent/server configuration • 524  
    configuration • 528  
    functions • 498, 527  
    remote submission agent • 498  
    remote tracking agent • 498  
job management autoscan  
    cleanup and backlogging • 519  
    definition • 518  
    new-day • 518, 522, 532  
    qualifications for • 519  
    undefined calendars during • 523  
    workload processing • 520  
Job Management calendar profile  
    calendar processing • 503  
    definition • 499  
    event-based scheduling • 500  
    job scheduling • 500  
    profile notebook page • 503  
    undefined calendars • 523  
Job Management environment variables  
    authorized remote manager nodes • 530  
    autoscan hour • 530  
    CAISCHD0009 • 502, 532  
    CAISCHD0011 • 532  
    CAISCHD0012 • 532  
    CAISCHD0014 • 532

- 
- cycle count precedence? • 530
  - default calendar • 530
  - for jobs and actions • 533
  - interval between autoscans • 530
  - post on cancel? • 530
  - Job Management maintenance
    - CLEAN parameter • 523
    - database definitions, change to text file • 523
    - log files • 521
    - purge old history records • 523
    - submit jobs for another user • 524
    - tracking file • 522
    - undefined calendars • 523
  - Job Management Option
    - agent/server configurations • 524
    - agents • 497, 498
    - autoscan • 518
    - CPU station • 500
    - event-based scheduling • 500
    - job demand option • 516
    - Job Management manager • 527
    - job server • 497, 498
    - jobflow • 497
    - jobsets • 504
    - log files • 521
    - maintenance • 521
    - multiple hosts • 525
    - overview • 497
    - POSTCPU station • 500
    - PRECPU station • 500
    - predictive scheduling • 500
    - profiles • 499
    - simulation report • 517
    - single server • 525
    - trigger profiles • 513
    - variables • 500
    - workload balancing • 501
  - Job Management Option tasks
    - clear undefined calendars • 523
    - form groups of related tasks (jobsets) • 504
    - identify resource requirements • 501
    - identify work to perform • 507
    - maintenance • 522
    - monitor workload status • 534
    - run a job on demand • 516
    - run simulation reports • 517
    - schedule work by dates • 502
    - schedule work by special events • 513
    - specify where to perform work • 500
  - Job Management option variables
    - early start time • 500
    - maximum time • 500
    - must complete time • 500
    - must start time • 500
  - Job Management predecessor profiles • 499
  - Job Management resource profiles • 499, 501
  - Job Management shell scripts
    - cleanlog • 521
    - schdcheck • 518
    - schdfore • 518
    - schdhist • 518, 523
    - schdpxr • 518
    - schdsimu • 518
    - wmmodel • 518
  - Job Management station group profiles • 499, 500
  - Job Management station profiles • 499, 500
  - job management triggers
    - as predecessors • 515
    - caevent • 514
    - profile • 513
  - job scheduling
    - by dates • 502
    - by special events • 513
    - cross-platform scheduling • 526
    - cross-platform scheduling implementation • 528
    - expanded calendar processing • 503
    - scheduling workstation usage • 530
    - types • 502
  - job server
    - components • 498
    - functions • 498
  - jobflow
    - .GBF files • 541
    - customizing • 538
    - design mode • 539
    - forecast files • 541
    - forecast views • 535
    - jobflow GUI • 535
    - levels, expanding/collapsing • 539
    - multiple views • 538
    - navigating • 539
    - page setup dialog • 541
    - refreshing • 540
    - run mode • 539
    - status color codes • 537
    - status views • 535
-

---

- view status • 537
- jobflow forecast views
  - definition • 535
  - example • 536
  - expanding/collapsing • 539
  - multiple views • 538
  - opening • 541
  - overview • 536
  - printing • 540
  - saving • 541
  - viewing dependencies • 537
- jobflow printing
  - adjusting pages/scale • 541
  - customizing output • 541
  - forecast views • 540
  - page setup for • 541
- jobflow status view
  - multiple views • 538
- jobs
  - autoscan • 518
  - cyclic • 511
  - DYNAMIC type • 516
  - early start time • 507
  - EXTERNAL setting • 508
  - in jobsets • 507
  - job profiles • 499
  - on demand • 516
  - overview • 507
  - password validation • 509
  - predecessors • 509
  - resource requirements • 501
  - resources • 508
  - scheduling types • 500
  - submission characteristics • 508
  - workload processing • 520
- jobset predecessors
  - canceled • 506
  - nonexistent predecessor evaluation • 506
- jobset resources
  - resource amount • 504
  - resource usage • 504
  - resource weight • 504
- jobsets
  - definition • 504
  - jobset profile • 504
  - jobsets, on demand • 516
  - predecessors • 505
  - resources • 504
  - workload processing • 520

## L

- link browser • 170
- log agent • 240

## M

- maintenance
  - Event Management • 469
- Management Command Center (Unicenter MCC)
  - 21
- Management Database (MDB) also see MDB • 19
- Management Information Base (MIB) • 299, 300
- manager layer browsers
  - Agent View • 250
  - DSM View • 251
  - Event Browser • 251
  - MIB Browser • 252
  - Node View • 253
  - overview • 250
  - Remote Ping • 253
  - Repository Monitor • 254
  - Service Control Manager • 254
  - SNMP Administrator • 254
- master catalog
  - description of • 22
- Maximum time • 500
- MDB
  - about • 19
  - database support • 19
  - managed objects • 150
  - securing • 39
  - users • 40, 42
- MDB server
  - operating system user • 42
  - remote connections • 43, 53
- mdbadmin Ingres user group • 40
- message actions • 282, 283, 284, 303
  - scheduling by • 500, 513
- message records • 282, 303
- messages • 280, 281, 284, 288
- metadata • 374
- methods.xml file • 106
- modp command • 50, 51
- MOM Management
  - about • 33
  - alert severities • 459
  - how it works • 458
  - resolution states • 460
  - status in WorldView • 460

---

Monitor in Job server • 498  
monitoring  
    enterprise • 221  
    remote • 222  
    resources • 227  
mtrsf log • 521  
multi-homed devices • 102  
multiple node configuration in Job Management  
    • 525  
Must complete time • 500  
Must start time • 500

## N

network topology • 151  
New-day autoscan • 534  
Nodes  
    multiple • 525  
    single • 525  
    UNIX/Linux node support • 387  
Non-CPU tasks • 500  
non-root Event Agent • 277  
notification events • 183  
notifications • 213

## O

objects  
    Discovery creation rules for • 104  
    importance • 161, 162, 163  
    managed • 150  
    setting policy using alarmsets • 164  
    severity levels • 161  
    viewing properties • 166  
ObjectView  
    customize chart • 168  
    dashboard monitor • 168  
    graph wizard • 168  
    overview • 167  
on demand scheduling • 516  
OpenSSL • 82

## P

Page Writer • 310  
PageNet® • 310  
PAT (protected asset table) • 385  
performance architecture • 368  
Performance Chargeback • 366  
Performance Configuration • 375  
Performance Data Grid (PDG) • 370  
Performance Distribution Servers • 370, 372

Performance Domain Servers • 370, 371  
Performance Reporting • 365  
Performance Trend • 365  
period cubes • 367  
policies • 276, 302  
policy packs • 302, 303, 304  
ports  
    about • 471  
    optional • 473  
    required • 472  
POSTCPU station • 500  
PRECPU station • 500  
Predecessors  
    evaluated • 520  
    for jobs • 509  
    for jobsets • 505  
    profiles • 499  
    triggers as • 515  
Predictive scheduling • 500  
prfAgent • 366  
priority in Alert Management System • 317  
priority of jobs • 520  
Profile Editor • 375  
Profiles • 513  
propagation thresholds rule • 182  
protected asset table (PAT) • 385  
pseudo-MIBs • 300  
Purging history records • 523

## Q

queues in Alert Management System • 316

## R

Real-Time Performance Agent • 366  
Remote Monitoring  
    about • 27  
    advantage and disadvantage • 222  
    architecture • 223  
    components • 223  
    resource types • 225  
    role-based security • 227  
remote repository, connecting • 54  
Reports  
    Configuration Manager • 273  
    creating • 30  
    in Unicenter MP • 214  
    Job Management • 518  
    report types • 377  
    Security Management • 399

---

- Simulation • 517
- templates • 378
- Repository Bridge
  - and notifications • 438
  - architecture types • 439
  - bridge configuration GUI • 448
  - bridged and non-bridged objects • 437
  - bridging rules • 448
  - components • 442
  - initialization • 438
  - on UNIX/Linux • 437
  - overview • 437
  - rule file parameters • 451
  - supported platforms • 444
  - troubleshooting • 445
  - view log files • 446
- Repository Bridge architecture
  - aggregation architecture • 440
  - duplicate objects • 440
  - fanout architecture • 439
  - which to use • 441
- Repository Bridge components
  - bridge configuration • 442
  - bridge control • 443
  - bridge instances • 444
  - creating a bridge configuration file • 446, 450
  - managing bridge instances • 449
- Repository Bridge uses
  - for problem notification • 445
  - in a distributed organization • 444
  - restricted view of resources • 445
- repository import export utility • 173, 174
- resource monitoring
  - call-back mechanism • 229
  - configuring auto discovery • 230
  - evaluation policy • 231
  - object instances • 234
  - overview • 227
  - polling method • 234
  - selection list • 234
  - support for SNMPv3 • 235
  - using Agent Technology • 221
  - using metrics • 232
  - watchers • 235
  - write-back periodic configuration • 234
- resource monitoring functions • 228
- resource monitoring system agents
  - Active Directory Services agent • 238
  - CICS agent • 238

- Log agent • 240
- Script agent • 240
- SystemEDGE agent • 241
- UNIX/Linux agent • 241
- Windows agent • 242
- Windows Management Instrumentation agent • 242
- z/OS agent • 243
- Resources for Job Management
  - evaluated • 520
  - for jobs • 508
  - for jobsets • 504
  - profiles • 499, 501
  - requirements • 501
- rules • 180

## S

- SAN Discovery • 133
- schdcheck shell script • 518
- schdchk command • 518
- schdfore command • 518
- schdfore shell script • 518
- schdhist command • 518
- schdhist shell script • 518, 523
- schdlcs log • 521
- schdpxr shell script • 518
- schdsimu command • 518
- schdsimu shell script • 518
- Scheduling work across platforms • 526
- SCOM management
  - how the integration works • 464
  - NSM integration with • 461
  - SCOM alerts • 465
  - SCOM entities status in WorldView • 466
  - software requirements • 462
  - terminology • 463
- scoreboards and dashboards • 191, 192
- script agent • 240
- Secure Sockets Facility (SSF)
  - ccicrypt utility • 87
  - compability with previous CAICCI versions • 83
  - encryption utility • 87
  - enum\_cert\_info function • 92
- Security Management
  - creating rules in WARN mode • 389
  - daemons, starting • 388
  - deactivate • 399
  - FAIL mode • 384

- 
- file access authorization • 46
  - functions • 380
  - implementation phases • 382
  - MONITOR mode • 384
  - node support implementation • 387
  - QUIET mode • 384, 386
  - remote CAISSF return codes • 386
  - setting options in FAIL mode • 398
  - WARN mode • 384, 389
  - Security Management access
    - access determination • 395
    - access modes • 393
    - access permissions • 394
    - access rule example • 394
    - access types • 392
    - asset permissions • 392
    - CAISSF scoping • 396
    - CONTROL mode • 393
    - CREATE mode • 393
    - date and time controls • 393
    - DELETE mode • 393
    - DENY type • 392
    - LOG type • 392
    - PERMIT type • 392
    - READ mode • 393
    - rule evaluation • 395
    - UPDATE mode • 393
    - violations • 399
    - WRITE mode • 393
  - Security Management asset groups
    - adding new • 392
    - defining • 391
    - description • 381
    - groups within groups • 391
    - nested • 391
  - Security Management assets
    - asset definitions • 387
    - asset names • 391
    - asset permissions • 392
    - asset types • 391
    - description • 381
  - Security Management commit process
    - description • 381
    - executing for production, FAIL mode • 398
    - in FAIL mode • 398
    - in WARN mode • 398
    - using the GUI • 398
  - Security Management options
    - authorized user list • 386, 388
    - automatic startup • 387
    - CAISSF command scope (Windows) • 397
    - CAISSF data scope • 397
    - CAISSF keyword scope • 396
    - CAISSF scoping • 396
    - command scoping options • 396
    - customizing • 383
    - default permission • 384
    - DEFAULT\_PERMISSION • 384
    - implicit DENY permission for rules • 392
    - rule server support • 386
    - security management options, UNIX/Linux
      - location • 383
    - setting to absolute values • 388
    - SSF\_AUTH • 386, 388
    - system violation mode • 384, 395
    - SYSTEM\_MODE • 384
    - USE\_PAT • 385
    - user group server report • 386
  - Security Management reports
    - access violations • 399
    - overview • 399
    - what-has reports • 402
    - whohas reports • 400
  - Security Management user groups • 381
    - defining • 389
    - description • 381
    - nested • 390
    - server • 386
  - security policies • 381
    - testing during implementation • 388
  - Service Desk integration in Alert Management System • 321
    - about the integration • 321, 322
    - integrating with Unicenter MP • 216
  - set up read-only users for WorldView • 53
  - severity level of an object • 161, 164
  - severity propagation service • 164, 165
  - Simulation Report for Job Management • 517
  - single node configuration for Job Management • 525
  - SmartBPV
    - about • 26, 184
    - benefits • 185
    - create Business Process Views • 184
    - examples • 186
    - how it works • 185
    - integration with Event Management • 183
  - SNMP traps • 292, 293, 296, 297, 298, 299
-



- 
- source repositories, multiple • 440
  - source repositories, single • 439
  - special event scheduling in Job Management • 513
  - SPECTRUM
    - in Unicenter MP • 219
    - integrating with • 35, 485
  - SPECTRUM integration kit
    - about • 485
  - SPO-Performance Scope • 364
  - SSF support
    - return codes • 386
  - state count rule • 181
  - station groups profile • 499
  - stations
    - profiles • 499
    - types • 500
  - submit process as another user • 524
  - Submitting jobs • 508
  - summary cubes • 373
  - support for version 3 traps • 293
  - system monitoring for z/OS • 32
  - SystemEDGE agent • 241
  - systems management
    - architecture • 244
    - DSM discovery process • 249
    - managed objects • 245
    - manager layer • 247
    - monitoring layer • 244
    - overview • 243
    - states • 246
    - threshold breach process • 246
    - WorldView layer • 250
  - systems management configuration sets
    - adaptive configuration • 256
    - benefits • 255
    - configuration file location • 256
    - distribute configset • 257
    - load configset • 257
    - write configset using mkconfig utility • 256
  - systems management DSM configuration
    - Agent Class Scoping • 258
    - Discovery Community Strings • 259
    - Discovery Pollset Values • 259
    - DSM Wizard • 261
    - IP Address Scoping • 259
    - Managed Object Scoping • 260
    - overview • 258
  - systems management DSM Monitor
    - DSM Monitor Dashboard • 263
    - DSM Monitor View • 262
    - DSM Node View • 263
    - using DSM • 261
  - systems management, manager layer components
    - Distributed Services Bus • 247
    - DSM • 248
    - DSM Monitor • 249
    - DSM Store • 249
    - Object Store • 248
    - Service Control Manager • 247
    - SNMP/DIA Gateways • 248
    - Trap multiplexer • 248
    - WorldView Gateway • 249
  - Systems Performance FIPS encryption
    - about • 410
    - adding the Performance Agent • 418
    - changing encryption key • 416
    - data encrypted • 410
    - data encryption key • 412
    - generate a new encryption key • 420
    - installation considerations • 413
    - migration considerations • 418
    - reencrypting data • 421
    - switching off FIPS mode • 417
    - switching to FIPS mode • 415
    - turning on • 412
    - updating user domain access file • 419
  - Systems Performance, introduction to • 29, 363
- ## T
- Technical Support • iv
  - time controls
    - user access • 393
    - workload • 502, 507, 511
  - timestamp for discovered devices • 102
  - tracking file in Job Management • 518, 522
  - Trap daemon • 292
  - Trap Daemon • 433
  - trap destinations • 296
  - trap filters • 434
  - trap formatting • 297
  - Trap Manager
    - about • 31
    - local versus remote installation • 434
  - TRAP tables • 299
  - trigger profiles
    - as predecessors • 515
-

---

- scheduling by • 500, 513
- types of • 514
- trix • 171, 172, 173, 174
- tuples • 385

## U

- uniadmin Ingres user group • 40
- Unicenter classic interface • 24
- Unicenter MCC
  - about • 21
  - controlling access to • 56
  - deactivating password caching • 59
  - FIPS encryption • 426
  - overriding user ID • 58
  - starting • 22
- Unicenter MCC FIPS encryption
  - about • 426
  - configuring CAM to use TLS encryption • 428
  - data encrypted • 426
  - data encryption key • 427
  - installation considerations • 427
  - migration considerations • 428
  - turning off password caching • 428
- Unicenter MP
  - about • 189
  - administration • 193, 194
  - components • 198
  - filters • 209
  - FIPS encryption • 429
  - managing component connections • 195
  - portal explorer • 201
  - reports • 214
  - scoreboards and dashboards • 191, 192
  - severity browser • 202
  - users • 191
  - workplaces • 196
- Unicenter MP components
  - about • 198
  - agent dashboards • 205
  - agent management • 203
  - agent scoreboards • 203, 204
  - alert actions • 212
  - Alert Management • 209, 211
  - alert scoreboard • 211
  - eHealth • 217, 218
  - event actions • 208
  - Event Management • 206, 208
  - event scoreboard • 207
  - notifications • 213
  - other components • 220
  - SPECTRUM • 219
  - Unicenter Service Desk • 216
  - Unicenter Service Metric Analysis • 215
  - WorldView • 198
  - WorldView scoreboards • 199, 200
- Unicenter MP FIPS encryption
  - about • 429
  - data encrypted • 429
  - data encryption key • 430
  - installation considerations • 430
- Unicenter Notification Services
  - about • 323
  - configuration and diagnostics • 329
  - email protocol • 330
  - features • 325
  - how it works • 324
  - instant message protocol • 331
  - page snpp protocol • 331
  - page tap protocol • 332
  - protocols • 329
  - script protocol • 338
  - short message protocol • 335
  - voice protocol • 336
  - wireless protocol • 331
- Unicenter NSM
  - about • 17
  - database support • 19
  - for UNIX and Linux • 18, 403
- Unicenter NSM security
  - administrator ID • 42
  - changing administrator password • 47
  - changing file privileges • 48
  - changing severity propagation password • 47
  - create Unicenter NSM administrators (Ingres) • 43
  - create Unicenter NSM administrators (Microsoft SQL Server) • 42
  - embedded • 44
  - MDB security • 39
  - Microsoft SQL Server roles • 40
  - role-based • 39
  - roles (user groups) • 44
  - security rules • 44
  - Windows Administrators and Power Users groups • 47
- Unicenter Registration Services • 177

---

Unicenter Service Desk. See Service Desk  
Integration in Alert Management System •  
321

Unicenter Service Metric Analysis, working with  
• 215

unicntrl command • 530

unishutdown command • 532

unistart command • 532

uniuser Ingres user group • 40

UNIX system agent • 241

UNIX/Linux support  
about • 18, 403  
database support • 19  
supported components • 404, 406

unmanaged devices, create with Continuous  
Discovery • 118

urgency in Alert Management System • 317

user actions in Alert Management System • 317

user data in Alert Management System • 317

user groups  
defining • 389  
groups within groups • 390  
nested • 390

user interfaces • 23

using Agent Technology • 221

## V

Virtual Node Names (VNODEs) • 43

virus scanning utility • 469

Vista, running utilities on • 49

VNODEs, Ingres • 43

## W

Web Reporting Server FIPS encryption  
about • 430  
data encrypted • 430  
data encryption key • 431  
installation considerations • 431

weighted severity • 162

whathas command • 402

whohas command • 400

Windows Administrators and Power Users  
groups • 47

Windows system agent • 242

Wireless Messaging • 306  
capagecl • 308  
command messaging • 310  
configuration files • 309  
environment variables • 310

message files • 308

Policy Writer GUI • 311

Reply Information Configuration file • 309

reply text format • 308

template files • 311

WMI agent • 242

wmmodel command • 517, 518

wmmodel shell script • 518

workload balancing in Job Management • 501

WorldView  
2D Map • 151, 152  
about • 25  
alarmsets • 164  
association browser • 169  
billboards • 153  
business process views • 157  
class editor • 166  
Common Discovery Import service • 147  
components • 149  
custom views • 154  
importance of an object • 161, 162  
importing and exporting objects • 171, 172,  
173, 174  
in Unicenter MP • 198  
IPv6 import tool • 175  
IPv6 topology • 147  
link browser • 170  
managed objects • 150  
ObjectView • 167  
scoreboards • 199, 200  
set up read-only users (Ingres) • 53  
set up read-only users (Microsoft SQL  
Servers) • 52  
set up read-only windows-authenticated  
users • 52  
severity levels • 161  
severity propagation service • 164  
WorldView classic • 24

wvadmin Ingres user group • 40

wvuser Ingres user group • 40

## Z

z/OS monitoring • 32

z/OS system agent • 32, 243