

# CA Virtual Assurance for Infrastructure Managers

## Administration Guide

Release 12.6



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

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

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

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

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

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

The manufacturer of this Documentation is CA.

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

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

## CA Product References

This document references the following CA products:

- CA eHealth®
- CA Embedded Entitlements Manager (CA EEM)
- CA IT Asset Manager
- CA IT Client Manager
- CA Network and Systems Management (CA NSM)
- CA Patch Manager
- CA Server Automation
- CA Service Desk Manager (CA SDM)
- CA Spectrum®
- CA SystemEDGE
- CA Virtual Assurance for Infrastructure Managers
- CA Software Delivery

# Contact CA

## Contact CA 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

## Providing Feedback About Product Documentation

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 available on the CA Support website at <http://ca.com/docs>.

## Documentation Changes

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

- Active Directory and Exchange Server AIM Chapter—The name of the AIM has been changed to Active Directory and Exchange Server AIM (ADES formerly ESAD). Some minor changes have been made to the content.



# Contents

---

## Chapter 1: Introduction 15

|                           |    |
|---------------------------|----|
| Audience .....            | 15 |
| Related Publications..... | 15 |
| Conventions .....         | 16 |

## Chapter 2: Overview 19

|                                |    |
|--------------------------------|----|
| Architecture .....             | 19 |
| Databases .....                | 22 |
| Management DB .....            | 22 |
| Performance DB .....           | 22 |
| User Interface.....            | 23 |
| Access the User Interface..... | 24 |

## Chapter 3: Managing Users 25

|                                                                                      |    |
|--------------------------------------------------------------------------------------|----|
| User Access Control.....                                                             | 25 |
| Active Directory.....                                                                | 25 |
| Native Security .....                                                                | 26 |
| Password Management.....                                                             | 29 |
| Change the CA Embedded Entitlements Manager Administrator Password (EiamAdmin) ..... | 29 |
| Change the Database Administrator (sa) Password .....                                | 30 |
| Change the System User Password for Native Security .....                            | 31 |
| Change the System User Password for Active Directory Security.....                   | 32 |
| Tasks for External Directory Users .....                                             | 33 |
| Create User Groups.....                                                              | 33 |
| Assign Users to Groups .....                                                         | 34 |
| Remove Users or User Groups from a User Group .....                                  | 34 |
| Assign User Groups Access Rights to Services .....                                   | 35 |
| Set User Group Permissions.....                                                      | 36 |
| Delete User Groups.....                                                              | 36 |
| Search for Users or User Groups.....                                                 | 36 |
| User-scoped Authentication for vCenter Server .....                                  | 37 |

## Chapter 4: Configuring Resources 41

|                                                            |    |
|------------------------------------------------------------|----|
| Discovery .....                                            | 41 |
| Configuration Requirements for Managing Environments ..... | 41 |

---

|                                                             |    |
|-------------------------------------------------------------|----|
| Discover Environments .....                                 | 42 |
| Manage Multiple vCenter AIMS .....                          | 49 |
| Manage a Renamed vCenter Server .....                       | 51 |
| Requirements for Solaris Zones Management.....              | 52 |
| Requirements for LPAR Management .....                      | 52 |
| Configure SNMPv1 Traps by Editing the sysedge.cf File ..... | 52 |
| Event Forwarding .....                                      | 54 |
| Configure Windows for SNMP .....                            | 54 |
| Configure CA Virtual Assurance to Forward Events .....      | 55 |
| SNMP V3 Engine ID .....                                     | 56 |

## Chapter 5: Managing Resources 57

|                                          |     |
|------------------------------------------|-----|
| Manager Components.....                  | 57  |
| Application Insight Modules (AIMs) ..... | 57  |
| Monitor Configuration .....              | 59  |
| Discover Resources .....                 | 59  |
| Discover a System .....                  | 59  |
| Discover a Network .....                 | 60  |
| Delete a Network .....                   | 61  |
| Cancel Network Discovery .....           | 62  |
| Manage a Resource.....                   | 62  |
| Create a Service.....                    | 62  |
| Edit a Service .....                     | 64  |
| Managed Virtual Environments .....       | 66  |
| VMware vCenter Server .....              | 67  |
| Solaris Zones .....                      | 82  |
| Microsoft Hyper-V Server .....           | 84  |
| IBM PowerVM (LPAR).....                  | 86  |
| Microsoft Cluster Service .....          | 90  |
| Cisco UCS.....                           | 93  |
| Cisco UCS Overview.....                  | 95  |
| Add a Cisco UCS Manager Server .....     | 96  |
| Register a UCS AIM Server .....          | 96  |
| Remove a UCS Server .....                | 97  |
| Configure the SNMP Data Poller .....     | 97  |
| Configure the Service Poller .....       | 98  |
| UCS Trap Management .....                | 98  |
| Service Profiles .....                   | 99  |
| Using Multiple Multi-instance AIMS.....  | 101 |
| Systems Management.....                  | 101 |
| SystemEDGE Features .....                | 103 |



---

|                                   |     |
|-----------------------------------|-----|
| Service Response Monitoring ..... | 113 |
| Remote Monitoring.....            | 117 |
| Agent Configuration .....         | 118 |
| Agent Visualization.....          | 119 |
| Security and Maintenance .....    | 122 |

## Chapter 6: Managing Policies 125

|                                                                             |     |
|-----------------------------------------------------------------------------|-----|
| Rules and Actions .....                                                     | 125 |
| Configure CA SDM .....                                                      | 126 |
| Configure the CA SDM Ticket Status Setting .....                            | 127 |
| Rule Planning.....                                                          | 128 |
| Create Automation Policy .....                                              | 128 |
| Create a Rule .....                                                         | 129 |
| Use a Predefined Action Type .....                                          | 131 |
| Create a Custom Action .....                                                | 135 |
| Define an Action Sequence .....                                             | 136 |
| Define a Schedule.....                                                      | 138 |
| Use Cases for Policies .....                                                | 139 |
| Use Case: Adding a Server to a Service .....                                | 139 |
| Use Case: Adding a New Rule to a Service .....                              | 140 |
| Use Case: Defining an Action .....                                          | 140 |
| Configuring Data Collection .....                                           | 141 |
| Key Points About Metrics Collection .....                                   | 141 |
| Configure Data Collection for a Data Center.....                            | 144 |
| Configure Data Collection for a Server.....                                 | 145 |
| Configure Data Collection for a Virtual Resource .....                      | 147 |
| Configure Performance Thresholds .....                                      | 149 |
| Configure the Metric Filter .....                                           | 150 |
| UCS Action Types for Policies .....                                         | 152 |
| Create a Blade Power Action .....                                           | 152 |
| Service Profiles .....                                                      | 153 |
| Policy Configuration .....                                                  | 155 |
| How to Create SystemEDGE Policy .....                                       | 156 |
| How to Create SRM Policy.....                                               | 166 |
| Apply Policy to Machines .....                                              | 169 |
| Review Policy Application Progress .....                                    | 170 |
| Configure and View Applied Policies.....                                    | 171 |
| Agent Policy Dashboard Views.....                                           | 173 |
| Example: How to Monitor User-specific Metrics (MIB Extensions) .....        | 173 |
| Example: How to Monitor a Specific Windows Performance Registry Metric..... | 175 |

---

## Chapter 7: Provisioning Resources

177

|                                                                                    |     |
|------------------------------------------------------------------------------------|-----|
| Imaging Services.....                                                              | 177 |
| VMware vCenter Provisioning.....                                                   | 178 |
| VMware vCenter and vSphere User Permissions.....                                   | 178 |
| Add a Virtual Machine (vCenter Server) .....                                       | 182 |
| Add a Virtual Machine (Hyper-V Server) .....                                       | 185 |
| Add a Solaris Zone.....                                                            | 186 |
| Add a Logical Partition for an IBM AIX Computer .....                              | 187 |
| Additional Provisioning Information .....                                          | 189 |
| Remote Deployment .....                                                            | 190 |
| Remote Deployment Architecture .....                                               | 191 |
| Deployment Components .....                                                        | 192 |
| Using Multiple Distribution Servers .....                                          | 192 |
| Change the Domain Server a Distribution Server Connects To.....                    | 193 |
| Scalability .....                                                                  | 193 |
| Audit Trail.....                                                                   | 194 |
| Deployment Restrictions.....                                                       | 195 |
| Deployment Credential Restrictions .....                                           | 195 |
| Remote Deployment to UNIX/Linux Using Non Privileged User Account .....            | 195 |
| Deployment Sizing Key Factors .....                                                | 196 |
| Deploying/Installing SystemEDGE Agents Using Custom Ports .....                    | 197 |
| Deployment Package Library .....                                                   | 198 |
| Deployment Package Configuration File .....                                        | 201 |
| Deployment Packages .....                                                          | 202 |
| Default Package Wrappers.....                                                      | 203 |
| Agent Configuration Without Write Community .....                                  | 204 |
| Create a New Package Wrapper .....                                                 | 205 |
| Modify a Package Wrapper .....                                                     | 205 |
| Deployment Jobs.....                                                               | 212 |
| Track Deployment Job Status.....                                                   | 212 |
| View Deployment History .....                                                      | 213 |
| Deployment Dashboard Views.....                                                    | 214 |
| Infrastructure Deployment Process .....                                            | 215 |
| Prerequisites for Automatically Deploying CA Virtual Assurance Infrastructure..... | 216 |
| Notes on Infrastructure Deployment Using IPv6 Addresses.....                       | 218 |
| Protocols for Transferring Packages Employed by IDManager .....                    | 219 |
| Manual Installation of the Infrastructure Deployment Primer Software .....         | 219 |
| Deployment Primer Installation on Windows.....                                     | 219 |
| Deployment Primer Installation on Linux or UNIX.....                               | 220 |
| Provide the Deployment Management Certificate to a Primer Installation .....       | 220 |
| Deployment Management Certificate on Windows .....                                 | 220 |

---

|                                                                                                    |     |
|----------------------------------------------------------------------------------------------------|-----|
| Deployment Management Certificate on Linux or UNIX .....                                           | 221 |
| Compatibility Libraries for Linux .....                                                            | 221 |
| Deployment to Windows Vista, Windows 2008 and Windows XP Computers Running Firewall Software ..... | 223 |

## **Chapter 8: Managing Traps** **225**

|                                             |     |
|---------------------------------------------|-----|
| SNMP Trap Receiver .....                    | 225 |
| SNMP Trap Receiver Configuration File ..... | 226 |
| SystemEDGE Trap Forwarding.....             | 229 |

## **Chapter 9: Reporting** **231**

|                           |     |
|---------------------------|-----|
| Report Availability ..... | 231 |
| Run Reports .....         | 231 |

## **Chapter 10: Remote Monitoring** **233**

|                                                                          |     |
|--------------------------------------------------------------------------|-----|
| Overview .....                                                           | 233 |
| Advantages of Remote Monitoring .....                                    | 233 |
| Features and Benefits .....                                              | 234 |
| Agent-less Monitored Systems .....                                       | 234 |
| Key Performance Indicator Metrics .....                                  | 235 |
| Visualization .....                                                      | 235 |
| Configuration .....                                                      | 235 |
| Access Control .....                                                     | 235 |
| Resilience .....                                                         | 235 |
| Scalability .....                                                        | 236 |
| Integration .....                                                        | 236 |
| Automation .....                                                         | 236 |
| Architecture .....                                                       | 236 |
| Use Case Scenario .....                                                  | 238 |
| Configuration Prerequisites .....                                        | 239 |
| Configuring Remote Monitor Systems .....                                 | 240 |
| Collection Engine and Report Support for Remote Monitoring Metrics ..... | 243 |
| Managing Systems Using Remote Monitoring .....                           | 243 |

## **Chapter 11: Active Directory and Exchange Server AIM** **245**

|                    |     |
|--------------------|-----|
| Overview .....     | 245 |
| Architecture ..... | 246 |
| Features.....      | 247 |
| Advantages.....    | 248 |
| Performance.....   | 248 |

---

|                                                              |     |
|--------------------------------------------------------------|-----|
| Install ADES AIM .....                                       | 248 |
| Configure the Environment to Enable ADES AIM Monitoring..... | 250 |
| Configure ADES AIM with NodeCfgUtil in Command Mode .....    | 251 |
| Configure ADES AIM with NodeCfgUtil in Dialog Mode .....     | 252 |
| Uninstall the ADES AIM .....                                 | 254 |

## **Chapter 12: Scenarios and Best Practices 255**

|                                                                                  |     |
|----------------------------------------------------------------------------------|-----|
| How to Configure the vCenter Server Management Components .....                  | 256 |
| Review Requirements .....                                                        | 257 |
| Review Interactions Between vCenter Server Management Components .....           | 258 |
| Add a New vCenter Server Connection to the Manager .....                         | 260 |
| Add the AIM Instance for the vCenter Server .....                                | 264 |
| Verify the vCenter Server Folder Appearance in the Resources Tree .....          | 270 |
| How to Use Policy Actions to Identify Performance Issues.....                    | 271 |
| Create an Action for CPU Metric.....                                             | 272 |
| Create a Rule for CPU Metric to Increase Allocation .....                        | 273 |
| Create a Rule for CPU Metric to Decrease Allocation .....                        | 273 |
| Scalability Best Practices .....                                                 | 274 |
| Remote Deployment and Policy Configuration Overview.....                         | 274 |
| Hardware Specifications .....                                                    | 276 |
| Database Considerations .....                                                    | 276 |
| Network Considerations .....                                                     | 277 |
| Scalability Recommendations .....                                                | 277 |
| How to Apply Policy and Layered Templates to Servers .....                       | 289 |
| Layered Templates Concept.....                                                   | 290 |
| Create a Policy.....                                                             | 291 |
| Create Templates for Server Workload .....                                       | 301 |
| Apply Policy and Templates to Servers and Verify Settings .....                  | 315 |
| (Optional) Manage the Base Policy and Templates for One or More Servers .....    | 317 |
| (Optional) Update the Policy or Templates.....                                   | 318 |
| (Optional) Apply Policy and Template Updates to Servers and Verify Updates ..... | 318 |

## **Appendix A: FIPS 140-2 Encryption 321**

|                    |     |
|--------------------|-----|
| FIPS Overview..... | 321 |
|--------------------|-----|

## **Appendix B: Troubleshooting 323**

|                                                                      |     |
|----------------------------------------------------------------------|-----|
| Adjusting Poll Interval Settings for Solaris Zones Environments..... | 323 |
| Attributes Show a Value of Zero .....                                | 324 |
| Browsers Do Not Display Consecutive Spaces in Events .....           | 324 |
| CA NSM Fails to Receive Traps on Windows .....                       | 324 |

---

|                                                                                          |     |
|------------------------------------------------------------------------------------------|-----|
| Cisco UCS Folder Does Not Display in UI .....                                            | 325 |
| DB Transaction Log Sizes Increase Unexpectedly .....                                     | 325 |
| Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero .....                    | 326 |
| Empty Task ID for the dpmvc virtualswitch Command .....                                  | 327 |
| Local and Remote Monitors Do Not Show the Same Values .....                              | 327 |
| Remote Deployment to Solaris Lists SPARC and x86 Systems .....                           | 328 |
| PMM Stops Polling an AIM .....                                                           | 328 |
| Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear ..... | 329 |
| Resetting the vCenter Server Password Causes Data Collection to Fail.....                | 329 |
| Solaris Zones AIM Reset if a Monitored System is Down .....                              | 329 |
| Unable to Connect to Microsoft SQL Server .....                                          | 330 |
| Upgrading SystemEDGE .....                                                               | 330 |
| User Interface is Unresponsive on Provisioning and Policy Screens .....                  | 330 |
| vCenter Server AIM Attributes Show Zero .....                                            | 331 |
| vCenter Server Connection Failed .....                                                   | 331 |
| vCenter AIM Instance Status Icon Shows Disabled .....                                    | 333 |
| vCenter AIM Instance Status Icon Shows Discovery in Progress .....                       | 333 |
| vCenter AIM Instance Status Icon Shows Error .....                                       | 334 |
| vCenter AIM Instance Status Icon Shows Multiple Instances .....                          | 335 |
| vCenter AIM Instance Status Icon Shows Warning .....                                     | 336 |
| VM Usage Values Do Not Update Immediately After Power Down .....                         | 336 |
| Blank Query Results Tab after Upgrade .....                                              | 337 |

|                 |            |
|-----------------|------------|
| <b>Glossary</b> | <b>339</b> |
|-----------------|------------|

|              |            |
|--------------|------------|
| <b>Index</b> | <b>349</b> |
|--------------|------------|



# Chapter 1: Introduction

---

This section contains the following topics:

[Audience](#) (see page 15)

[Related Publications](#) (see page 15)

[Conventions](#) (see page 16)

## Audience

This guide is intended for administrators who install, configure, and use CA Virtual Assurance to manage virtual environments. It assumes that you are familiar with the operating systems used in your environment, virtualization technologies, and SNMP.

## Related Publications

The CA Virtual Assurance documentation consists of the following deliverables:

### **Administration Guide**

Explores how to administer and use CA Virtual Assurance to manage virtual resources in your environment.

### **Installation Guide**

Contains brief architecture information, various installation methods, post-installation configuration information, and Getting Started instructions.

### **Online Help**

Provides window details and procedural descriptions for using the CA Virtual Assurance user interface.

### **Reference Guide**

Provides detailed information about AutoShell, CLI commands, MIB attributes, and performance metrics.

### **Release Notes**

Provides information about operating system support, system requirements, published fixes, international support, known issues, and the documentation roadmap.

### **Service Response Monitoring User Guide**

Provides installation and configuration details of SRM.

### SystemEDGE User Guide

Provides installation and configuration details of SystemEDGE.

### SystemEDGE Release Notes

Provides information about operating system support, system requirements, and features.

## Conventions

This guide uses the following conventions:

### Case-Sensitivity

All names of classes, commands, directives, environment parameters, functions, and properties mentioned in this guide are case-sensitive and you must spell them exactly as shown. System command and environment variable names *may* be case-sensitive, depending on your operating system's requirements.

### Cross-References

References to information in other guides or in other sections in this guide appear in the following format:

#### Guide Name

Indicates the name of another guide.

#### "Chapter Name"

Indicates the name of a chapter in this or another guide.

### Synonyms

Terms such as attribute, object, object identifier (OID) are synonymous to the term 'variable' in this document.

Terms such as SystemEDGE Agent, CA SystemEDGE are synonymous to SystemEDGE in this document.

### Syntax

Syntax and user input use the following form:

#### *Italic*

Indicates a variable name or placeholder for which you must supply an actual value.

#### {a|b}

Indicates a choice of mandatory operands, a or b.

#### [ ] or [[ ]]

Indicates optional operands.



### Syntax Example

The following example uses these conventions:

```
modify -t ZONE [-m zoneserver] -p psetname {-min mincpu|-max maxcpu} pset -session ssh
```

The operands `-min` and `-max` are mandatory, but you can only use one of them depending on what you want to define, the minimum number of CPUs in the processor set or the maximum number. The operand `-m` is not required for this command to function. All other parts of the command must be entered as shown.

### Default Directory

*sysedge\_home* used in path statements indicates the directory in which SystemEDGE is installed. **Default:** C:\Program Files\CA\SystemEDGE.

### Installation Path

*Install\_Path* used in path statements indicates the directory in which CA Virtual Assurance or components of CA Virtual Assurance are installed.

#### Defaults:

- Windows x86: C:\Program Files\CA
- Windows x64: C:\CA, C:\Program Files (x86)\CA, or C:\Program Files\CA
- UNIX, Linux: /opt/CA



# Chapter 2: Overview

---

This section contains the following topics:

[Architecture](#) (see page 19)

[Databases](#) (see page 22)

[User Interface](#) (see page 23)

## Architecture

CA Virtual Assurance is a policy-based product that automatically monitors physical and virtual resources to meet the load demands of complex data centers dynamically. CA Virtual Assurance uses a Service Oriented Architecture (SOA) and continuously analyzes your data center to verify that your servers are optimally provisioned to perform required tasks. You can manage your data center and obtain detailed information about each managed computer in your data center using the web-based CA Virtual Assurance user interface.


CA Virtual Assurance integrates with the following external technologies:

- VMware vCenter Server for virtual machine (VM) management
- Microsoft Hyper-V Server for virtual machine (VM) management and optional integration with Microsoft System Center Virtual Machine Manager (SCVMM) for provisioning
- IBM Logical Partitioning (LPAR) for AIX virtualization management
- Solaris Zones virtualization management
- Cisco Unified Computing System
- Microsoft Cluster Service

CA Virtual Assurance leverages the following existing technologies:

- CA SDM integration for request escalation and resolution
- CA Embedded Entitlements Manager for security
- Common discovery for lightweight stand-alone discovery capabilities

The following diagram represents the product architecture:



Each component registers with the service controller on startup. The service controller is a central component for identifying the location and status of all components. After a component is registered, it provides its Web Services Description Language (WSDL) location and publishes its events for other components to subscribe to and receive real-time notification of changes to the data center.

The Provisioning component provides provisioning capabilities for vCenter Server, Hyper-V, and Solaris Zones. Additionally, you can remotely deploy SystemEDGE and the AIMs provided by CA Virtual Assurance to remote servers.

The Resource Manager component is responsible for updating all resources in the product, such as creating and updating user-defined services.

The Initiation component provides integration with Microsoft Scheduler for job scheduling. Long running maintenance tasks and actions can be scheduled as jobs, for example.

CA Embedded Entitlements Manager is used for all security and role-based management.

CA Virtual Assurance Event Manager captures all events generated by CA Virtual Assurance components and provides SNMP forwarding. SNMP forwarding can be used to forward events to any CA or third-party product capable of receiving SNMP traps.

The Policy component analyzes the collected data.

The Performance Monitor component integrates with SystemEDGE to collect system performance data. SystemEDGE must be installed on any server from which you want to collect the base system metrics, unless you are using the Remote Monitoring AIM to monitor Windows servers remotely without installing any additional software on them (zero footprint). All performance metrics are stored in the Performance Database.

After the Management DB is populated with the managed servers, the Performance Monitor begins gathering information to determine if the servers can collect performance metrics.

The Policy component analyzes the collected data to determine which user-defined business rules have been breached, and runs actions on the target servers or services. You define the rules and actions to resolve a particular problem in advance. The policy component uses your parameters to make intelligent decisions. After the server or service is identified, you can perform various actions to resolve the situation. For example, run custom scripts, provision new systems, and more.

Finally, you can monitor the data center performance through customizable graphs and charts directly from the graphical user interface or through reports generated by CA Virtual Assurance. CA Virtual Assurance lets you also view its current component status, configure components, validate settings, or create user access lists.

## Databases

The product uses both a management database and a performance database.

### Management DB

The Management DB is a common data repository for all managed objects, based on a model for describing management data. The Management DB stores information about servers, services, rules, actions, virtual platform objects, events, alerts, and relationships among these objects.

CA Virtual Assurance uses the Management DB to store the following information:

- Server information
- Service relationships
- Service thresholds
- Rules and actions
- Events
- Credentials for other components

**Note:** For more information about configuring the Management DB, see the chapter "Command Line Utilities" in the *Reference Guide*.

### Performance DB

The Performance DB is a repository that stores all the metrics collected from the servers in your data center.

CA Virtual Assurance uses the Performance DB to store the following information:

- Which metrics are collected from which servers
- Values of those metrics (aggregated over time)
- Server-level recording interval
- Server-level thresholds (overall server utilization)
- Data center-level recording interval
- Data center-level thresholds

The data stored in this database is used for various functions. For example, this DB is the source of the data used to create historical reports. CA Virtual Assurance also uses the data in this database and user-created rules to make logical business decisions.

**Note:** For more information about configuring the Performance DB, see the section `dpmutil set | get perf Command`—Configure the Performance Database.

## User Interface

You can manage your data center from a central location using the CA Virtual Assurance web-based user interface. You can use the functions of the embedded components in CA Virtual Assurance without having to open the component interfaces separately.

For example, you can use CA SDM for issue management and JasperReports (default reporting engine) for reporting from the CA Virtual Assurance web-based user interface. You can also use CA Embedded Entitlements Manager functionality to take advantage of active directory and manage your users and permissions from the user interface without having to open CA EEM.

You always have the option to directly access component user interfaces to perform more advanced functions. For example, you may want to open CA Embedded Entitlements Manager to use native security. You can log in to the server where the component is installed to directly access its user interface or you can go to the Administration tab, Configuration page in the CA Virtual Assurance UI, select a component, and launch the product Home Page from the Actions drop-down menu.

## Access the User Interface

Access the user interface to discover and provision virtual and physical systems, to create policy, to schedule jobs, and so on. The Start menu shortcut is only available on the CA Virtual Assurance server. You must access the manager server directly, and use the Start menu to use some product functions such as the CLI and Autoshell. Users accessing the interface from a separate server must enter the URL in a web browser.

### To access the user interface

1. Select Start, Programs, CA, CA Virtual Assurance, Launch CA Virtual Assurance from the CA Virtual Assurance server.

The CA Virtual Assurance login page appears at the following URL:

`https://servername:port/UI`

***servername***

Identifies the name of the server where the graphical user interface is installed.

***port***

Specifies the port that the server is listening on.

**Default:** 8443

2. Enter your administrator login credentials and click Login.

The Dashboard appears.



# Chapter 3: Managing Users

---

This section contains the following topics:

[User Access Control](#) (see page 25)

[Password Management](#) (see page 29)

[Tasks for External Directory Users](#) (see page 33)

[User-scoped Authentication for vCenter Server](#) (see page 37)

## User Access Control

CA Embedded Entitlements Manager secures all communication between CA Virtual Assurance components. You can select one of the following configurations:

- Active Directory
- Native Security

**Note:** For more information about configuring external directories, see the *CA EEM Getting Started* and *Online Help* installed with CA Embedded Entitlements Manager. Locate the documentation from Start, Programs, CA, Embedded Entitlements Manager, Documentation where CA Embedded Entitlements Manager is installed or on the CA Support Online website at <http://ca.com/support>.

## Active Directory

When you connect to an existing Active Directory configuration, your predefined users and user groups remain consistent with your central repository of users. CA Technologies recommends that you create and modify users in Active Directory instead of using CA Virtual Assurance or CA Embedded Entitlements Manager.

CA Virtual Assurance uses the Lightweight Directory Access Protocol (LDAP) to read from and write to the Microsoft Active Directory server. By default, LDAP traffic is transmitted unsecured. This results in unsecured communication between the server and Microsoft Active Directory. To make Microsoft Active Directory secure, use LDAP over Secure Sockets Layer (SSL)—LDAPS. In this case, install a properly formatted certificate from either a Microsoft certification authority or another certification authority.

**Note:** For more information about configuring Active Directory to transmit data securely, see the Microsoft website. Search for the Knowledge Base article "How to enable LDAP over SSL with a third-party certification authority." After you configure Active Directory to use LDAPS, you can transmit data securely.

## Security Considerations for Active Directory

The Lightweight Directory Access Protocol (LDAP) is used to read from and write to the Microsoft Active Directory server. LDAP traffic is transmitted unsecured by default. This results in unsecured communication between the server and Microsoft Active Directory. You can make Microsoft Active Directory secure by using LDAP over Secure Sockets Layer (SSL)—LDAPS. You must install a properly formatted certificate from either a Microsoft certification authority or a non-Microsoft certification authority.

The requirements are described in a Microsoft Knowledge Base article.

**Note:** For more information about configuring Active Directory to transmit data securely, see the Knowledge Base article "How to enable LDAP over SSL with a third-party certification authority" on the Microsoft website. After you configure Active Directory to use LDAPS, you can transmit your data securely.

## Native Security

Native Security lets the CA EEM administrator create users, user groups, and policies specifically for CA Virtual Assurance because this information resides in the local store. Native Security requires you to define your own set of users and user groups manually. Those users and user groups may not be consistent with what is currently defined in the directory service.

## How CA Embedded Entitlements Manager Works with CA Virtual Assurance

CA Embedded Entitlements Manager includes the following key objects:

- Identities (users and user groups)
- Resources
- Policies

CA Embedded Entitlements Manager provides the following capabilities:

### Authentication

Authenticates the user. The authenticated user can then be used in subsequent authorization processing.

### Authorization

Permits a user to access a particular resource. A resource can be any logical or physical entity. In CA Virtual Assurance, the typical resource is a user interface component (for example, tab, command, drop-down list, and so on). A set of policies associated with a resource class control authorization. These policies are the primary way to integrate CA Embedded Entitlements Manager with CA Virtual Assurance.

## Access the CA EEM User Interface

Log in to the CA Embedded Entitlements Manager home page to use native security. The CA Embedded Entitlements Manager documentation is also available from the Start menu, and Online Help is available on the home page after you log in.

### To access the CA Embedded Entitlements Manager user interface

1. Select Start, Programs, CA, Embedded Entitlements Manager, EEM UI.

The CA Embedded Entitlements Manager Log In window appears.

**Note:** At logon, if you receive a security certificate request, bypass it and continue. To eliminate security certificate messages, you can acquire a certificate from the vendor of your choice and apply it to the server. For information about installing security certificates, see the Apache Tomcat website.

2. Select AIP from the application drop-down list.

The User Name field is populated with EiamAdmin.

3. Enter your password in the Password field and click Log In.

The CA Embedded Entitlements Manager Home Page appears with the home page displayed by default.

## Create CA EEM Users

To give a user access to CA Virtual Assurance, create a CA EEM user. This procedure describes how to add CA EEM users manually to the common data store used by CA EEM for CA Virtual Assurance. You can also add users by referencing an external directory.

**Note:** For more information about adding users by referencing an external directory, see the *CA EEM Getting Started* and *Online Help* that is installed with CA EEM.

### To create CA EEM users

1. Click Manage Identities on the Home Page of the CA Embedded Entitlements Manager home page.

The Users page is selected by default.

2. Select the Application User Details option in the Search Users section.

3. Leave User Name selected in the Attribute drop-down list, leave LIKE selected in the Operator drop-down list, leave the Value field blank, and click Go.

All CA Virtual Assurance users are listed in a hierarchical tree in the Users pane.

4. Click the New User icon in the left pane.

The New User pane appears on the right.

5. Enter the user ID for this user in the User Name field and click Add Application User Details in the User Details pane.
6. Select the application group from the Available User Groups box in the Application Group Membership pane, and click the right arrow .

The application group is added to the Selected User Groups.

**Note:** You can also add this user to one or more dynamic groups or global groups. For more information, see the CA EEM documentation.

7. Enter the password for the user in the New Password and Confirm Password fields on the Authentication pane, and click Save.

A confirmation message appears below the Users pane.

## Create Default User Groups

User groups let you group users logically by business function. You can create a user group to give multiple users the same access rights. Although this procedure only describes creating an application group, subsequent procedures describe policy creation for that application group. You can also create policies for global groups, dynamic groups, and individual users.

### To create user groups

1. Click Manage Identities on the Home tab of the Access CA Embedded Entitlements Manager home page.

The Users page is selected by default.

2. Click Groups, select the Show Application Groups check box, and click Go.

All available application groups are listed under Application Groups in the User Groups pane.

3. Click New Application Group in the left pane.

The New Application User Group page appears in the right pane.

4. Enter a name for the new application group and click Save.

The new Application User Group is created.

# Password Management

User credentials are essential for the communication between CA Virtual Assurance components. CA Virtual Assurance stores user and password information internally. When you change passwords of external components or applications CA Virtual Assurance integrates with, change these passwords in CA Virtual Assurance for consistency. Otherwise, CA Virtual Assurance does not work properly.

Consider the following areas:

- Active Directory security
- Native security
- CA Embedded Entitlements Manager administrator
- Database sa user (SQL authentication)

## Change the CA Embedded Entitlements Manager Administrator Password (EiamAdmin)

If you intend to change the CA Embedded Entitlements Manager administrator password (EiamAdmin), change the password in CA Embedded Entitlements Manager and also in CA Virtual Assurance.

### **To change the administrator password (EiamAdmin) in CA Embedded Entitlements Manager**

1. Navigate to Start, Programs, CA, Embedded Entitlements Manager, EEM UI and open the user interface.

The login dialog appears.

2. Log in with the current EiamAdmin password.

The user interface opens.

3. Click Configure and EEM Server.

The EEM Server pane appears.

4. Click EiamAdmin Password.

The New Password and Confirm Password fields appear.

5. Enter your password and click Save.

The new EiamAdmin password can now be used to log in CA Embedded Entitlements Manager.

### To change the administrator password (EiamAdmin) in CA Virtual Assurance

1. Navigate to Start, Programs, CA, CA Virtual Assurance, CA Virtual Assurance Command Prompt.

The command prompt appears.

2. Enter the following command:

```
dpmutil -set -eiam
```

The dpmutil command prompts you for the required credentials.

Complete the command.

3. Recycle the CAAIPApache and CAIPTomcat services.

The credentials are now consistent and CA Virtual Assurance works as expected.

**Note:** In both cases the Apache log file, located at *Install\_path*\Apache\logs\error.log, can confirm proper product startup. If the last entry is "Validating EEM is available," then there is still a credential problem. Verify that the credentials used for '-set -eiam' and '-set -sysuser' can be used to log in to the CA Embedded Entitlements Manager UI. Then, retry the dpmutil commands using valid credentials.

## Change the Database Administrator (sa) Password

If you use Microsoft SQL Authentication and you change the password for the Microsoft SQL user (typically the 'sa' user), change the CA Virtual Assurance password also.

### To change the database administrator (sa) password in Microsoft SQL Server

1. Open Microsoft SQL Server Management Studio and log in.
2. In the Object Explorer expand Security, Logins.
3. Open sa and change the password in the right pane.

**Note:** For further details, see the Microsoft SQL Server documentation.

### To change the database administrator (sa) password in CA Virtual Assurance

1. Navigate to Start, Programs, CA, CA Virtual Assurance, CA Virtual Assurance Command Prompt.

The command prompt appears.

2. Enter the following command:

```
dpmutil -set -mgmtadb
```

The dpmutil command prompts you for the appropriate credentials.

Complete the command.

3. If the performance database uses the same server and database user (sa), enter the following command:

```
dpmutil -set -perfdb
```

The dpmutil command prompts you for the appropriate credentials.

Complete the command.

4. Recycle the CAAIPApache and CAIPTomcat services.

The credentials are now consistent and CA Virtual Assurance works as expected.

## Change the System User Password for Native Security

CA Virtual Assurance requires the `sys_service` system user to function correctly, for example, to start or stop the Apache service. You specify the `sys_service` system user and its password during an installation with native security. The installation program stores the `sys_service` credentials in CA Embedded Entitlements Manager and CA Virtual Assurance. If you change the password for `sys_service` in CA Embedded Entitlements Manager later, also change it in CA Virtual Assurance to help ensure that all CA Virtual Assurance services continue running.

### To change the `sys_service` password in CA Embedded Entitlements Manager

1. Navigate to Start, Programs, CA, Embedded Entitlements Manager, EEM UI and open the user interface.

The login dialog appears.

2. Log in with the current EiamAdmin password.

The user interface opens.

3. Click Manage Identities and Search Users.

The users appear in the Users pane.

4. Click the `sys_service` user.

The user properties appear in the right pane.

5. Scroll down to the Authentication section and click Reset Password.

The New Password and Confirm Password fields appear.

6. Enter your password and click Save.

The new password is now stored in CA Embedded Entitlements Manager.

### To change the `sys_service` user password in CA Virtual Assurance

1. Navigate to Start, Programs, CA, CA Virtual Assurance, CA Virtual Assurance Command Prompt.

The command prompt appears.

2. Enter the following command:

```
dpmutil -set -sysuser
```

The dpmutil command prompts you for the required credentials.

Complete the command.

3. Recycle the CAAIPApache and CAIPTomcat services.

The credentials are now consistent and CA Virtual Assurance works as expected.

## Change the System User Password for Active Directory Security

If your CA Virtual Assurance installation is configured to connect to Active Directory, the user who installs CA Virtual Assurance is automatically registered with CA Embedded Entitlements Manager. This registration allows CA Virtual Assurance to authenticate users from the Active Directory domain. If the user password changes, users cannot log in to the CA Virtual Assurance user interface because CA Embedded Entitlements Manager can no longer authenticate them. Change the user password as follows:

### To change the user password for Active Directory

1. Navigate to Start, Programs, CA, Embedded Entitlements Manager, EEM UI and open the user interface.

The login dialog appears.

2. Log in with the current password.

The user interface opens.

3. Click Configure and EEM Server.

The EEM Server pane appears.

4. Click Global Users/Global Groups in the left pane and retain default option "Reference from an external directory" selected.

5. Retain default Type as Microsoft Active Directory and enter a new password in the Password and Confirm Password fields and click Save.

6. Close CA Embedded Entitlements Manager

7. Navigate to Start, Programs, CA, CA Virtual Assurance, CA Virtual Assurance Command Prompt.

The command prompt appears.



8. Enter the following command:

```
dpmutil -set -sysuser
```

Sysuser is the same user who installs CA Virtual Assurance. The dpmutil command prompts you for the required credentials specified in Step 5.

Complete the command.

9. Recycle the CAAIPApache and CAIPTomcat services.

The credentials are now consistent and CA Virtual Assurance works as expected.

**Note:** In both cases the Apache log file, located at *Install\_path*\Apache\logs\error.log, can confirm proper product startup. If the last entry is "Validating EEM is available," then there is still a credential problem. Verify that the credentials used for '-set -eiam' and '-set -sysuser' can be used to log in to the CA Embedded Entitlements Manager UI. Retry the dpmutil commands using valid credentials.

## Tasks for External Directory Users

You can complete user access control tasks from the CA Virtual Assurance user interface regardless of the external directory that you are using.

### Create User Groups

User groups let you group users logically according to business functions. You can create a user group to give multiple users the same access rights.

#### To create user groups

1. Click Administration.

The Administration page appears.

2. Click User Groups.

The User Groups page appears.

3. Type a Name for the user group. The name can be based on a business function or service.

4. (Optional) Type a Description.

5. Click Save.

The new user group appears in the left pane.

## Assign Users to Groups

Users inherit the access privileges assigned to their user group. You can add new users to an existing user group when you want to grant its access rights to them. The administrator user group is a predefined group and appears in the list by default.

### To assign users to groups

1. Click Administration.  
The Administration page appears.
2. Click User Group.  
The User Groups page appears.
3. Expand User Groups and select a user group from the list.  
A submenu appears.
4. Select the Membership submenu.  
A series of membership panes appears.
5. Enter the user name to add in the Value text box, and click Search.  
The search results appear in the Available User/User Group pane or a message notifies you that no match was found. If you are unsure of the user name, you can [search for users or user groups](#). (see page 36)
6. Select the user to add from the Available User/User Group pane, and click the right arrow .  
The user name moves to the Selected User/User Group pane.
7. Click Save to finish adding users.  
Users are granted the access privileges of their user group.

## Remove Users or User Groups from a User Group

You can remove users and user groups from an existing CA Virtual Assurance user group. The administrator user group is a predefined group and appears in the list by default.

### To remove users or user groups from a user group

1. Click Administration, then Configuration.  
The Configuration page appears.
2. Select User Groups.  
The User Groups menu appears on the left pane.

3. Expand User Groups and select a user group from the list.  
A submenu appears on right pane.
4. Select the Membership submenu.  
A series of membership panes appears.
5. Select the user or user group to remove from the Selected User/User Group pane and click the left arrow .  
The user or user group is moved to the Available User/User Group pane.
6. Click Save when you finish removing users and user groups.

## Assign User Groups Access Rights to Services

In environments with multiple groups of users, it is typically necessary to prevent one group from viewing the resources of another. Administrators can assign specific resources to groups of users. Some administrators can assign resources only for groups in which they are members. Administrators in the group *AlPAdmins*, however, have full access for assigning resources.

### To assign user groups access rights to services

1. Click Administration.  
The Administration page appears.
2. Click User Group.  
The User Groups page appears.
3. In the left pane, select a user group for which to set permissions, and click Service Access.  
A tree listing of services defined to the system appears.
4. Select the services for which you want to grant or restrict access, and click Save.  
User groups are granted the access privileges that are assigned to their associated services.

## Set User Group Permissions

You can control user group access to functional areas and specific functions in the user interface. The AIPAdmins user group has access to all functional areas and functions by default.

### To set user group permissions

1. Click Administration.  
The Administration page appears.
2. Click User Group.  
The User Groups page appears.
3. Select a user group for which to set permissions, and click Privileges.  
The Privileges page appears.
4. Select the functional areas or specific functions for which you want to grant or restrict access, and click Save.  
The user permissions are updated.

## Delete User Groups

You can delete user groups that you no longer need.

### To delete user groups

1. Click Administration.  
The Administration page appears.
2. Click User Group.  
The User Groups page appears.
3. Right-click a User Group, and select Delete User Group.  
The user group is removed.

## Search for Users or User Groups

You can search for users or user groups that you want to add or delete.

### To search for users or user groups

1. Click Administration.  
The Administration page appears.

2. Click User Group.

The User Groups page appears.

3. Expand User Groups and select a user group from the list.

The user group page appears in the right pane.

4. Click Membership.

The User/User Group page appears.

5. Select Users or User Groups in the Identity drop-down list. Select the attribute to search for in the Attribute drop-down list, and leave the LIKE operator selected. Enter the value (or a partial value with a wildcard) in the Value field, and click Search.

A list of matching user or user group names appears in the Available User/User Groups list.

## User-scoped Authentication for vCenter Server

You can enable user-scoped authentication for vCenter Server environments by adding a configuration entry to the `caaipconf.cfg` file located in the `Install_Path\VirtualAssurance\conf` directory. Because this entry does not exist after installation, user-scoped authentication is disabled by default. In this case CA Virtual Assurance uses the user for vCenter Server authentication who is specified in the vCenter Server configuration pane under Administration.

In contrast, enabled user-scoped authentication uses the currently logged in user (user interface) for authenticating vCenter Server environment operations. User-scoped authentication implies that appropriate users and their permissions are specified in vCenter Server. The same users must also be specified in CA EEM to log in the CA Virtual Assurance user interface.

### To enable user-scoped authentication

1. Specify the required users and their permissions (administrator or read-only) in vCenter Server.
2. Specify the same users in CA EEM.
3. Change to the CA Virtual Assurance manager server and navigate to the *Install\_Path*\VirtualAssurance\conf directory.
4. Open the caaipconf.cfg file with a text editor and add the following entry to the AIP product section:

```
<property name="USER_SCOPED_AUTHENTICATION">
  <value>VC</value>
  <displayName>The vCenter PMM component uses the currently logged in user for authenticating vCenter
Server platform operations.</displayName>
</property>
```

Result:

```
<properties targetNamespace="http://www.ca.com/cfg/types/2008/05">
  <product name="AIP">
    ...
    <property name="USER_SCOPED_AUTHENTICATION">
      <value>VC</value>
      <displayName>The vCenter PMM component uses the currently logged in user for authenticating
vCenter Server platform operations.</displayName>
    </property>
    ...
  </product>
  ...
</properties>
```

5. Save the file.  
CA Virtual Assurance automatically detects the change.
6. Verify that a currently logged in user has the same permissions in CA Virtual Assurance for managing the VMware environment as specified in vCenter Server.

**Note:** If you want to disable user-scoped authentication, remove the entry from the caaipconf.cfg file.

### Example

**Initial scenario:** During a CA Virtual Assurance installation, the user *CA* has been configured to log in the CA Virtual Assurance user interface. On vCenter Server, *administrator* is the user with administrator permissions. CA Virtual Assurance is configured to use *administrator* for authenticating vCenter Server environment operations (see Administration tab, vCenter Server Configuration pane in the user interface). User-scoped authentication is disabled by default.

This scenario conforms for example to a full custom installation with an appropriate vCenter Server configuration.

Assume the following:

- Two additional users are configured on the vCenter Server: *Superuser* (administrator) and *Reader* (read-only permissions)
- *Superuser* and *Reader* are added to CA EEM
- User-scoped authentication is enabled

When you log in CA Virtual Assurance as *Superuser*, then you have administrator permissions for managing vCenter Server.

When you log in CA Virtual Assurance as *Reader*, then you have read-only permissions for monitoring vCenter Server.

When you disable user-scoped authentication, everyone who logs in CA Virtual Assurance has vCenter Server administrator permissions. In case of disabled user-scoped authentication, CA Virtual Assurance uses the user *administrator* specified in the vCenter Server Configuration pane under Administration in the user interface (see also the initial scenario of this example).





# Chapter 4: Configuring Resources

---

Verify whether you must perform the tasks described in this section after installing CA Virtual Assurance. Depending on your installation, additional configuration steps can be necessary.

This section contains the following topics:

[Discovery](#) (see page 41)

[Configuration Requirements for Managing Environments](#) (see page 41)

[Configure SNMPv1 Traps by Editing the sysedge.cf File](#) (see page 52)

[Event Forwarding](#) (see page 54)

## Discovery

After CA Virtual Assurance installation, add the SNMP read and write credentials used with your network devices and monitored servers to the SNMP Settings panel. Open the Administration tab, Configuration to access the SNMP Settings panel.

**Note:** Add the SNMP credentials before starting discovery.

If discovery was already done, discover the system again.

## Configuration Requirements for Managing Environments

If you have not entirely discovered your environments and configured all AIMs during installation, install required programs or configure your AIMs as described in this section.

## Discover Environments

You can discover virtual environments and Microsoft Clusters during the installation of AIMs when using a custom CA Virtual Assurance manager installation. If possible, provide the required data about managed servers in the installation dialogs. You can skip the configuration during installation and perform the configuration later.

To configure the AIMs and discover virtual environments after a custom installation, do one of the following:

- Open the Administration tab from the user interface, change to Configuration, Provisioning and select the appropriate server type to add credentials and configure the AIM. CA Virtual Assurance automatically discovers the physical and virtual components and populates the Management Database.
- Use NodeCfgUtil.exe utility on a Windows AIM Server to add the required data for managing virtual environments. The utility is located in the *SystemEDGE\_install\_path\plugins\AIPCommon* directory. Then rediscover the AIM Server from the CA Virtual Assurance manager. This option lets you manually perform the required steps.

Consider the following guidelines

- The users specified for accessing virtual environments and Microsoft Clusters must have sufficient privileges to allow remote access.
- To manage Hyper-V Servers, install SystemEDGE and the Hyper-V AIM on the Hyper-V Server. SystemEDGE and the Hyper-v AIM must run on the same Hyper-V Server. Then configure the AIM and discover the Hyper-V Server.
- Verify that SystemEDGE and the AIMs which do not run on the CA Virtual Assurance manager server use the same SNMP settings as their associated CA Virtual Assurance manager.
- To manage VMware Infrastructure or vSphere, enter the credentials for the corresponding vCenter Servers.
- Install VMware Tools on all VMware VMs. If VMware tools are not installed on a VM, some features do not work properly.

VMware Tools optimize the virtualization of VMs. Without these tools, many features are not available. For further information, see the VMware documentation.

**To discover vCenter Servers, Solaris Zones Servers, HMC/IVM Servers, Microsoft Clusters, Cisco UCS Servers, Active Directory, or Exchange Servers from an AIM Server**

1. Run the NodeCfgUtil.exe utility on a Windows AIM Server to update the configuration data for the corresponding AIMS.  
  
The NodeCfgUtil.exe utility stores the data in zone.cfg, vc.cfg, ucs.cfg, lpar.cfg, mscs.cfg, or esad.cfg.
2. Open the user interface on the Manager Server and click Resources, Data Center in the navigation pane.
3. Right-click and select Management, Discover.  
  
The discovery options appear.
4. Select one of the following:
  - Discover a system
  - Discover a network  
The corresponding dialog opens.
5. Enter a system name of a vCenter Server, Solaris Zones Server, Microsoft Cluster, Cisco UCS Server, or HMC/IVM Server that you want to manage. Alternatively, you can enter network properties for the discovery process. Click OK.  
  
CA Virtual Assurance starts the discovery process.  
  
The discovered resources appear in the Explore pane.

**More Information**

[How to Configure the vCenter Server Management Components](#) (see page 256)

## Configure AIMS with NodeCfgUtil in Dialog Mode

NodeCfgUtil lets you modify the AIM configurations for vCenter Server, IBM LPAR, Solaris Zones, Microsoft Clusters, Cisco UCS Server, Active Directory and Exchange Server (ADES). When you use the utility in dialog mode, you can configure which nodes the appropriate AIMS manage.

**Note:** NodeCfgUtil.exe must be run as Windows Administrator.

**Follow these steps:**

1. Log in as Administrator, open Windows Explorer on the system on which the AIM resides, change to the *SystemEDGE\_InstallPath\plugins\AIPCommon* directory, and start NodeCfgUtil.exe. NodeCfgUtil discovers the installed AIMS and lists only those in subsequent dialogs.
2. Enter **1** to add a new managed node.  
  
NodeCfgUtil can only discover installed AIMS.

3. Follow the on-screen instructions to complete the configuration. Each node requires a valid user name and password for authentication.
4. After a successful configuration, enter 0 to return to previous menus, or to exit the utility.

NodeCfgUtil writes a configuration file for Solaris Zones (zone.cfg), vCenter Server (vc.cfg), Microsoft Clusters (mscs.cfg), UCS (ucs.cfg), LPAR (lpar.cfg), or ADES (esad.cfg) to the *SystemEDGE\_InstallPath\plugins\AIPCommon* directory. You can also use the NodeCfgUtil utility to edit or remove existing entries. The corresponding dialogs are self-explaining.

**Note:** When removing a server entry, NodeCfgUtil marks this server for deletion, but it does not remove the server from the user interface or database. To remove the server from the user interface and database, rediscover the Windows server on which the AIM runs. The new discovery process detects the deletion mark and removes the corresponding server from the user interface and database.

## Examples

The following example shows the Install Managed Node dialog for the myvc5 server that has been successfully added to the configuration of the vCenter AIM. The AIM is now ready to manage the vCenter Server. The vCenter AIM is a multi-instance AIM. So you can repeat this procedure and add more vCenter Servers that you want to manage with this AIM.

\*\*\*\* Choose Managed Node \*\*\*\*

1. IBM LPAR
2. Solaris Zones
3. VMware vCenter
4. Cisco UCS
5. Microsoft Cluster
0. Go Back to Previous Menu

\*\*\*\*\*

Enter choice: 3

Enter following information for the VMware vCenter Node...

(At any point to go back to previous menu, Enter 'CTRL Q')

1. Server Name: myvc5
2. User Name: administrator
3. Password: \*\*\*\*\*
4. Port [default=443]:
5. Protocol [default=https]:

CAAC1016 Authenticating, please wait...

CAAC1019 Authentication SUCCESSFUL.

CAAC1023 Added Node Successfully.

Press any key to continue . . .

The following example shows the Install Managed Node dialog for mydomain that has been successfully added to the configuration of the ADES AIM. Management Entity is set to Active Directory. Management Mode is set to Entire Domain. For details, see the NodeCfgUtil command mode. The ADES AIM is a multi-instance AIM. So you can repeat this procedure and add more entities that you want to manage with this AIM.

\*\*\*\* Choose Managed Node \*\*\*\*

1. Microsoft Cluster
2. Microsoft Exchange Server and Active Directory
0. Go Back to Previous Menu

\*\*\*\*\*

Enter choice: 2

Enter following information for the Microsoft Exchange Server and Active Directory Node...

(At any point to go back to previous menu, Enter 'CTRL Q')

1. Domain Name: mydomain
2. User Name: administrator
3. Password: \*\*\*\*\*
4. Management Entity: 0
5. Management Mode: 0

CAAC1016 Authenticating, please wait...

CAAC1018 Credential authentication SUCCESSFUL.

Press any key to continue . . .

## Configure AIMs with NodeCfgUtil in Command Mode

NodeCfgUtil.exe lets you modify the AIM configurations for IBM LPAR, Solaris Zones, VMware vCenter, Microsoft Clusters, Cisco UCS, Active Directory and Exchange Server (ADES). When you use the utility in command mode, you can only add managed nodes to an AIM configuration.

**Note:** NodeCfgUtil.exe must be run as Windows Administrator.

This command has the following format:

- (1) nodecfgutil -help
- (2) nodecfgutil {lpar|zone|mscs} -u *user* -p *password* -h {pvmname|hostname|cluster\_name}
- (3) nodecfgutil lpar -u *user* -p *password* -h *pvmname* -s *serial*
- (4) nodecfgutil {vc|ucs} -u *user* -p *password* -h *hostname* -t *port* -c *protocol*
- (5) nodecfgutil esad -u *user* -p *passwd* -d *domainname* -e *entity* -o *option*

### **-help**

Displays the usage information about the console.

### **lpar|ucs|vc|zone|mscs|esad**

Specifies the virtual or physical environment.

### **-u *user***

Specifies the name of an administrative user.

### **-p *password***

Specifies the password of that user.

### **-h *hostname***

Specifies the name of the server that is managed through the corresponding AIM.

### **-d *domainname***

Specifies the name of the domain that is monitored through the ADES AIM.

**-h *pvmname***

Specifies the name of the IBM PowerVM server (HMC or IVM) that is managed through the LPAR AIM.

**-h *cluster\_name***

Specifies the name of the MSCS cluster.

**-s *serial***

Specifies the HMC/IVM serial number.

**-t *port***

(vCenter, Cisco UCS only) Specifies the port number.

**-c *protocol***

(vCenter, UCS only) Specifies the protocol (HTTP, https).

Return codes: 0 success, -1 failure

**-e *entity***

Specifies the managed entity.

**0**

Specifies the Active Directory for monitoring.

**1**

Specifies the Exchange Server for monitoring.

**2**

Specifies both the Active Directory and Exchange Server for monitoring.

**-o *option***

Specifies the option for providing management.

**0**

Specifies the entire domain for monitoring.

**1**

Specifies a specific host of the domain for monitoring.



**To configure AIMs with NodeCfgUtil in command mode**

1. Open a command prompt on the system on which the AIM resides.

The command prompt appears.

2. Enter one of the following commands:

```
(1) nodecfgutil -help
(2) nodecfgutil {[par|zone|mscs]} -u user -p password -h {pvmname|hostname|mscs}
(3) nodecfgutil lpar -u user -p password -h pvmname -s serial
(4) nodecfgutil {vc|ucs} -u user -p password -h hostname -t port -c protocol
(5) nodecfgutil esad -u user -p password -d domainname -e {0|1|2} -o {0|1}
```

- (1) Displays the usage information about the console.
- (2) Authenticates and stores the passed credentials for Solaris Zones, LPAR, or MSCS.
- (3) Authenticates and stores the passed credentials for LPAR.
- (4) Authenticates and stores the passed credentials for vCenter or Cisco UCS.
- (5) Authenticates and stores the passed credentials for Active Directory and Exchange Server.

The utility writes a configuration file for vCenter Server (vc.cfg), Solaris Zones (zone.cfg), Cisco UCS (ucs.cfg), or LPAR (lpar.cfg) or ADES (esad.cfg) to `sysedge_InstallLpath\plugins\AIPCommon` directory.

## Manage Multiple vCenter AIMs

If you have multiple AIMs with the same configuration, the vCenter PMM selects the first detected AIM configured with the same credentials. If the selected, managed AIM is not the desired AIM, the vCenter Server can be switched. The vCenter PMM performs the following steps depending on whether you add, remove, or switch vCenter Servers:

**Adding vCenter Server**

1. Validating and saving vCenter Server credentials.
2. Detecting AIM with vCenter Server configuration.
3. Saving AIM configuration, and creating required objects.

**Removing vCenter Server**

1. Removing vCenter Server credentials.
2. Deleting objects.
3. Removing AIM configuration.

### Switching vCenter Server

1. Saving AIM configuration to both AIM and database.
2. Deleting and recreating objects based upon the new configuration.

**Note:** A small delay can occur between adding or removing vCenter Server credentials and the time the configuration entry is added or removed.

For a successful deployment, verify that the AIMs are configured and discovered properly prior to updating.

If you did not configure a vCenter Server AIM during installation, perform the following procedure that enables you to manage vCenter Server resources.

### To manage multiple vCenter AIMs

1. Install the vCenter AIM on a server and configure the AIM during installation or later with the NodeCfgUtil utility from the command-line available on that server.
2. Let CA Virtual Assurance discover the system on which the vCenter Server AIM runs.
3. Change to Administration, Configuration, vCenter Server.

The vCenter Server pane and the vCenter AIM Servers pane appear.

4. Add vCenter Servers you want to manage into the vCenter Server pane. Use valid credentials of your vCenter Server environment.

In the vCenter AIM Servers pane the user interface displays a list of one-to-one relationships between available vCenter AIMs and vCenter Servers. If a displayed, managed AIM is not the desired AIM, you can switch the vCenter Server.

## Manage a Renamed vCenter Server

If you have renamed a managed vCenter Server, configure CA Virtual Assurance and the vCenter AIM accordingly in the Administration tab:

### To manage a renamed vCenter Server

1. Discover the renamed vCenter Server.
2. Change to Administration, Configuration, vCenter Server.  
The vCenter Server pane and the vCenter AIM Servers pane appear.
3. Click the plus (+) button to add the renamed vCenter Server to the vCenter Server pane. Use valid credentials of your vCenter Server environment.  
The renamed vCenter Server appears in the list in normal state.
4. Change to the vCenter AIM Servers pane and select the 1:1 relationship between the vCenter AIM Server and the vCenter Server (previous name). Click the Configure button.  
A dialog appears displaying the server names of that relationship.
5. Select the renamed vCenter Server from the drop down menu, and click OK.  
The new relationship is active and CA Virtual Assurance starts to manage the vCenter Server with the new name.

**Important!** The data related to the previous name is removed from the database.

## Requirements for Solaris Zones Management

Verify if the user account that CA Virtual Assurance requires for Solaris Zones management meets the following settings and permissions on the Solaris Server:

- The prompt of the user on the Solaris Server must be "#" (default).
- The Solaris user requires privilege to execute the following commands:
  - zlogin
  - zoneadm
  - zonecfg
- From the global zone the user must have the permission to login to individual Solaris Zones with zlogin and run the following commands:
  - uname -a
  - sar
  - prstat
  - netstat

You add this user name and the corresponding password to CA Virtual Assurance during the product installation or through the NodeCfgUtil.exe utility on the Managed Node on which the Solaris Zones AIM resides.

## Requirements for LPAR Management

Verify that the following components are installed:

- SystemEdge with LPAR AIM on a Windows Server with configured connection to HMC/IVM Server.
- HMC/IVM Server connected to one or more Power Systems

## Configure SNMPv1 Traps by Editing the sysedge.cf File

The sysedge.cf file located in the SystemEDGE\data\port<n> directory contains definitions for SNMPv1 trap communities, which tell SystemEDGE where to send SNMPv1 trap messages. You can configure trap destinations and SNMPv1 communities during SystemEDGE installation or by editing the sysedge.cf file. You can configure the agent to send traps to any number of management systems.

**To configure SNMPv1 traps by editing the sysedge.cf file**

1. Navigate to the *SE\_Install\_Dir\data\port<num>* directory (Windows) or to the *SE\_Install\_Dir/config/port<num>* directory (UNIX, Linux) and create a backup copy of the sysedge.cf file.
2. Open sysedge.cf in a text editor and locate the trap destination section in the upper part of the file.

The trap destination section contains a brief description of trap destinations and communities.

3. Add a line at the end of the trap destination section for each management system to which you want to send SNMPv1 traps. Use the following syntax:

```
trap_community community-name [IP-address | hostname] [port-number]
```

**community-name**

Specifies the SNMP community; for example, public or admin.

**IP-address**

(Optional) Specifies the IP address of the target system.

**Default:** 127.0.0.1

**hostname**

(Optional) Specifies the name of the target system.

**Default:** localhost

**port-number**

(Optional) Specifies a port to send the trap to.

**Default:** 162

4. Save and close the file.
5. Do one of the following:
  - Restart SystemEDGE to activate the changes by restarting the SystemEDGE service from the Windows service control.
  - Restart SystemEDGE to activate the changes by running the following commands from a Windows command prompt.

```
net stop sysedge  
net start sysedge
```

- Restart SystemEDGE to activate the changes by running the following commands from a UNIX or Linux terminal window.

```
/etc/init.d/sysedge restart (Linux, Solaris)  
/etc/init.d/sysedge restart (HP-UX)  
/etc/rc.d/sysedge restart (AIX)
```

### Example

Add the following lines to `sysedge.cf` to send traps with a community-name of `mycommunity` to two systems. The first system has the IP address `192.168.5.26`. The second system is the host `atlanta-noc` and listens on port number `1692`.

```
trap_community mycommunity 192.168.5.26
trap_community mycommunity atlanta-noc 1692
```

**Note:** `sysedge.cf` only defines SNMPv1 trap communities. For information about configuring SNMPv2c or SNMPv3 traps, see the *SystemEDGE User Guide* installed in the `SE_Install_Dir\doc` directory (Windows) or in the `/opt/EMPsysedge/doc` directory (UNIX, Linux).

## Event Forwarding

This section describes how to configure CA Virtual Assurance to forward events to Simple Network Management Protocol (SNMP) management servers or third-party SNMP management servers.

### Configure Windows for SNMP

The SNMP Service and the SNMP Trap Service are installed with Windows, but are not a part of the typical setup. Verify that the SNMP Service and Trap Service are running.

#### To configure Windows for SNMP

1. Click Start, Control Panel, Administrative Tools, Services.  
The Services dialog appears.
2. Perform *one* of the following:
  - If SNMP Service and SNMP Trap Service are listed, continue the Event Forwarding configuration.
  - If SNMP Service and SNMP Trap Service are *not* listed, continue with step 3.
3. Click Start, Control Panel, Add or Remove Programs, and Click Add/Remove Windows Components.  
The Windows Components Wizard dialog appears.
4. Scroll down and select Management and Monitoring Tools, and click Next.  
You are prompted for the Windows installation media.
5. Follow the on-screen instructions to complete the installation.
6. Repeat step 1 to verify that SNMP Service and SNMP Trap Service are running.

## Configure CA Virtual Assurance to Forward Events

Configure CA Virtual Assurance to forward events to a CA or third-party Event Management server.

### To configure CA Virtual Assurance to forward events

1. Open the CA Virtual Assurance user interface.
2. Click Administration.

The Administration page appears.

3. Click Configuration.

The Configuration page appears.

4. Click Event in the left pane.

The Event pane appears.

5. Click + (Add).

The Forwarding and Type fields are automatically populated.

**Note:** If these fields do not populate, restart Apache Tomcat.

6. Select SNMP v1 or v3 from the Type drop-down list, then select Enabled from the Forwarding drop-down list.

If you select SNMP v3 Trap, a panel for additional configuration parameters appears.

7. Enter the management server name in the Server field.
8. Enter a different port number for SNMP or leave the default port 162, which is automatically populated.
  - If you selected SNMP v3, continue with step 9.
  - If you selected SNMP v1, go to step 13.
9. Select Security Level from the drop-down list. The following types are valid options:

#### **NoAuthNoPriv**

No additional security parameters.

#### **AuthNoPriv**

Authentication settings only.

#### **AuthPriv**

Authentication and encryption.

10. Enter the user name for the remote management server in the Username field.
11. Select an Authentication protocol from the options in the drop-down list and specify a password.

### **MD5**

Defines the Message-Digest algorithm 5 cryptographic hash function with 128-bit hash value.

### **SHA**

Defines the Secure Hash algorithm cryptographic hash function.

**Note:** This option is valid for security level AuthNoPriv and AuthPriv only.

12. Select Encryption protocol from the drop-down list. DES and AES are valid options. Specify a password.

**Note:** This option is valid for security level AuthPriv only.

13. Click OK.

A confirmation message appears.

14. Click Save to save the updated Event Forwarding record.

Your settings are updated and the configuration information appears. CA Virtual Assurance is now configured to forward events.

## SNMP V3 Engine ID

The SNMPv3 standard requires each engine (trap sender) to have an ID. Each management platform (target application) honors the engine ID differently. Some management platforms, such as CA Virtual Assurance, require that the actual hexadecimal engine ID is properly configured.

For CA Virtual Assurance, the hexadecimal engine ID is built using a combination of the computer name and the string DCAMTrap as the seed. For example, if the computer name is COMP999, the seed of the engine ID is COMP999-DCAMTrap. CA Virtual Assurance uses the seed in an algorithm to compute the engine ID. The computer name is used to help ensure uniqueness. The hexadecimal value is written to an output file under the CA Virtual Assurance installation directory. The name of the file is dem\_snmp3\_engine\_id.dat. The file is created after you configure CA Virtual Assurance to forward events to another management server.



# Chapter 5: Managing Resources

---

This section contains the following topics:

[Manager Components](#) (see page 57)

[Discover Resources](#) (see page 59)

[Managed Virtual Environments](#) (see page 66)

[Cisco UCS Overview](#) (see page 95)

[Using Multiple Multi-instance AIMs](#) (see page 101)

[Systems Management](#) (see page 101)

## Manager Components

CA Virtual Assurance provides the capability of managing the environments that it monitors and performing actions on their platforms. Data collectors control monitoring and threshold settings for all server agents. Platform Management Modules (PMMs) are web services that poll the data collectors on all managed systems using SNMP. PMMs convert all information received from the data collectors, including MIB object values and traps, into properties that fit into the CA Virtual Assurance architecture and create corresponding managed object representations in the database. Each PMM is specific to a virtual platform. To run a command, a PMM instructs its associated AIM to perform actions.

### More Information

[Application Insight Modules \(AIMs\)](#) (see page 57)

[Monitor Configuration](#) (see page 59)

## Application Insight Modules (AIMs)

Application Insight Modules (AIMs) adds the capability to monitor and manage application-specific events and processes. AIMs are functional extensions to the SystemEDGE.

### AIM for VMware vCenter Server

Provides capabilities to manage and monitor systems that are under VMware vCenter Server control. This AIM can run on any Windows system where SystemEDGE is installed. The AIM communicates with vCenter Server software to get an entire view of all ESX Servers that the associated VMware vCenter Server manages.

### **AIM for IBM PowerVM (LPARs)**

Provides capabilities to manage and monitor the entire system, including LPARs. This AIM can run on any Windows system where SystemEDGE is installed. The AIM communicates with the HMC/IVM through a Secure Shell (SSH) connection, so that the AIM can communicate with the LPARs on POWER systems through the associated HMC/IVM system. Verify that SSH is enabled on the HMC/IVM system and on the Windows sever on which the AIM runs.

### **AIM for Microsoft Hyper-V**

Provides capabilities to manage and monitor Microsoft Hyper-V environments. The SystemEDGE AIM for the Hyper-V server runs on the Hyper-V server.

### **AIM for Microsoft Cluster Services**

Provides capabilities to monitor Microsoft clusters. This AIM can run on any Windows system where SystemEDGE is installed. The AIM communicates with Microsoft Cluster Service to get an entire view of the monitored clusters, nodes, services and applications.

### **AIM for Solaris Zones**

Provides capabilities to manage and monitor Solaris 10 systems configured to run zones. This AIM can run on any Windows system where SystemEDGE is installed. The AIM communicates with the managed Solaris Zones Servers through SSH connections. Verify that SSH is enabled on the managed Solaris servers and on the Windows sever on which the AIM runs.

### **AIM for Cisco Unified Computing System (UCS)**

CA Virtual Assurance interacts with Cisco UCS to query devices and collect statistics. Instead of managing resources as disparate systems, Cisco UCS unifies networking, hardware, storage, and virtualization resources into one cohesive system.

### **AIM for Service Response Monitor**

Provides capabilities to monitor the health and responsiveness of services running on Windows, UNIX, or Linux servers.

### **AIM for Remote Monitoring**

Provides capabilities to manage and monitor remote (agent-less) Windows systems.

### **AIM for Active Directory and Exchange Server**

Provides capabilities to monitor the Active Directory and Exchange Server. The AIM enables Domain and Exchange server Management, maintenance and upgrade.

### **More Information**

[Discover Environments](#) (see page 42)

## Monitor Configuration

You can specify whether to monitor a virtual resource or attribute by changing its monitored state. If you select an object and change its configuration to Unmanaged, the PMM processes the request and sets the value to Unmanaged in the AIM. The current monitor configuration is preserved in the MIB attribute, and the child objects also change to Unmanaged. Traps are generated for the configuration change and the state change for the parent. No traps are generated for the parent object and its children in subsequent polling and recording cycles.

If you select an object and change its configuration to Managed, the PMM processes the request and sets the value to Managed in the AIM. The current monitor configuration is preserved in the MIB attribute and the child objects also changes to Managed. A configuration change trap is generated for the parent. The state of the parent object and its children are evaluated in the next polling and recording cycle and state change traps are generated as needed.

### More information:

[Manage a Resource](#) (see page 62)

## Discover Resources

You can discover and add servers or entire subnets that you want to manage. Servers that were previously unmanaged or new servers that were added to the data center can be discovered.

**Note:** CA Virtual Assurance discovery requires hostname resolution. If the IP address for a discovered server changes, CA Virtual Assurance does not automatically resolve the IP address. As a result, the discovery profile fails to update the Management DB. If the IP address changes, rediscover the server.

## Discover a System

You can specify a single system to be discovered, managed, and optionally assigned to a service.

**Note:** SystemEDGE or CA Performance Lite Agent cannot be installed on HMC/IVM servers and therefore no HMC/IVM system data can be collected for IBM Power Series servers. See also the "Configuring Resources" chapter in the *Administration Guide*.

### To discover a system

1. Select Resources, Manage, Discover System.

2. Complete the following fields:

**System Name**

Specifies the name or IP address of the server.

3. (Optional) Click Next.

The Enhanced Discovery and SNMP Information dialog appears.

4. (Optional) Enable Enhanced Discovery to perform detailed discovery using SoftAgent technology.
5. (Optional) Enable Override SNMP Defaults to perform detailed discovery using SNMP information.
6. Click Finish.

A success message appears when the system is discovered. Discovered servers are automatically managed, but servers in a subnet discovery are not managed.

## Discover a Network

You can specify a segment of networks to be discovered. CIDR (Classless Inter-Domain Routing) notation is used when specifying IP addresses for network discovery in the user interface. This notation consists of an address and the number of bits to use as the subnet prefix, as shown in the following example:

172.24.143.0/24

You can also use wildcards and ranges:

172.24.143.\*

172.24.143.{1-255}

**To discover a network**

1. Select Resources, Manage, Discover Network.
2. Complete the following fields.

**Network Name**

Specifies the name of the network.

**Network Address**

Specifies the network IP address. Hover the cursor over this field to display address examples.

**Exclude Address**

(Optional) Specifies the network address that you want to exclude from discovery.

3. Select *one* of the following options for Discovery Method and complete the corresponding fields:

**Ping Sweep**

Discovers all IP addresses in the network.

**DNS**

Discovers host names registered in the Domain Name System (DNS) server.  
Type the domain name and the DNS server name in the fields.

4. (Optional) Click Next.  
The Enhanced Discovery and SNMP Information dialog appears.
5. (Optional) Enable Enhanced Discovery to perform detailed discovery using SoftAgent technology.
6. (Optional) Enable Override SNMP Defaults to perform detailed discovery using SNMP information.
7. Click Finish  
A success message appears when the network is discovered. Discovered servers are automatically managed, but servers in a subnet discovery are not managed.

## Delete a Network

You can delete a discovered network. However, systems that are already discovered retain their status.

**To delete a network**

1. Select Resources, Manage, Manage Discovered Networks.  
A list of discovered networks appears in the right pane.
2. Select a network to delete and click - (Delete) on the Network List toolbar.  
A message prompts for confirmation.
3. Click OK.  
The network is deleted from the network list.

## Cancel Network Discovery

You can cancel a network discovery that is in progress.

### To cancel network discovery

1. Click Resources and open the Manage pane.
2. In the Management section, click Manage Discovered Networks.
3. Select the In Progress network for which to cancel discovery, then click — (Cancel) on the Network List toolbar.

A message confirms that discovery of the selected network is canceled.

## Manage a Resource

You can monitor the performance of discovered resources by adding them to the list of Managed resources.

### To manage a resource

1. Click Resources, and open the Manage pane.
2. In the Management section, click Manage Systems.
3. Select Unmanaged Servers from the drop-down list.

The list of unmanaged resources appears.

4. Select the resource that you want to manage, then click Manage in the Actions drop-down menu.

A message confirms that the resource is Managed.

5. Expand the Managed folder.

The resource appears in the Managed list, and metric collection starts with the next recording cycle.

## Create a Service

You can organize the servers that you monitor into logical services that reflect the resources required by your business needs.

### To create a service

1. Click Resources, and open the Explore pane.

2. Select a parent service node, such as Data Center or CA Virtual Assurance Services.
3. Right-click Management, New Service.

The Service: New dialog appears.

4. Enter a name for the new service in the Service Name field and set a priority level in the Service Priority field.

**Note:** The following characters are not supported for service names: % " " ' ' < > / \ : ` ~ ;

#### Service Priority

Specifies the order in which to run actions in a single poll cycle.

##### Example:

ServiceA: Priority 3

ServiceB: Priority 1

ServiceC: Priority 2

When all of their respective rules evaluate as true, the actions run in the following order: ServiceB, ServiceC, ServiceA.

5. Change the Lag occurrence or accept the default provided.

#### Lag

Defines how often the rule must evaluate as true before the action triggers.

6. Change the Lower and Upper Threshold percentages or accept the defaults.

#### Lower and Upper Threshold %

Specifies the lower and upper thresholds of the entire service.

**Limits:** Only the overall utilization metric can be evaluated at the service level.

7. You can assign the service to a CCA server. A CCA service with the same list of servers is created automatically.

**Note:** Servers that the CCA server does not discover are not added to the CCA service. Review the Events table for possible issues.

You can also select a management profile to be applied to all servers within the service.

8. Select the servers for the new service from the Available Servers list in the Servers section, then click the right arrows.

**Note:** If your list of available servers is lengthy, filter the list to reduce the set of servers. To do so, click the Filter arrow, enter your filter criteria, and click Search.

The servers are added to the Selected Servers section.

9. Click Save on the Actions drop-down menu.

The new service is saved and appears in the Explore pane.

On a service level, you can take snapshots, view components, run discovery or change detection. Right-click a service and select the relevant option.

## Edit a Service

You can edit an existing service to rename it, to change settings, or to add or remove resources in the group.

### To modify a service

1. Click Resources, and open the Explore pane.
2. Select the service , and right-click Management, Edit Service.

The Service: Edit dialog appears.

3. Change the priority level in the Service Priority field and the Lower and Upper Threshold percentages, as necessary.

#### Service Priority

Specifies the order in which to run actions in a single poll cycle.

##### Example:

ServiceA: Priority 3

ServiceB: Priority 1

ServiceC: Priority 2

When all of their respective rules evaluate as true, the actions run in the following order: ServiceB, ServiceC, ServiceA.

4. Change the Lag occurrence, or accept the default.

#### Lag

Defines how often the rule must evaluate as true before the action triggers.

#### Lower and Upper Threshold %

Specifies the lower and upper thresholds of the entire service.

**Limits:** Only the overall utilization metric can be evaluated at the service level.

5. Select the servers to add to the service from the Available Servers list in the Servers section, then click the right arrows.

**Note:** If your list of available servers is lengthy, filter the list to reduce the set of servers. To do so, click the Filter arrow, enter your filter criteria, and click Search.

The servers are added to the Selected Servers section.



6. Select the servers to remove from the service from the Selected Servers list in the Servers section, then click the left arrows.

The servers move from the Selected Servers section to the Available Servers section.

7. Click Save.

The Servers lists are updated.


## Managed Virtual Environments

CA Virtual Assurance is designed to manage virtual environments, but it also discovers and manages hosts running Linux, UNIX, or Windows operating systems. The following list outlines the management capabilities of this product:

- Management of vCenter Server, Solaris Zones, Hyper-V, Cisco UCS, and LPAR environments. The required CA Virtual Assurance components reside on one or more Windows hosts (separate Windows AIM Servers or CA Virtual Assurance manager). Required components on these hosts are:
  - Application Insight Modules (AIMs) for vCenter Server, Solaris Zones, Hyper-V, Cisco UCS, or LPAR
  - SystemEDGE
- Management of hosts which run AIX, HP-UX, Linux, Solaris, or Windows operating systems. Recommended components on these hosts are:
  - SystemEDGE, Advanced Encryption, Remote Monitoring AIM (Windows Servers only), Service Response Monitoring (SRM) AIM

**Note:** For details about platform-specific system requirements, see the *Release Notes*.

### Supported Virtual Environments and Operating Systems



#### More Information

[VMware vCenter Server](#) (see page 67)  
[Solaris Zones](#) (see page 82)  
[Microsoft Hyper-V Server](#) (see page 84)  
[IBM PowerVM \(LPAR\)](#) (see page 86)  
[Microsoft Cluster Service](#) (see page 90)  
[Cisco UCS](#) (see page 93)

## VMware vCenter Server


VMware vCenter Server provides the central point of control for configuring, provisioning, and managing virtualized IT environments. SystemEDGE and the AIM for VMware vCenter Server run on the CA Virtual Assurance manager server or on an arbitrary Windows server. The AIM remotely communicates with one or more VMware vCenter Servers through web-services and with the vCenter Server PMM on the manager server through SNMP.

The VMware vCenter Server PMM provides connection and operational support for all VMware vCenter Server operations. The PMM is responsible for managing connections, performing VM-related operations, and populating the database with data retrieved from VMware vCenter Server. The provisioning service performs VMware vCenter Server operations including cloning, power operations, resource and share adjustments, and snapshot management.

## Interactions Between vCenter Server Management Components

The following diagram illustrates how the components involved in vCenter Server management interact. SystemEDGE and the vCenter Server AIM run on a Windows server. The AIM communicates with one or more remote vCenter Servers to manage the virtual environment. The vCenter Server AIM collects the data for an entire view of the physical and virtual resources associated with the vCenter Server.

**Interaction Between vCenter Server Management Components**



You can configure vCenter Server management during installation or through the Administration tab of the user interface.

**Note:** VMware Tools optimize the virtualization of VMs and it is recommended that they are installed on each VM in your VMware environment. Some features of this product are not available or do not function correctly for VMs that do not have VMware Tools installed. For this reason, VMs that do not have VMware tools installed are not supported.

### More Information

[Manage Multiple vCenter AIMs](#) (see page 49)

## Remote and Multi-instance vCenter Server Support

The vCenter Server AIM communicates with one or more remote vCenter Server instances. However, when you use a CA Virtual Assurance manager with multiple remote vCenter Server AIMs to manage multiple vCenter environments, consider the following:

- Verify that each vCenter Server is uniquely associated with one AIM.
- Use Administration on the CA Virtual Assurance manager user interface or NodeCfgUtil on the appropriate AIM server to remove any ambiguous associations.

If more than one AIM manages a particular vCenter Server, management, discovery, or monitoring issues can occur.

**Note:** When you run the vCenter Server AIM without the CA Virtual Assurance manager in an eHealth, Spectrum Infrastructure Manager, or NSM environment, the AIM supports single-instance mode only.

## Hot-plug Support for VMs

CA Virtual Assurance detects if the hot plug option is enabled for VMs. CA Virtual Assurance supports the following adjustments for hot plug-enabled VMs while the VM is powered on.

- Adding vCPU
- Adding vRAM

**Note:** How to enable or disable the hot plug option, see the *VMware vSphere Virtual Machine Administration Guide*.

## Device Management for VMs

Device management includes the following:

- Adding and removing vDisks
- Adding and removing vNICs

## vNetwork Standard Switches (vSwitch)

CA Virtual Assurance monitors policies and properties of standard vSwitches which are abstracted network devices. A vSwitch can route traffic internally between VMs and link to external networks. vSwitches combine the bandwidth of multiple network adapters and balance communications traffic among them. A vSwitch can handle physical NIC failover.

A vSwitch models a physical Ethernet switch. The default number of logical ports for a vSwitch is 120. You can connect one network adapter of a VM to each port. Each uplink adapter associated with a vSwitch uses one port. Each logical port on the vSwitch is a member of a single port group. Each vSwitch can also have one or more port groups assigned to it. When two or more VMs are connected to the same vSwitch, network traffic between them is routed locally. If an uplink adapter is attached to the vSwitch, each VM can access the external network that the adapter is connected to.

## Distributed Virtual Switches (DVS)

CA Virtual Assurance supports the following Distributed Virtual Switches in a vSphere environment:

- VMware vNetwork Distributed Switch (vDS, vSphere component)
- Cisco Nexus 1000V Switch (integrates with vSphere)

CA Virtual Assurance discovers Distributed Virtual Switches in a vSphere environment and monitors its policies and properties through events. CA Virtual Assurance VM provisioning supports vNetwork Distributed Switches and Cisco Nexus 1000V Switches.

A DVS operates as a single virtual switch that spans across all hosts which are associated with that switch. A DVS represents the same switch (same name, same network policy) and port group for these hosts. These properties allow VMs to maintain a consistent network configuration as they migrate among multiple hosts.

Like a vNetwork Standard Switch, each DVS is a network hub that VMs can use. A DVS can forward traffic internally between VMs or link to an external network by connecting to physical NICs (uplink adapters).

Distributed Virtual Port Groups (dvPort Groups) are port groups associated with a DVS and specify port configuration options for each member port. dvPort Groups define how a connection is made through the DVS to the network.

Distributed Virtual Uplinks (dvUplinks) provide a level of abstraction for the physical NICs (vmnics) on the ESX or ESXi hosts. Each physical NIC is mapped to a dvUplink. The mapping from the dvPort Group to the dvUplink defines which physical NICs on ESX or ESXi hosts are used by VMs to get access to the network through the DVS.

The Cisco Nexus 1000V Switch consists of the Virtual Ethernet Module (VEM) and the Virtual Supervisor Module (VSM). On each ESX or ESXi host associated with a Cisco Nexus 1000V Switch, the VEM replaces the VMware vSwitch and runs as a module in the hypervisor kernel. The VSM controls multiple VEMs as one logical switch and runs in a VM on an ESX or ESXi host.

For further details, see the VMware vNetwork Distributed Switches documentation at <http://pubs.vmware.com> or the Cisco Nexus 1000V Switch documentation at <http://www.cisco.com/go/1000vdocs>.

**Note:** If you use the Cisco Nexus 1000V Switch, the VSM VM does not appear as a special VM in the CA Virtual Assurance user interface. Verify that your rules and actions that you apply to the VSM VM do not affect the Cisco Nexus 1000V Switch.



## Monitor Distributed Virtual Switches Through Events

You can monitor Distributed Virtual Switches through the following events:

- Add switch:  
Distributed Virtual Switch VM-dvSwitch added to Datacenter MyDC. vSphere: vcserver.mycomp.com
- Delete switch:  
Distributed Virtual Switch VM-dvSwitch removed from Datacenter MyDC. vSphere: vcserver.mycomp.com
- Add Port Group:  
Distributed Virtual Port Group VM dvPortGroup added to Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com
- Remove Port Group:  
Distributed Virtual Port Group VM dvPortGroup removed from Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com
- Add Uplink:  
Distributed Virtual Uplink VM DVUplink added to Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com
- Remove Uplink:  
Distributed Virtual Uplink VM DVUplink removed from Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com

### To monitor Distributed Virtual Switches through events

1. Click the Dashboard tab, scroll to the Events panel, and click the Show Table Filter icon.  
The Filter panel opens.
2. Specify an appropriate filter for the Distributed Virtual Switch events that you want to monitor and click Apply.  
The Events panel lists the filtered events.

### vApp Support

A *vApp* is a specific resource pool which treats a collection of VMs as a single unit. vApp uses the Open Virtualization Format (OVF). OVF is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels associated with it. CA Virtual Assurance can perform operations on a vApp. An operation on a vApp is propagated to all VMs in the vApp.

You can partition any vApp into smaller vApps to divide and assign resources to specific groups or for specific purposes. You can add resources like VMs, Resource Pools, or vApps to an existing vApp. You can also hierarchically organize and nest vApps.

CA Virtual Assurance supports the following operations on vApps:

- Provision VMware vApp
- Add Resource
- Clone vApp
- Power On vApp
- Power Off vApp
- Suspend vApp
- Delete from VMware vCenter
- Unregister from VMware vCenter

### Monitor vApps Through Events

You can monitor vApps through the following events:

- Add vApp:  
vApp *MyvApp* added to parent resource pool resources. vSphere  
*vcserver.mycomp.com*
- Delete vApp:  
vApp *MyvApp* removed from parent resource pool resources. vSphere,  
*vcserver.mycomp.com*

The following traps are available:

- ResPoolvAppAddedTrap: Add vApp to resource pool or vApp.
- ResPoolvAppRemovedTrap: Remove vApp from resource pool or vApp.
- ResPoolvAppVCConfigChangeTrap: Configuration data for vApp entity in vApp has changed.
- VMAddedTovAppTrap: VM added to vApp.
- VMRemovedFromvAppTrap: VM removed from vApp.
- VMvAppVCConfigChangeTrap: Configuration data for VM entity in vApp has changed

#### To monitor vApps through events

1. Click the Dashboard tab, scroll to the Events panel, and click the Show Table Filter icon.

The Filter panel opens.

2. Specify an appropriate filter for the vApp events that you want to monitor and click Apply.

The Events panel lists the filtered events.

## vCenter Server in a Cluster

If vCenter Server resides in a cluster, the vCenter Server AIM must run outside of this cluster. Configure the vCenter Server AIM to point to the cluster host. The AIM can detect a failover and repopulates its internal cache when vCenter Server is successfully started.

## Fault Tolerance

VMware vSphere lets you enable *Fault Tolerance (FT)* on a VM defined to a cluster which is configured for High Availability (HA). Fault Tolerance creates a secondary VM on another ESX Server in the cluster. The secondary VM operates in lock-step mode with the primary VM that is executing the workload. If there is a failure, the secondary VM immediately takes over the workload execution from the point of failure. CA Virtual Assurance discovers and manages primary and secondary VMs in a cluster.

Regarding VM management, CA Virtual Assurance treats the primary and secondary VM as a single VM, with fault tolerance enabled, and displays its fault tolerant properties. The primary VM appears on the left pane (first class object) and provides its FT properties in the right pane. The secondary VM properties (second class object) are listed in the right pane only. You cannot perform VM operations like start, stop, or clone on secondary VMs.

The number of VMs represented in the General Information panel is based on the running count of non-FT VMs plus primary FT VMs. Secondary FT VMs are not included in the overall total count of VMs.

CA Virtual Assurance gathers FT VM data on various levels in the environment.

### Fault Tolerance Requirements

When a VM is Fault tolerant, the following operations must be disabled:

- Clone VM
- Remove from Inventory (unregister)
- Snapshot
- Convert to template

### VM Fault Tolerance Properties

For each VM CA Virtual Assurance displays:

#### **Fault Tolerance Status**

Indicates the VM fault tolerance status.

##### **Not Fault Tolerant**

Indicates that the VM is not fault tolerant.

##### **Protected**

Indicates that the VM is fault tolerant and protected.

##### **Not Protected (Starting)**

Indicates that the fault tolerance is starting and the VM is not protected.

##### **Not Protected (Need Secondary VM)**

Indicates that the fault tolerance is enabled but needs secondary VM.

##### **Not Protected (Disabled)**

Indicates that the fault tolerance is disabled and the VM is not protected.

##### **Not Protected (VM Not Running)**

Indicates that the fault tolerance is enabled but the VM is not running.

#### **Secondary VM Location**

Identifies the secondary host location.

## ESX Host Fault Tolerance Attributes

ESX Host Fault Tolerance attributes are as follows:

### **Fault Tolerance**

Identifies whether the host has the fault tolerance enabled.

### **Fault Tolerance version**

Identifies the version of Fault Tolerance running on the host.

**Note:** Only hosts with the same version of Fault Tolerance are compatible with one another.

### **Total Primary VMs (calculated by the AIM)**

Indicates the total number of primary VMs configured to this host.

### **Total Secondary VMs (calculated by the AIM)**

Indicates the total number of secondary VMs configured to this host.

### **Powered on Primary VMs (calculated by the AIM)**

Indicates the total number of primary VMs running (powered on) on this host.

### **Powered on Secondary VMs (calculated by the AIM)**

Indicates the total number of secondary VMs running (powered on) on this host.

## Logical Volume Management in VMs

CA Virtual Assurance supports management of logical volumes in virtual disks. For example, you can manage the C: drive in a VM.

## VMware vCenter Server Architecture

The following major components comprise the VMware vCenter Server architecture:

### **VMware vCenter Server**

Provides the central point of control for configuring, provisioning, and managing virtualized IT environments. vCenter Server runs as a service on Microsoft Windows platforms.

**Note:** For information about vCenter Server system requirements, see the *Release Notes*.

### **Database**

Stores persistent information about the physical servers, resource pools, data centers, and virtual machines managed by the vCenter Server.

### **VMware vCenter Server Agent**

Connects ESX Servers with a vCenter Server.

### **Virtual Infrastructure Client**

Provides remote connections to the vCenter Server or to individual ESX Servers from any Windows computer.

### **Virtual Infrastructure Web Access**

Connects with the vCenter Server without installing a client.

In this environment, the CA Virtual Assurance AIM detects the logical and physical relationships between the components of the virtualized environment. The AIM provides a view of the entire virtualized environment.

The AIM for vCenter Server monitors the following resource types:

#### **vCenter Server**

Provides information about the health status of the vCenter Server computer. For example, status and data about CPU, datastore, and memory usage.

#### **Data Center**

Serves as a container for your hosts, virtual machines, resource pools, or clusters. Data centers can represent organizational structures if their virtual configurations meet the requirements of specific departments. You can also use data centers to create isolated virtual environments for testing or organization purposes.

#### **ESX Host**

Represents all computing and memory resources of a physical server on which an ESX Server runs.

#### **Resource Pool**

Defines partitions of physical computing and memory resources of a single host or a cluster. You can partition any resource pool into smaller resource pools to divide and assign resources to specific groups or for specific purposes. You can also hierarchically organize and nest resource pools.

#### **Virtual Machine**

Specifies virtualized x86 environments in which guest operating systems and applications can run. When you create a virtual machine, it is assigned to a particular host, cluster, or resource pool, and to a datastore. A virtual machine consumes resources dynamically on its physical host, in the same manner a physical device consumes energy dynamically depending on its workload.

**Datastore**

Specifies a virtual representation of combinations of underlying physical storage resources in a data center. Local disks on a server provide the physical storage resources, using SAN disk arrays, and so on.

**Virtual Disk**

Defines the disk drive in a virtual guest operating system. A virtual disk is a specific file or a set of files that reside on the local host or on a remote file system. The virtual disk behaves like a physical disk drive in an operating system.

**Virtual Switch**

Works like a physical switch. Each ESX Server has its own virtual switches that connect to virtual machines through port groups. These virtual switches also have uplink connections to the physical Ethernet adapters on the ESX server. Virtual machines communicate with the outside world through physical Ethernet adapters connected to virtual switch uplinks.

**Physical NIC**

Specifies a physical Ethernet adapter on an ESX Server.

**Virtual NICs**

Specifies a virtual Ethernet adapter on a virtual machine. The guest operating system communicates with the virtual Ethernet adapter through a device driver as if the virtual Ethernet adapter was a physical Ethernet adapter. The virtual Ethernet adapter has its own MAC address, one or more IP addresses, and responds to the standard Ethernet protocol exactly like a physical NIC.

**Hardware Sensors**

Provide physical information about the CPU, memory, fan, voltage, storage, temperature, and power. Hardware sensors can be accessed in ESX servers through vCenter Server.

## Resource Allocation

If available resource capacity does not meet the demands of the resource consumers, customize the amount of resources for virtual machines, vApps, and resource pools.

Use the settings for shares, reservation, and limit to determine the amount of CPU and memory resources provided for virtual machines, resource pools, or vApps.

## Resource Allocation Shares

Shares specify the relative priority or importance of a virtual machine, resource pool, or vApp regarding to its siblings. If a virtual machine has twice as many shares of a resource as another competing virtual machine, it can consume twice as much of that resource.

Shares are typically specified as natural numbers. You can use defaults or assign a specific number of shares (proportional weight) to each virtual machine.

Specifying shares makes sense only with regard to sibling virtual machines, vApps, or resource pools. Sibling virtual machines or resource pools have the same parent in the hierarchy. Siblings share resources according to their relative share values, bounded by the reservation and limit. When you assign shares to a virtual machine, you always specify the priority for that virtual machine relative to other powered-on virtual machines.

For example, when competition occurs, a virtual machine with 2000 shares receives more CPU time than a virtual machine with 1000 shares. Shares are configured relative to the other shares; thus, only the proportion of shares matters, not the values of the shares. Three virtual machines with share values of 1000, 2000, 3000 act the same as three virtual machines with share values of 1, 2, 3. You can use any number scheme you prefer. If you leave ample space between the numbers, you can easier add resources to your resource pool in the future.

When there is no competition between resources, shares do not affect the operations of the virtual machines. Specifying shares help you to balance out your resource pools or vApps.

## Resource Allocation Reservation

A reservation specifies the guaranteed minimum CPU or memory allocation for a virtual machine, resource pool, or vApp. vSphere allows you to power on a virtual machine only if there are enough unreserved resources available for the virtual machine. The server guarantees that amount of reserved resources even when the physical server is heavily loaded. The reservation is defined in megahertz or megabytes.

For example, assume you have 2GHz CPU available. Then specify a reservation of 1000 MHz for VM1 and 1000 MHz for VM2. Now each virtual machine is guaranteed to get 1GHz if necessary. However, if VM1 is using only 500MHz, VM2 can use 1.5GHz.

The reservation default is 0. You can specify a reservation to guarantee that the minimum required amounts of CPU or memory are always available for the virtual machine.



## Resource Allocation Limit

A limit specifies the maximum for CPU or memory allocation for a virtual machine, resource pool, or vApp. A server can allocate more than the reservation to a virtual machine, but never allocates more than the limit. Unutilized CPU or memory on the system is not allocated beyond the limit. The limit is defined in megahertz or megabytes.

CPU and memory limit defaults are set to unlimited. When the memory limit is set to unlimited, vSphere effectively determines the amount of memory when it creates a virtual machine. Usually, it is not necessary to specify a limit.

## Resource Allocation Best Practices

Specify resource allocation settings (shares, reservation, and limit) that are appropriate for your ESX/ESXi environment.

The following guidelines can help you achieve better performance for your virtual infrastructure.

- If you expect frequent changes to the total available resources, use shares to allocate resources across virtual machines. If you use shares, and you upgrade the host, the number of shares do not change. For example, each virtual machine stays at the same priority even though each share represents a larger amount of memory or CPU.
- Use reservation to specify the minimum acceptable amount of CPU or memory, not the amount you want to have available. The host assigns additional resources as available based on the number of shares, estimated demand, and the limit for your virtual machine. The amount of resources specified by a reservation does not change when you modify the environment, such as by adding or removing virtual machines.
- When specifying the reservations for virtual machines, do not commit all resources and plan to leave an appropriate part unreserved. When moving closer to fully reserving all system capacity, it becomes increasingly difficult to change reservations and the resource pool hierarchy.
- For further details, see the vSphere documentation at [www.vmware.com](http://www.vmware.com).

## Solaris Zones

A Solaris Zone defines a virtualized operating system that provides an isolated, secure environment in which to run applications. This environment 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 resource management of the operating system. The Solaris Zones PMM provides health monitoring, management, and provisioning of Solaris Zones environments.

Solaris Zones Container resources can be managed at three levels:

### **Solaris Zones Zone Management**

Solaris servers use *zones* to run applications in isolated environments to make it appear as if they are running on physically separate computers. Each zone on a server takes its resources from a resource pool and includes virtual network interfaces, file systems, memory, and other dedicated units.

### **Solaris Zones Project Management**

A *project* is an application or set of applications that you want to divide into a separate workload entity. A zone allocates resources to a project separately from other resources or projects in the zone, according to workload and configuration settings.


### **Solaris Zones Resource Pool Management**

*Resource pools* provide a persistent configuration mechanism for processor set configuration and scheduling class assignment. Resource pools can dynamically allocate resources to projects and tasks in a zone according to how they are configured.

## Interaction Between Solaris Zones Management Components

The following diagram illustrates how the components involved in Solaris Zones management interact. The managed node is a Windows server on which SystemEDGE and the Solaris Zones AIM run. The communication between the AIM and the Solaris Zones servers is based on SSH (Secure Shell).

### Interaction Between Solaris Zones Management Components



To add the required connection information for each Solaris Zones Server, use a custom installation, the Administration tab of the user interface, or the NodeCfgUtil.exe utility on the managed node. The connection information is written to the configuration file on the managed node. The AIM polls the configuration file and starts monitoring your Solaris Zones environment.

## Microsoft Hyper-V Server

Windows Server 2008 R2 Hyper-V, the hypervisor-based server virtualization technology, is available as an integral feature of Windows Server 2008 R2 that enables you to implement server virtualization. The SystemEDGE AIM for Hyper-V server runs on the Hyper-V Server computer.

The Hyper-V Server PMM provides connection and operational support for all Hyper-V Server operations. The PMM is responsible for managing connections, performing VM-related operations, and populating the database with data retrieved from Hyper-V Server.

The AIM for Hyper-V Server monitors the following resource types:

### **Hyper-V Server**

Represents all computing and memory resources of a physical server on which Hyper-V runs. The Hyper-V AIM provides information about the health status of the Hyper-V Server computer. For example, status and data about CPU and memory usage.


### **Virtual Machine**

Specifies virtualized x86 environments in which guest operating systems and applications can run. When you create a virtual machine, it is assigned to a particular host, cluster, or resource pool, and to a data store. A virtual machine consumes resources dynamically on its physical host, in the same manner as a physical device consumes energy dynamically depending on its workload.

## Interactions Between Hyper-V Server Management Components

The following diagram illustrates how the components involved in Hyper-V management interact. SystemEDGE and the Hyper-V AIM run on the Windows 2008 (Hyper-V) Server to manage the virtual environment. The Hyper-V AIM collects the data for an entire view of the physical and virtual resources associated with the Hyper-V Server.

**Interaction Between Hyper-V Server Management Components**



You can configure Hyper-V management during installation or through the Administration tab of the user interface.

## IBM PowerVM (LPAR)

On PowerVM systems, IBM AIX provides the ability to divide systems into logical partitions (LPARs). Each logical partition runs as an independent system, and you can distribute resources among partitions. Typically each system has a specialized partition named Virtual I/O Server (VIOS) which virtualizes disk resources and network interfaces. Partitioning a system lets you account for separate computing needs while sharing virtualized resources dynamically. PowerVM systems have a Virtualization Manager Component that can either be the Hardware Management Console (HMC) or the Integrated Virtualization Manager (IVM). HMC is an appliance that runs on a separate system and is used to manage multiple PowerVM servers. IVM is an extension to the Virtual I/O Server and can only manage the local PowerVM server.

The LPAR AIM enables SystemEDGE to monitor PowerVM resources.

The LPAR Performance Management Module (PMM) provides connection and operational support for all LPAR operations. The PMM is responsible for managing connections and retrieving data from the Hardware Management Console (HMC) or Integrated Virtualization Manager (IVM), performing various LPAR-related operations, populating the database, and providing web services/ssh for all HMC/IVM interaction.

You can retrieve managed system and LPAR data from the HMC/IVM and perform the following LPAR-related operations:

### Server level

On the server level, you can perform the following tasks:

- Provision LPARs
- Delete LPARs

### Power operations level

On the power operations level, you can perform the following tasks:

- Activate LPARs
- Shutdown LPARs
- Restart LPARs

### Resource adjustments level


On the resource adjustments level, you can perform the following tasks:

- Add LPAR processor and memory units
- Subtract LPAR processor and memory units

## Interaction Between AIX LPAR Management Components

The following diagram illustrates how the components involved in IBM LPAR management interact. The AIM Server is a Windows Server on which SystemEDGE and the LPAR AIM run. The communication between the AIM and the HMC/IVM Server is based on SSH (Secure Shell). Because CA Virtual Assurance can connect to multiple HMC or IVM Servers, CA Virtual Assurance gains an overall view of your LPAR environment.

### Interaction Between PowerVM Management Components



To add the required connection information for each required HMC/IVM and Virtual I/O Server, use one of the following methods:

- During custom installation
- Administration tab of the user interface (see *Online Help*)
- NodeCfgUtil.exe utility on the AIM Server

The connection information is written to the configuration file on the managed node. The LPAR AIM polls the configuration file and starts monitoring your LPAR environment through HMC/IVM.

## LPAR Monitoring

To monitor LPAR resources, create SystemEDGE monitors based on the LPAR AIM MIB and the SystemEDGE Component Object Model in the sysedge.cf file without using UI functionality. Use appropriate object classes and specify object instances according to LPAR resources. The created monitored LPAR objects propagate their state to the computer system where the LPAR AIM is installed. We recommend that you provide HMC, POWER5/POWER6/POWER7 and LPAR system information in the monObjInstance attribute, similar to the following example.

### Example

The following monitor definitions for the sysedge.cf file are set up to watch the Alive status of a POWER5 or POWER6 system named *powersys*. An LPAR named *lpar01* is set to be greater than 2, that is, warning-3, minor-4, and so on.

```
monitor oid monCurrState.53001 98 0x0 60 absolute > 2 'Lpar System status' "'System' 'hmc/powersys/Total'
Alive critical
monitor oid monCurrState.53006 99 0x0 60 absolute > 2 'Lpar01 System status' "'System'
'hmc/powersys/lpar01/Total' Alive critical
```

**Note:** The instance name of a monitor must not begin with lpar://

The following table shows an example of the Self Monitor table that corresponds with the monitor definition examples for the sysedge.cf file.

| mon Index | monOID                           | mon ObjClass | monObjInstance             | mon ObjAttribute | mon Severity | mon CurrState |
|-----------|----------------------------------|--------------|----------------------------|------------------|--------------|---------------|
| 530001    | lparAimStatSys Status.1          | System       | lpar://hmc/powersys /Total | Alive            | critical     | ok            |
| 530002    | lparAimStatSys CPUUsage PerMil.1 | CPU          | lpar://hmc/powersys /Total | PercentUsed      | warning      | ok            |




| <b>mon<br/>Index</b> | <b>monOID</b>                             | <b>mon<br/>ObjClass</b> | <b>monObjInstance</b>                | <b>mon<br/>ObjAttribute</b> | <b>mon<br/>Severity</b> | <b>mon<br/>CurrState</b> |
|----------------------|-------------------------------------------|-------------------------|--------------------------------------|-----------------------------|-------------------------|--------------------------|
| 530003               | lparAimStatSys<br>CPUUsage<br>PerMil.1    | CPU                     | lpar://hmc/powersys<br>/Total        | PercentUsed                 | minor                   | ok                       |
| 530004               | lparAimStatSys<br>MemoryUsage<br>PerMil.1 | Memory                  | lpar://hmc/powersys<br>/Total        | PercentUsed                 | warning                 | warning                  |
| 530005               | lparAimStatSys<br>MemoryUsage<br>PerMil.1 | Memory                  | lpar://hmc/powersys<br>/Total        | PercentUsed                 | minor                   | minor                    |
| 530006               | lparAimStatLP<br>Status.1.1               | System                  | lpar://hmc/powersys<br>/lpar01/Total | Alive                       | critical                | critical                 |
| 530007               | lparAimStatLPCPU<br>Usage.1.1             | CPU                     | lpar://hmc/powersys<br>/lpar01/Total | PercentUsed                 | warning                 | ok                       |
| 530008               | lparAimStatLPCPU<br>Usage.1.1             | CPU                     | lpar://hmc/powersys<br>/lpar01/Total | PercentUsed                 | minor                   | ok                       |
| 530009               | lparAimStatLP<br>MemoryUsage.1.1          | Memory                  | lpar://hmc/powersys<br>/lpar01/Total | PercentUsed                 | warning                 | ok                       |
| 530010               | lparAimStatLP<br>MemoryUsage.1.1          | Memory                  | lpar://hmc/powersys<br>/lpar01/Total | PercentUsed                 | minor                   | ok                       |
| 530011               | lparAimStatLP<br>Status.1.2               | System                  | lpar://hmc/powersys<br>/lpar02/Total | Alive                       | critical                | critical                 |
| 530012               | lparAimStatLPCPU<br>Usage.1.2             | CPU                     | lpar://hmc/powersys<br>/lpar02/Total | PercentUsed                 | warning                 | ok                       |
| 530013               | lparAimStatLPCPU<br>Usage.1.2             | CPU                     | lpar://hmc/powersys<br>/lpar02/Total | PercentUsed                 | minor                   | ok                       |
| 530014               | lparAimStatLP<br>Memory<br>Usage.1.2      | Memory                  | lpar://hmc/powersys<br>/lpar02/Total | PercentUsed                 | warning                 | ok                       |

## Microsoft Cluster Service

The Microsoft Cluster Service (MSCS) connects two or more servers together so that they appear as a single computer to clients. Clustering helps you to have a fail-safe application. A cluster-aware application like Microsoft SQL Server runs on a node at a time. If that node goes down, some other node takes over the service. Clustering also helps in making sure that your application is up all the time.

### Interaction Between MSCS Management Components



Performance monitoring requires remote access to clusters and individual cluster nodes for metric collection such as CPU and memory use. The cluster-specific information is available on each node. The MSCS AIM uses WMI (port 135) to communicate with clusters.

## Microsoft Cluster Overview

The Microsoft Cluster Service (MSCS) connects two or more servers together and shows them as a single computer to clients. Clustering helps you to have a fail-safe application. A cluster aware application such as Microsoft SQL Server runs on one node at a time. If that node goes down, another node takes over the service. Clustering helps ensure that your application is up all the time.

If the Microsoft cluster component is installed with CA Virtual Assurance, an administrator can register and manage clusters using the Administration tab.

### More Information

[Register a Cluster](#) (see page 91)

[Remove a Cluster](#) (see page 92)

[Modify Cluster Properties](#) (see page 92)

[Register a Cluster](#) (see page 91)

[Remove a Cluster](#) (see page 92)

[Modify Cluster Properties](#) (see page 92)

## Register a Cluster

You can register a Microsoft cluster using the Administration page of the user interface.

### To register a Microsoft cluster from the user interface

1. Click Administration.

The Administration page appears.

2. In the Provisioning section of the Configuration pane, click Microsoft Clusters.

The Microsoft Clusters section appears on the right.

3. Click + (Add) on the Registered Microsoft Clusters toolbar.

The Register New Cluster dialog appears.

4. Enter the required cluster name and access identification information, and click OK.

The Microsoft cluster is registered.

**Note:** Use the cluster hostname when you register a cluster.

## Remove a Cluster

You can remove a Microsoft cluster using the Administration page of the user interface.

### To remove a cluster

1. Click Administration.  
The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Microsoft Clusters.  
The Microsoft Clusters section appears on the right.
3. Select the cluster that you want to remove.
4. Click - (Delete) on the Registered Microsoft Clusters toolbar.  
A confirmation prompt appears.
5. Click OK.  
The cluster is removed.

## Modify Cluster Properties

You can modify Microsoft cluster properties using the Administration page of the user interface.

### To modify cluster properties

1. Click Administration.  
The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Microsoft Clusters.  
The Microsoft Clusters section appears on the right.
3. Select the cluster that you want to edit.
4. Click the Edit icon on the Registered Microsoft Clusters toolbar.  
The Modify Cluster Properties dialog appears.
5. Edit the required properties and click OK.  
The cluster properties are modified.

## Cisco UCS

The Cisco Unified Computing System (Cisco UCS) is the Cisco data center solution. The solution integrates a pair fabric interconnect switch with up to two switches, 40 chassis, and 320 blade servers (blades). A Cisco UCS Manager running on the switch provides management functionality for networking, storage, and blades, and also supports virtualization. CA Virtual Assurance interacts with Cisco UCS to query UCS device information including hardware resource, and health and device statistics. CA Virtual Assurance supports Cisco UCS using a UCS AIM and PMM. For information about the Cisco UCS interfaces and their operations, see the Cisco UCS documentation.

An administrator can register UCS Managers using either the Administration user interface or the dpmutil CLI command. If dpmutil is used, run nodecfgutil.exe to configure the UCS AIM.

**Note:** For CLI command information, see the *Reference Guide*.


**More information:**

[Interaction Between Cisco UCS Management Components](#) (see page 94)

## Interaction Between Cisco UCS Management Components

Cisco UCS integration requires the UCS AIM for SystemEDGE to provide SNMP get/set requests for retrieving UCS devices and statistic data and configuring devices. The UCS Platform Management Module (PMM) also queries UCS devices and statistic information and stores the data in the Management DB. Cisco provides an XML API for interaction with the Cisco UCS Manager. The API allows CA Virtual Assurance to gain access to the hardware, statistics, pools (UUID, MAC, WWPN, WWNN), and service profiles information provided by the UCS Manager.

### Interaction Between Cisco UCS Management Components



The XML API also provides the ability to configure certain device properties and perform pools and service profile management. Pools and Service Profile Management are one of the use cases that CA Virtual Assurance manages across multiple UCS Managers to detect pool range conflicts.

The diagram shows components used for the Cisco UCS integration. The communication protocol between the UCS Adapter and Cisco UCS manager is HTTP or HTTPS.

## Cisco UCS Overview

The Cisco Unified Computing System (Cisco UCS) is the Cisco data center solution. The solution integrates a pair fabric interconnect switch with up to two switches, 40 chassis, and 320 blade servers (blades). A Cisco UCS Manager running on the switch provides management functionality for networking, storage, and blades, and also supports virtualization. CA Virtual Assurance interacts with Cisco UCS to query UCS device information including hardware resources, and health and device statistics. CA Virtual Assurance supports Cisco UCS using a UCS AIM and UCS PMM. For information about the Cisco UCS interfaces and their operations, see the Cisco UCS documentation.

Administrators can add UCS Managers and register UCS AIMS using either the Administration user interface or the `dpmutil` CLI command.

**Note:** For CLI command information, see the *Reference Guide*.

### More information:

[Add a Cisco UCS Manager Server](#) (see page 96)

[Register a UCS AIM Server](#) (see page 96)

[Remove a UCS Server](#) (see page 97)

[Configure the SNMP Data Poller](#) (see page 97)

[Configure the Service Poller](#) (see page 98)

[UCS Trap Management](#) (see page 98)

[Service Profiles](#) (see page 99)

[Configure the SNMP Data Poller](#) (see page 97)

[Configure the Service Poller](#) (see page 98)

[UCS Trap Management](#) (see page 98)

[Service Profiles](#) (see page 99)

[How to Create or Update a Service Profile](#) (see page 100)

## Add a Cisco UCS Manager Server

If the Cisco UCS component is installed, you can add a Cisco UCS Manager server using the `dpmutil` CLI command or the Administration page of the user interface.

**Note:** For CLI command information, see the *Reference Guide*.

### To add a UCS Manager from the user interface

1. Click Administration.  
The Administration page appears.
2. Click Configuration.  
The Configuration page appears.
3. Click Cisco UCS Servers.  
The Cisco UCS Servers page appears.
4. Click + (Add) on the Cisco UCS Servers toolbar.  
The Add Cisco UCS Server dialog appears.
5. Enter the required server identification information, and click OK.  
The server is added.

## Register a UCS AIM Server

If the Cisco UCS component is installed, you can register a UCS AIM server using the `dpmutil` command or the Administration page of the user interface.

**Note:** For CLI command information, see the *Reference Guide*.

### To register a UCS AIM server from the user interface

1. Click Administration.  
The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Cisco UCS Servers.  
The UCS AIM Servers section appears on the right.
3. Click + (Add) on the UCS AIM Servers toolbar.  
The Add Cisco UCS AIM Server dialog appears.
4. Enter the required server and SNMP access identification information, and click OK.  
The UCS AIM server is registered.



## Remove a UCS Server

You can remove a Cisco UCS server using the Administration page of the user interface.

### To remove a UCS server

1. Click Administration.  
The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Cisco UCS Servers.  
The Cisco UCS Servers page appears.
3. Select the server that you want to remove.
4. Click - (Delete) on the server toolbar.  
A confirmation prompt appears.
5. Click OK.  
The server is removed.

## Configure the SNMP Data Poller

The SNMP data poller retrieves Cisco device information from the UCS AIM. Polled elements include:

- Switch
- Chassis (fan, PSU)
- Blade (main logic board, memory)
- Power usage
- Temperature

### To set the polling interval

1. Edit the `\conf\caucsconf.cfg` file as follows:

```
<property name="CONFIG_KEY_UCS_AIM_POLL_INTERVAL">  
  <!-- UCS AIM polling interval -->  
  <value>300</value>  
  <displayName>UCS AIM Polling Interval</displayName>  
</property>
```
2. Save the file.

## Configure the Service Poller

The *service poller* retrieves pool and service profile information directly from the UCS Manager. Polled elements include:

- UUID pools
- MAC pools
- World Wide Node Name (WWNN) pools
- World Wide Port Name (WWPN) pools
- Server pools
- Service profiles

The default service polling interval is 300 seconds.

### To reset the polling interval

1. Edit the `\conf\caucsconf.cfg` file as follows:

```
<property name="CONFIG_KEY_UCS_SERVICE_POLL_INTERVAL">
    <!-- UCS service interval in seconds -->
    <value>300</value>
    <displayName>UCS Manager Polling Interval</displayName>
</property>
```

2. Save the file.

## UCS Trap Management

The UCS PMM listens for UCS trap indications. All UCS traps are forwarded as events.

To forward UCS traps to your trapreceiver on default port 162, configure SystemEDGE.

## Service Profiles

A *service profile* contains configuration information about Cisco UCS hardware, including interfaces, fabric connectivity, and network and server identity. You can create a service profile for a specific UCS manager or centrally in the CA Virtual Assurance Management DB. Service profiles on a UCS manager can be *imported* into the Management DB, from which they can be *exported* to other UCS managers.

Service profiles allow supporting Cisco UCS hardware to be abstracted from the operating system. By addition or removal of a service profile, services can be brought online or taken offline. By reassignment of a service profile from one blade to another, a set of services (operating system and applications) can be moved to different servers.

Service profile information can include:

- Blade servers by device UUID or virtual UUID pool
- Storage in any configuration of local storage, RAID, or SAN (HBA, WWNN, and WWNN pools)
- Networking (MAC, vNIC 0, vNIC 1, and MAC pools)
- Server boot order or other policies
- Server assignment type (assign later, provision a slot in advance, provision an existing server, select server from pool)

## How to Create or Update a Service Profile

CA Virtual Assurance provides a wizard that administrators can use to create service profiles, and you can update service profiles with predefined policy options. System, network, and storage administrators can collaborate to create service profiles with unique identity characteristics, and required connectivity characteristics.

The service profile wizard provides a subset of the Cisco UCS Manager interface to take advantage of knowledge and experience in the Cisco environment.

### Example: Create a Service Profile

Using the CA Virtual Assurance user interface, select the create-service-profile option and specify the server profile name to create and choose whether the server profile is to be:

- Hardware-based
- Simple server profile with default networking and storage connectivity
- Based on an existing service profile template
- Custom service profile which must be created explicitly

Based on the option selected, the wizard leads you through an interview process to obtain the required identity and connectivity information. You can take the defaults or specify identity (UUID), network (MAC/VLAN), storage (WWN/vHBA), and boot policy information explicitly.

## Using Multiple Multi-instance AIMs

When you use a CA Virtual Assurance manager with multiple remote AIMs of the same type (vCenter Server, Solaris Zones, LPAR, MSCS, or Cisco UCS) to manage a particular virtual environment, consider the following:

Verify that each vCenter Server, Solaris Zones Server, HMC/IVM Server, or Cisco UCS Server is uniquely associated with one AIM. Use Administration on the CA Virtual Assurance manager user interface or NodeCfgUtil on the appropriate AIM server to remove any ambiguous associations.

If more than one AIM is managing such a server, this causes management, discovery, or monitoring issues.

### Example

Three vCenter Servers are managed through two AIMs:

- vCenter Server 1 and vCenter Server 2 are managed through vCenter Server AIM 1.
- vCenter Server 2 and vCenter Server 3 are managed through vCenter Server AIM 2.

In this case, vCenter Server 2 is associated with vCenter Server AIM 1 and vCenter Server AIM 2 which is not unique. You can remove one of these two associations either from AIM 1 or AIM 2 to get unique conditions.

## Systems Management


CA Virtual Assurance is designed to manage virtual environments, but it also discovers and manages systems (managed nodes). CA Virtual Assurance supports the following operating systems on managed nodes:

- AIX
- HP-UX
- Linux, zLinux
- Solaris (Intel, SPARC)
- Windows

Available management components for managed nodes are:

- SystemEDGE
- Advanced Encryption AIM
- Remote Monitoring AIM (Windows Servers only)
- Service Response Monitoring (SRM) AIM
- CA Performance Lite Agent

### Supported Virtual Environments and Operating Systems



SystemEDGE is the base for systems management in CA Virtual Assurance and provides the following benefits:

- Centralized remote agent deployment to all managed systems
- Centralized remote agent configuration

- Visualization of all monitored metrics, including status information from the object model of the agent
- Remote deployment and configuration of the Service Response Monitor AIM
- Enhanced agent security options

For details about agent functionality, see the *SystemEDGE User Guide*.

#### **More Information**

[SystemEDGE Features](#) (see page 103)

[Service Response Monitoring](#) (see page 113)

[Remote Monitoring](#) (see page 117)

[Agent Configuration](#) (see page 118)

[Agent Visualization](#) (see page 119)

[Security and Maintenance](#) (see page 122)

## **SystemEDGE Features**

SystemEDGE is a lightweight agent that provides SNMP-based monitoring of physical and virtual systems. Use the agent to access important system information such as system configuration, performance, users, file systems, and so on. Monitor this information based on specified thresholds or conditions; and create objects based on monitors to maintain aggregate object states.

SystemEDGE supports monitoring metrics from the following MIBs:

- MIB-II (RFC 1213)
- Host Resources MIB (RFC 1514)
- Systems Management MIB (CA proprietary)
- IF-MIB (partial) (RFC 2233)
- IP-MIB (partial) (RFC 4293)
- TCP-MIB (partial) (RFC 4022)
- UDP-MIB (partial) (RFC 4113)

You can use the monitoring tables in the Systems Management MIB to enable the following types of intelligent monitoring:

### **Self monitoring**

Provides monitoring of any integer-based MIB object that the agent supports. Create entries in the Self Monitor table to specify objects to monitor, comparison operators, threshold values, and severities. The agent automatically monitors the objects according to your entries. The agent monitors the objects, maintains a current state according to specified threshold and severity values. The agent sends a state change trap when thresholds are breached.

### **Process and service monitoring**

Provides monitoring of any process, Windows service, or application. Create entries in the Process Monitor table to monitor whether a process or service is running or to monitor process table objects against specified thresholds. The agent monitors the processes, maintains a current state according to specified threshold and severity values. The agent sends a state change trap when thresholds are breached or the state of a process (running or stopped) changes.

### **Process group monitoring**

Provides the ability to define a group of processes and monitor that group for changes. Create entries in the Process Group Monitor table defining process groups, and the agent monitors the groups. If a process group changes, the agent sends a trap.

### **Log file and directory monitoring**

Provides monitoring of any UTF-8 encoded system or application log file by searching for strings specified as regular expressions. Create entries in the Log Monitor table, and the agent monitors the specified log file for lines matching user-defined regular expressions. The agent sends a trap when a match occurs. You can associate a severity with the monitor, which is included with the sent trap.

### **Windows event monitoring**

Provides monitoring of Windows event log entries using different filters, such as event source. Create entries in the NT Event Monitor table, and the agent monitors the event log for events matching user-defined regular expressions. The agent sends a trap when a match occurs.

### **History collection**

Provides historical data collection for manager-side baselining and trend analysis. Create entries in the History Control table, and the agent collects metrics over time. Use the metrics to provide a picture of average system performance during a specific time interval.

For more information about monitoring functionality and SystemEDGE architecture, see the *SystemEDGE User Guide*.



**More Information**

[Systems Management MIB](#) (see page 105)  
[State Management Model](#) (see page 107)  
[Configure Object Aggregation](#) (see page 108)  
[Stateless Monitoring](#) (see page 109)  
[Managed Mode and Legacy Mode](#) (see page 110)  
[SystemEDGE AIMs](#) (see page 111)

## Systems Management MIB

The Systems Management MIB is a private-enterprise MIB that includes objects for monitoring the health and performance of the underlying system and its applications.

The groups and tables with objects that you can monitor in the Systems Management MIB are as follows:

**System Group (sysedgeSystem)**

Contains basic system information such as host name, CPU type, and operating system version.

**Mounted Devices Table (devTable)**

Contains information about devices and file systems mounted on the host. You can create monitors for values such as file system space or unmount a mounted device by setting a column value in this table.

**Kernel Configuration Group (kernelConfig)**

Contains kernel information such as number of CPUs, amount of virtual memory, and clock rate. You can monitor how the kernel is configured and the kernel version using this group.

**Boot Configuration Group (bootconf)**

Contains information about the root file system, dump file system, and swap space. Monitor this table to track values such as root file system name, file system blocks, and file system type.

**Streams Group (streams)**

Contains information about the streams I/O subsystem. You can monitor the health of the subsystem by monitoring objects in this group such as number of streams in use, number of stream allocation failures, and number of streams in queue.

**User Table (userTable)**

Contains information about the user accounts on the system.

**Group Table (groupTable)**

Contains information about the user groups on the system.

**Process Table (processTable)**

Contains information about running processes. You can monitor this table to track the processes that are currently running, and you can also control processes by setting certain attributes. For example, you can kill a process by setting the value of the processkill column to 9.

**Who Table (whoTable)**

Contains information about the users currently logged on to the system. You can monitor attributes in this table to track who is using a system at any particular time.

**Remote Shell Group (remoteshell)**

Contains attributes for running shell scripts and programs on the remote system. Set the attributes in this table to specify a command, its arguments, and the name of an output file.

**Kernel Performance Group (kernelperf)**

Contains information about the health and performance of the host operating system. You can monitor attributes such as the number of current processes and open files, the number of active jobs, and the number of jobs in the scheduler queue.

**Interprocess Communication Tables (msgqueueTable, shmемTable, semTable)**

Contains information about message queues, shared memory, and semaphores in separate tables. Monitor these tables to coordinate communication between processes.

**Message Buffers Allocation Table (mbufAllocTable)**

Contains information about how your system is using message buffers. Monitor attributes in this table to track information such as the number of times buffer requests were denied or delayed.

**Streams Buffers Allocation Table (strbufAllocTable)**

Contains information about buffer allocation and usage statistics for buffers used by the Streams subsystem.

**I/O Buffer Cache Group (ioBufferCache)**

Contains information about I/O buffer allocation and usage for basic disk I/O. Monitor this table to track information such as peak periods of I/O buffer activity.

**Directory Name Lookup Cache Group (dnlc)**

Contains information about directory and file name cache performance.

**AIX Logical Partition Group (logicalPartition)**

Contains information about IBM AIX logical partitions (LPARs). You can monitor attributes such as physical or logical CPU for each partition and the number of CPUs for each partition.

**Trap Community Table (trapCommunityTable)**

Contains SNMP information such as configured communities, users, and trap destinations.

**NT System Group (ntSystem)**

Contains information specific to Windows systems. This group contains System, Thread, Registry, Service, System Performance, Cache Performance, Memory Performance, Page File Performance, and Event Monitor groups for monitoring attributes for these areas on Windows systems.

**RPC Statistics Group (rpc)**

Contains information about kernel remote procedure calls. Monitor this table to track attributes such as counters and statistics for detecting peak periods of RPC activity.

**NFS Statistics Group (nfs)**

Contains information about the kernel's NFS facility. Monitor this table to track attributes such as statistics and counters for detecting peak periods of NFS activity.

**Disk Statistics Table (diskStatsTable)**

Contains information about disk I/O.

**CPU Statistics Table (cpuStatsTable)**

Contains performance statistics for each CPU. You can monitor attributes such as time spent in Idle mode and time spent in Wait mode.

The Systems Management MIB also contains the monitoring tables and tables to support object aggregation.

## State Management Model

The SystemEDGE agent supports a state management model for self monitors and process monitors fully integrated with the overall CA Virtual Assurance Management Model. The agent aggregates multiple monitors of different severities into a single Managed Object. This object has a state corresponding to the breached monitor with the worst severity.

The agent calculates individual monitor states according to an assigned severity value. The resultant states can be one of the following:

- unknown (1)
- ok (2)
- warning (3)
- minor (4)

- major (5)
- critical (6)
- fatal (7)
- up (11)
- down (12)

**Note:** If a monitor has a severity of none, the state toggles between up and down.

The Aggregate table of the Systems Management MIB uses the object class, instance, and attribute values to aggregate monitors with the same values into one entry. This entry represents a monitored object, for which it maintains an aggregate state.

**Note:** If you do not enter values for the object class, instance, and attribute in a monitor, the agent populates them with meaningful default information. Default self monitor values are based on the monitored OID using a sysedge.oid file that maps a monitored OID to instance, class, and attribute values. Default process monitor values are based on the process regular expression and monitored attribute.

The Aggregate table updates the current state in the table and sends a state change trap only when a threshold breach creates the worst state of all monitors for an object. For example, assume that you are monitoring CPU usage with three monitors; one for 60 percent (assigned a warning severity), one for 80 percent (critical severity), and one for 100 percent (fatal severity); and the agent returns 82 percent CPU usage. This value causes a threshold breach for the 60 percent and 80 percent monitors. However, the agent only sends one state change trap for the 80 percent monitor and changes the aggregate state to critical.

## Configure Object Aggregation

By default, SystemEDGE aggregates monitors into a managed object that contain the same values for the object class, instance, and attribute properties. For example, all monitors with a class of SysHealth, an instance of CPU, and an attribute of SysTime are combined into an aggregate managed object.

You can configure the agent to aggregate objects on higher levels when defining SystemEDGE policy. You can also configure other aspects of agent behavior related to object aggregation and the state management model.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Control Settings.

The Controls page appears.

4. Click Aggregate Monitors.

The Aggregate Monitors page appears.

5. Select one or more of the check boxes to specify aggregation levels.

These represent higher aggregation levels than the default, up to aggregating all monitors into one top-level agent object. Specifying aggregation levels lets you create a tiered object architecture that propagates status up to the level you specify.

6. Configure the following additional settings, and click Save Policy:

**Send legacy traps for all aggregated monitors**

Specifies whether to send legacy traps for all monitors that make up a managed object. By default, the agent only sends a state change trap for the monitor with the highest severity, even if other monitors in the object experience threshold breaches.

**Execute commands of all aggregated monitors**

Specifies whether to execute action commands for all monitors that make up a managed object. By default, the agent only runs an action command for the monitor with the highest severity, even if other monitors in the object experience threshold breaches.

Aggregation settings are configured. Apply or reapply the policy for the changes to take effect.

**More Information:**

[Define SystemEDGE Policy Control Settings](#) (see page 158)

## Stateless Monitoring

Stateless monitors do not derive object status information or use the object model to maintain an overall object state. These monitors do maintain a severity value, but this severity is for tracking the importance of the individual monitor and is not used to calculate object state. The following tables support stateless monitoring:

- Process Group Monitor
- Log File Monitor
- NT Event Monitor

You can configure these monitors from the CA Virtual Assurance user interface, but you cannot visualize the resultant data. You must rely on traps that the agent sends when one of the following is detected based on defined monitors:

- Process group change
- A log file message matching a specified regular expression
- A directory threshold breach
- A Windows event log event that matches specified criteria

For more information about creating process group, log file, and Windows event monitors, see the *SystemEDGE User Guide*.

### Managed Mode and Legacy Mode

When you deploy SystemEDGE (or install it on a standalone basis), you can specify to run the agent in managed mode. In managed mode, the agent is managed by the CA Virtual Assurance Manager node from which you deployed the agent (or a Manager node that you specify in a standalone agent installation). Operating the agent in managed mode enables all CA Virtual Assurance agent management functionality, such as remote configuration and advanced visualization from the CA Virtual Assurance user interface. Managed mode also establishes CA Virtual Assurance as the primary source of agent configuration. If an agent in managed mode is modified outside of CA Virtual Assurance, the CA Virtual Assurance administrator can block or overwrite the change.

You can also operate SystemEDGE in legacy mode, or without a CA Virtual Assurance Manager controlling its configuration. An agent running in legacy mode is not restricted to legacy monitors, or monitors that do not maintain and calculate state.

When deploying an agent from CA Virtual Assurance, you specify whether to run it in managed mode in the package wrapper settings using the 'Run in Managed Mode' check box. When installing an agent separately from CA Virtual Assurance, provide a CA Virtual Assurance Manager node for the agent to run in managed mode.

## SystemEDGE AIMs

Application Insight Modules (AIMs) extend the SystemEDGE agent into application-specific monitoring or provide specialized monitoring capabilities that do not exist in the base agent. The following AIMs are provided with CA Virtual Assurance:

- Remote Monitoring AIM
- Service Response Monitor (SRM) AIM
- VMware vCenter Server AIM
- Solaris Zones AIM
- Hyper-V AIM
- MSCS AIM
- Cisco UCS AIM
- LPAR AIM

You can configure and deploy these AIMs and SystemEDGE from the CA Virtual Assurance user interface.

For more information about SystemEDGE, see the *SystemEDGE User Guide*.

## Remote Configuration and Deployment

CA Virtual Assurance provides end-to-end management of SystemEDGE from a centralized location through file-based configuration and deployment. File-based configuration does the following:

- Allows for changes that SNMP managers are unable to perform through SNMP Sets, like configuring read-only areas of the Systems Management Empire MIB such as community strings, operational parameters, trap communities, AIM loading, and so on
- Enables verifiable delivery and creates fewer security vulnerabilities than SNMP
- Enables complete remote agent configuration and delivery and minimizes or even eliminates the requirement for local manual file manipulation

The CA Virtual Assurance manager provides installation packages for SystemEDGE for each supported platform. The deployment solution lets you specify credentials and installation parameters, and when you have specified all required information, you can save the deployment package and deploy it to managed systems in your enterprise.

From the CA Virtual Assurance user interface, you can set all agent configuration parameters in the sysedge.cf file and create self-monitoring entries using a graphical representation of the monitor tables. You can save an agent configuration as a profile and apply that profile to a deployment package for deployment to managed systems. You can also make point configuration changes to installed agents and apply those changes instantly.

When you deploy a configuration change, CA Virtual Assurance compiles the `sysedge.cf` file using the template directives in the file to apply the information you entered in the CA Virtual Assurance user interface for configuration and self monitoring. It sends the compiled `sysedge.cf` file to the specified systems using CAM (CA Messaging) as the transport mechanism. CAM sends the file to the data directory specified during installation. The SystemEDGE agent listens for CAM messages containing configuration files, and when a file is detected, the agent checks the file for correct syntax and then automatically loads and uses the file. Configuring the agent through CA Virtual Assurance triggers an agent warm start, so that you can make changes to the agent that it can apply automatically without requiring a restart.




## Service Response Monitoring

The Service Response Monitoring Application Insight Module (SRM AIM) is a functional extension (plug-in) for SystemEDGE. SRM retrieves the responsiveness of a logical or physical service that runs on the local or on a remote system. SRM is Java-based and multi-threaded and handles multiple test configurations across multiple servers. SRM executes preconfigured or custom tests to measure the elapsed time and throughput of execution.

The following diagram illustrates these relationships.

### Interaction Between Service Response Monitoring Components



The `svcrsp.cf` configuration file contains the test specifications. The SRM AIM reads this configuration file and makes the test specifications available in the shared memory segment. The SRM Jcollector component reads each test configuration from the shared memory. Jcollector executes the tests, collects the results of this timing process, and propagates the results to the SRM AIM. SystemEDGE sends these results and associated status information to CA Virtual Assurance.

The Service Response Monitor (SRM) AIM monitors the availability and response time of critical system services, such as DNS, DHCP, or SQL-based on defined thresholds. Enable this functionality by creating SRM tests. SRM tests let you do the following:

- Test system service availability and response time
- Gain visibility into the complex, multi-tier infrastructure to pinpoint problems before users are affected
- Obtain real-time notification of delays, outages, and performance problems
- Confirm that services such as DNS and DHCP are performing well against service level agreements
- Maintain historical data for capacity planning, troubleshooting, or analyzing trends in long-term behavior

CA Virtual Assurance provides the following functionality for the SRM AIM:

- Remote deployment with the SystemEDGE agent
- Remote test configuration
- Test visualization

For more information about the SRM AIM architecture, see the *SRM User Guide*.

### More Information

[SRM Tests](#) (see page 114)

## SRM Tests

The SRM AIM provides the following response time tests:

### Active Directory

Verifies that Windows Active Directory Services are working properly to manage shared files and resources.

### Custom

Verifies that important custom services or other tasks are working efficiently.

**DHCP**

Verifies that Dynamic Host Configuration Protocol servers are responding to address requests.

**DNS**

Verifies the Domain Name System servers are processing hostname to address resolution requests.

**File I/O**

Verifies that operations such as read, write, and compare work across file systems.

**FTP and TFTP**

Verifies that users can log in to specified servers to upload and download files.

**HTTP and HTTPS**

Verifies that users can connect to your business web servers and determines whether specific text displays on a web page.

**LDAP**

Verifies the connection to LDAP servers to verify access for user requests and LDAP queries.

**NIS**

Verifies that NIS map requests are being processed.

**NNTP**

Verifies that users can connect to their Usenet newsgroup servers and company bulletin boards.

**Ping**

Verifies that network devices exist and are reachable across the network.

**Email**

Verifies that email servers are available and processing email effectively. SRM supports tests for IMAP, MAPI, POP3, SMTP, and round-trip email that originates from an SMTP server.

**SNMP**

Verifies that SNMP agents are responding to SNMPv1 GET requests.

**SQL Query**

Verifies that SQL database servers are available and processing short queries.

**TCP**

Verifies that systems are listening for and processing connection requests.

**Virtual User**

Obtains continuous response time and availability data for actual user transactions (keyboard entry and mouse clicks) that can be recorded (typically with WinTask) to confirm that business tasks run successfully.


## Remote Monitoring

The Remote Monitoring AIM accesses an RM system through a WMI connection to the root\CIMV2 namespace utilizing DCOM. DCOM requires local system administrator user and password credentials. If you want to monitor a Windows computer, you must provide these credentials which the RM AIM stores in a file. The password is encrypted.

Remote Monitoring collects and provides Windows system information (Win32\_ComputerSystem) performing WMI queries (port 135) on the monitored RM systems. WMI uses port 135 (default).

The following diagram illustrates these relationships.

### Interaction Between Remote Monitoring Components



## Agent Configuration

The following two types of SystemEDGE configuration are available from the CA Virtual Assurance user interface:

### Point Configuration

Lets you make singular, temporary changes to an agent without having to deploy policy. For example, you can change a self monitor threshold, add a temporary process monitor, or create a self monitor for an SRM test. Policy deployments override point configuration changes.

### Policy Configuration

Lets you create agent configuration policy that you can deploy to sets of managed machines in one operation. For example, you can define a policy containing a set of common monitors and SRM tests and deploy that policy to all systems in your enterprise to ensure that the same important system metrics are being monitored.

Configuring agents in managed mode from the CA Virtual Assurance user interface takes precedence over all other forms of configuration. If a user makes manual changes to a local agent through the sysedge.cf configuration file or through SNMP Sets, CA Virtual Assurance policy configuration overrides these changes after the policy is applied.

### More Information

[Perform a Point Agent Configuration](#) (see page 118)

[Policy Configuration](#) (see page 155)

## Perform a Point Agent Configuration

CA Virtual Assurance provides the ability to make single or point configuration changes to a single agent without creating and applying policy. This functionality is meant for temporary changes to the monitoring configuration of a single system. The following scenarios provide examples of when a point configuration change may be useful or necessary:

- Any change considered temporary that is specific to an individual system
- A change to address a temporary aberration
- Changes to experiment with different monitoring severities and thresholds before committing these to a general monitoring policy

When you make a point configuration change, CA Virtual Assurance applies the change to the system on top of any existing policy or local configuration. However, the next time you apply policy to the system, the policy overwrites the point configuration change. Point configuration changes are reported as policy exceptions until they are merged into the base policy or overwritten by a policy application.

Point configuration is available for self and process monitors.

**To perform a point agent configuration**

1. Click Resources, and select the system to configure in the Explore pane.  
System information appears in the right pane.
2. Click Configuration in the right pane, and select Self Monitors or Process Monitors.  
The existing self or process monitors appear.
3. Click + (New) on the toolbar.  
Fields appear for creating a new self or process monitor.
4. Complete the necessary fields, and click Save.

**Note:** For more information, see the *SystemEDGE User Guide*.

The monitor is saved and appears in the updated list of self or process monitors.

You can also modify, delete, or copy an existing self or process monitor.

## Agent Visualization

The CA Virtual Assurance user interface displays monitoring information for systems with agents in managed mode. Platform management models (PMMs) interpret and transform agent information so that it fits in the underlying CA Virtual Assurance AIP architecture and can be represented in the AOM database. PMMs are available for the base SystemEDGE agent and the SRM AIM.

Agent data that you can visualize in the CA Virtual Assurance user interface includes the following:

- Managed objects created using the state management model
- The state of all managed objects
- Individual monitors
- SRM tests

### More Information

[View Managed Object States](#) (see page 120)

[View SystemEDGE Monitors](#) (see page 120)

[View SRM Tests](#) (see page 121)

[Dashboard Status Views](#) (see page 122)

## View Managed Object States

The CA Virtual Assurance user interface displays all SystemEDGE managed objects for systems with agents running in managed mode.

To view managed object states, click Resources, expand Managed, and select a system.

The system Summary page appears in the right pane.

The Machine Status Information pane contains the total number of managed objects and the maximum object severity.

The Managed Objects table contains the following information about each managed object:

- Health state
- Operating status (active or inactive)
- Object class, instance, and attribute
- Current monitored value and threshold value

From this table, you can select a managed object and click Actions, Go to Definition to view the monitors that make up the managed object.

## View SystemEDGE Monitors

The CA Virtual Assurance user interface displays all defined self and process monitors for systems running SystemEDGE in managed mode. You can view details about each monitor and [perform point configuration](#) (see page 118) such as adding, deleting, modifying, or copying a monitor.

### To view SystemEDGE monitors

1. Click Resources, expand Managed, and select a system.  
The system Summary page appears in the right pane.
2. Click Configuration, and click Self Monitors or Process Monitors.  
The Self Monitors or Process Monitors pane appears.



The Self Monitors and Process Monitors panes contain a table listing the following monitor properties:

- Index
- State
- Status

**Note:** This state may not be the same as the state of any managed object with which the monitor is associated. The managed object state is the worst current state of all monitors that make up the object.

- Class, Instance, and Attribute of the object

**Note:** Monitors with the same values in these columns are part of the same managed object.

- Value, Operator, and Threshold of the object currently monitored
- Severity
- Trap #
- Last Trap
- Flags

## View SRM Tests

The CA Virtual Assurance user interface displays SRM tests for systems running SystemEDGE in managed mode with the SRM AIM.

### To view SRM tests

1. Click Resources, expand Managed, and select a system.  
The system Summary page appears in the right pane.
2. Click Details, and click SRM Tests.  
The SRM Tests pane appears.

The SRM Tests pane contains a table listing the following test properties:

- Index number
- Object class name
- Test name and type
- Test destination
- Interval
- Status
- Last Results
- Total errors

## Dashboard Status Views

The following views are available on the Dashboard for tracking agent status:

### SystemEDGE Machines Status

Displays a table that lists all systems with SystemEDGE running in managed mode and the worst current state of all managed objects on each system. By default, only systems with at least one managed object with a current state of critical or higher are displayed. Click Show Filters to filter the view contents based on state and machine name. Click a system name in the table to view details about all managed objects on the system.

### SystemEDGE Objects Status

Displays a table that lists all SystemEDGE managed objects with a current state of critical or higher. Click Show Filters to filter the view contents based on state, machine name, class, instance, and attribute. Click a system name to view details about all managed objects on the system.

## Security and Maintenance

CA Virtual Assurance offers the following enhanced security and maintenance options for the SystemEDGE agent:

- Maintenance mode configurable from the user interface
- A single point of configuration for the SystemEDGE agent
- The ability to block changes performed outside of CA Virtual Assurance
- Notification of changes performed outside of CA Virtual Assurance, and the opportunity to override or reject unwanted changes

### More Information

[Enable Maintenance Mode](#) (see page 122)

## Enable Maintenance Mode

You can enable SystemEDGE maintenance mode in CA Virtual Assurance, in which the agent stops processing all monitor entries and sending traps. Maintenance mode is useful if the agent's system is undergoing a planned outage and you want to avoid receiving false alarm traps.

While in maintenance mode, the agent continues to collect metrics and respond to SNMP requests, but it suspends processing all monitors and history collections. The agent saves the current value of all monitors at the beginning of the maintenance window, compares it to the current value at the end of the maintenance window, and sends traps in response to the current value as necessary.

**Follow these steps:**

1. Open the Explore pane.  
Available groups, services, and systems appear.
2. Expand Managed, and select a system.
3. Click the Resources tab, and then Monitoring Software.  
The Machine Details pane appears.
4. Set the Maintenance Mode option to Enabled, and click Apply.  
The agent performs a warm start and enables maintenance mode.

To take the agent out of maintenance mode, simply clear the Maintenance Mode check box, and click Apply.



# Chapter 6: Managing Policies

---

This section contains the following topics:

[Rules and Actions](#) (see page 125)

[Use Cases for Policies](#) (see page 139)

[Configuring Data Collection](#) (see page 141)

[UCS Action Types for Policies](#) (see page 152)

[Policy Configuration](#) (see page 155)

## Rules and Actions

To configure rules and actions, you must first understand what they are and how they interact with each other and other components. By understanding these interactions, you can best decide how to set up your rules and actions to manage your data center efficiently.

CA Virtual Assurance collects and analyzes metrics and then makes intelligent decisions based on the analysis about how to distribute resources. For example, if CA Virtual Assurance determines that a server or a service is overutilized or underutilized, it can provision a new computer.

Usage is monitored at the server level and the service level. Server level monitoring involves diagnosing problems with a specific server and only key performance indicators are used. Service level monitoring diagnoses problems with the service as a whole and overall usage is used as the performance indicator.

Rules can be created at the server level or the service level. You create rules to evaluate performance metrics and generated events. Rules are composed of individual or combinations of conditions which must evaluate overall to a true state for an action to be taken. You can create your own rules or you can select a set of rule templates to generate rules using automation policy.

**Note:** For a list of performance metrics and descriptions, see the *Reference Guide*.

By default, the rules are evaluated at the recording interval defined in the collection settings at the data center level (default = 300 seconds) or when events occur because of monitored metric values. You can configure specific servers to override the default data center recording interval when you want to set an interval that differs from the data center. Server level rules are evaluated at the configured server level recording interval. Service level rules are evaluated at the shortest recording interval among all the servers within that service. When you change the recording interval, stop and restart the Policy Manager service to retrieve and use the updated interval for rule evaluation.

Metrics are the source of the evaluation data. When a metric rule evaluates to true, the action is triggered. The lag must be exceeded for a rule to evaluate to true. In some scenarios, you would want a one-time breach of a rule to trigger an action so you would set your lag to one, but in other instances you would not want a one-time event to trigger a rule.

For example, CA Virtual Assurance is integrated with CA SDM, which is a customer support application that manages calls, tracks problem resolution, shares corporate knowledge, and manages IT assets. If you want to open tickets automatically when your action is triggered, you can set your actions to interact with CA SDM. This arrangement is useful for actions requiring third-party approval. After the third party approves your ticket in CA SDM, the action will automatically run.

You can also schedule your actions to run at specified times using the initiation component. The current parameters for the action are saved when you create the job. If you change the action details after the job has been submitted, it will not have an impact on jobs that you have already scheduled to run. If you must change the action details of a job that has already been scheduled, open the job that uses the action and save it again to update it with the new action details.

## Configure CA SDM

For CA SDM releases before Version 12.5, configure CA SDM properly with the appropriate ticket status codes, so that you can set your actions to open issues automatically when necessary.

**Note:** The release number of CA Virtual Assurance and CA SDM need not be the same, as long as the two products do not share a database.

### To configure CA SDM

1. Log in to your CA SDM server by typing the following information in your web browser:

`http://servicedesk_servername:8080`

The CA SDM splash screen appears.

2. Enter your user name and password, and click Log In.

The CA SDM main page appears.

3. Click Administration and expand the Service Desk tree node in the left pane.

4. Select Requests\Incidents\Problems and then Status.

The Request\Incident>Status List appears.

5. Click Create New.

A Create New Request Status window opens.

6. Type **Approved** in the Symbol text box, select Active from the Record Status drop-down list, type **APP** in the Code text box, and click Save.

The new request status appears in the list.

7. Type **Rejected** in the Symbol text box, select Active from the Record Status drop-down list, type **REJ** in the Code text box, and click Save.

The new request status appears in the list.

CA SDM setup is complete and you can now automatically open requests when an action is triggered.

## Configure the CA SDM Ticket Status Setting

CA SDM versions before 12.5 used default status code settings of APP (Approved) and REJ (Rejected) for help desk tickets. CA Virtual Assurance uses and searches for these approval codes to run operations that are started upon approval of help desk tickets. These operations include but are not limited to running actions, reserving systems, and so on. If you are using CA SDM Version 12.5, new ticket status codes are supported. PRBAPP (Approved) and PRBREJ (Rejected) must be associated to the existing approval codes in CA Virtual Assurance. To support the new codes and for the product to work properly, update the configuration file as shown in the following steps.

### To change the ticket status setting

1. Open the caaipconf.cfg file located in the CA Virtual Assurance *Install\_Path*\conf directory with a text editor, and scroll to the Help Desk section.
2. Locate the special status code property as shown:

```
<property name="SPECIAL_STATUS_CODE">
  <!-- APP_CODE=PRBAPP;REJ_CODE=PRBREJ;(each code must be terminated by a semicolon) -->
  <value/>
  <displayName>type of code that added in SD R12.5 and later</displayName>
</property>
```

3. Uncomment and change the code as shown:

```
<property name="SPECIAL_STATUS_CODE">
  <value>APP_CODE=PRBAPP;REJ_CODE=PRBREJ;</value>
  <displayName>type of code that added in SD R12.5 and later</displayName>
</property>
```

CA Virtual Assurance is configured to use the CA SDM 12.5 status codes.

4. Save and close the file to enable the configuration change.

## Rule Planning

Consider the following points when setting up rules and actions:

- Which VMs, servers, and services do you want to analyze?
- What actions do you want to take when CA Virtual Assurance discovers violations?
- Which rules can be generic and which ones should be specific? Carefully consider the impact on your environment when planning generic rules that include scripts or batch files.
- Which metrics are you interested in evaluating?
- How many times should a rule be breached before an action is triggered? Consider that excessive executions of actions have a negative impact on performance in your environment.

**Note:** Actions that specify a help desk approval requirement cannot be used for action scheduling. If you need the same action for a scheduled action, create a second action that does not include the help desk approval requirement.

## Create Automation Policy

You can use the Create Automation Policy wizard to create automation rules based on two predefined policy types:

- Virtual Machine Dynamic Resource Brokering – CPU and memory allocation is dynamically changed based on defined utilization thresholds.
- Overall Utilization Metric Threshold Monitoring – Health state is set according to overall utilization.

**Follow these steps:**

1. Open the Manage pane, and click Create Automation Policy.

The Create Automation Policy wizard appears.

2. Select a Policy Type, and click Next to select target resources and set conditions for rules.

The Policy Summary displays the result.

3. Click Finish.

The policy is confirmed and the corresponding rules are created.



## Create a Rule

A rule functions as a trigger that runs your action when the rule condition is evaluated as true.

**Note:** Only the original creator or an administrator can edit or delete a rule.

**Follow these steps:**

1. Click Resources and select a server or service in the Explore tree.
2. Click the Policy tab, and then the Rules tab.

The Rules page appears.

3. Click + (Add new rule).

The Rule/Template wizard appears.

4. Type a meaningful name for the rule in the Identification section, and then select Rule to create a rule.

**Note:** Select Template to create a rule template that can be used with multiple rule definitions.

5. Select Enable to make the rule active.
6. Select Unlimited or Maximum (with number of retries) as the Number of Executions Allowed.

**Note:** Setting a limit on the number of times the rule can run prevents excessive retries that slow down system response time.

7. Click Next.

The Template Modeling and Action Selection section appears.

8. Define whether to model the rule on a template. Select an existing template or enter a name for a new template and select Enable to inherit any changes to the template.

9. Select the action for your rule from the list. Click Next.

The Define Rule Formula section appears.

10. Create the condition formula for your rule by completing the following fields in the Rule Evaluation Formula section:

**Source**

Specifies the source for the data that the rule evaluates, which can be Overall Utilization, Event, or specific server metrics.

### Operator

Specifies how to evaluate the source data against the value you enter in the Value field. The valid operators depend on the source. For example, if you select Overall Utilization, the following operators are valid:

"=" "!=" "<" "<=" ">" ">="

If you choose Event, the values are as follows:

#### contains

Matches an exact string or substring. Wildcards are not permitted in the Value field.

#### RegEx (Regular Expression)

Returns a value of true when strings matching the specified regular expression are found. Returns a value of false when no strings matching the specified regular expression are found.

#### NotRegEx

Returns a value of true when no strings matching the specified regular expression are found. Returns a value of false when strings matching the specified regular expression are found.

**Important!** Verify that the rule and action name does not contain the string that you want to match. This best practice helps to avoid incremental firing of actions when events are matched in the next rule evaluation cycle.

**Example:** If the Value field contains *threshold* as the matching string, the following events are matched:

Event A: The memory *threshold* has been breached!

Event B: threshold

### Value

Specifies the numeric value or alphanumeric string against which the selected operator evaluates the source data.

### Lag

Defines how often the rule must evaluate as true before the action triggers. Some actions that you define should trigger after a single occurrence. Other actions should trigger only after a number of occurrences signal a persistent problem. **Note:** When Source is set to Event, Lag is disabled by default.

**Logic Op**

Defines multiple formulas by using the logical operators AND or OR. Click New to complete each definition and add the formula to the list of defined formulas. The last formula that you define is set to NOOP by default.

Your condition formula will be used to trigger the action when the rule evaluates to true. The Confirm Configuration section appears.

11. Review the details of your rule, and then click Next at the top of the page.
12. Click Finish to commit the update.

Your rule or template is added to the Rules list.

13. Click the Return to Rules List link to verify that the rule has been added.

**Example: Set a Server Level Rule**

This example sets a rule for a server that exceeds CPU and memory thresholds more than three times, or when an event occurs indicating that a server is discovered.

**Rule formulas:**

1. CPU Utilization % > 80 (Lag 3) AND
2. Memory Utilization % > 50 (Lag 3) OR
3. Event RegEx .\*discovered
4. Event NotRegEx .\*discovered NOOP

**Action:** Add 200 CPU Shares, Max 8000

## Use a Predefined Action Type

You can select a predefined action type for your rule. If the conditions for a rule evaluate to true, the action that you defined runs.

**Follow these steps:**

1. Click the Policy tab, and then click the Actions & Rules tab.  
The Actions & Rules page appears.
2. Click the Actions tab.  
The Actions page appears.
3. Click + (Add new action).  
The Action Definition: New page appears.

4. Enter a meaningful name for the action in the Name text box, and select a predefined action type using the following menus:
  - Category - Product functional area filter. To list all action types, select All Categories.
  - Type - Available action types
  - Environment - Applicable platforms (for example, VMware vCenter or Microsoft Hyper-V)

The Details section appears. The options that appear in the section depend on the action type that you selected.

5. Select one of the following settings in the Action Start drop-down menu:

**No Delay**

Specifies that the same action can be rerun immediately when a rule using that action is triggered again.

**Delay For**

Specifies the time in seconds that must elapse before the same action can be rerun when a rule using that action is triggered again.

**Note:** The Action Start setting has no effect when the action is run by a scheduled job.

6. Select one of the following settings in the Action Completion drop-down list:

**No Wait**

Specifies not to wait for the action to complete before running succeeding actions in an action sequence.

**Wait No Longer Than**

Specifies to wait no longer than a specified value in minutes for the action to complete before running succeeding actions in an action sequence.

**Wait Indefinitely**

Specifies to wait for the action to complete. The succeeding actions in an action sequence run only after this action has been completed.

**Note:** The Action Completion drop-down list appears only for long-running actions.

7. Complete the fields for the requested information.
8. Select the Help Desk Approval check box if the ticket requires approval by a third party.

**Note:** CA SDM must be configured to use this option.

The Ticket Types and Templates fields become enabled.

**Note:** Actions that specify a help desk approval requirement cannot be used for action scheduling. If you need the same action for a scheduled action, create a second action that does not include the help desk approval requirement.

9. Select Auto close ticket on approval if you want to close the ticket automatically after it is approved.

10. Select a ticket type from the Ticket Types drop-down list. The following types are valid options, but depend on your configuration:

- Default
- Incident
- Problem
- Request

The Templates drop-down list is updated with the templates associated with the ticket type you selected.

11. Select a template from the Templates drop-down list.

The fields are populated with predetermined values depending on the ticket model you are using.

12. Click Save.

A confirmation message notifies you that your save was successful.

For testing purposes, you can run the action from the Actions page by selecting the action and clicking the Run action icon.

## Action Types

Several categories of action types are available.

**Note:** When using special or reserved characters in any operation, consider operating system and shell behavior. Behaviors include, but are not limited to, the invocation of custom scripts run by the operating system shell. For more information about shell behavior and how to escape special characters, see the Microsoft TechNet website at <http://technet.microsoft.com/en-us/library/cc723564.aspx>.

### Predefined Action Types

Predefined action types are commonly used actions that are available for you to use when creating actions for your rules. Action types are calling command-line utilities. All action types are listed in a drop-down list in the Policy, Actions & Rules pages of the user interface.

**Note:** For detailed descriptions of action types, see the *Online Help*.

### Custom Action Types

You can create custom action types using substitution strings rather than typing the full command line. The custom action types are added to the drop-down list of predefined action types. You can control user access to custom actions, in general, or you can control access to individual custom actions through the Administration page in the user interface.

The Run Command Script action type provides string substitutions that let you perform multiple actions on servers. String substitutions provide more flexible rules and reduce the need for custom scripts. The following string substitutions are available:

- %ACTIONNAME%
- %EVENTMESSAGE%
- %EVENTSOURCE%
- %RULENAME%
- %SERVER%
- %SERVICE%

The following string substitutions are only valid for actions running in an action sequence:

- %STDOUT% - standard output
- %STDERR% - standard error
- %EXITCODE% - action exit code

### Action Sequences

Action sequencing is treated as an action type and is listed in the drop-down with the other action types in the Policy page. Action sequencing lets you define multiple actions for a rule in a specified sequence and run them as a single action. You can save the sequence of actions you specified with a name and that sequence is saved to the Management DB for repeat usage. You can schedule your action sequences as a job using the Policy, Actions & Rules pages in the user interface. CA SDM support for action sequencing is handled differently from other action types. You can set help desk approval for individual actions running in a sequence, but you cannot set help desk approval for the overall action sequence.

Consider these key points when using action sequences:

- Do not configure sequences that create infinite loops. The action sequence is performed synchronously, but some actions are performed asynchronously. Therefore, if you are expecting certain actions to have completed their tasks when they return, use care. Some actions that are typically long running and asynchronous have a -wait parameter that causes them to wait until their task is complete before returning or after a specified timeout.
- If you attempt to delete an action that is associated with an action sequence, the product prevents you from deleting that action.

- If your action sequence terminates abnormally, it restarts at the last known sequence when the Policy Manager restarts. You can manually cancel an action sequence in progress through the user interface or from a web service.
- When you specify the %STDOUT% (Standard Output), %STDERR% (Standard Error), or %EXITCODE% (Action Return Code) substitution string actions in a custom action running in an action sequence, the standard output/standard error/exit code of the previous action can be piped into the current action. Piping uses the output of the first action as input for the next action. If you redirect the output in your action, then it cannot be piped to the next action. For example, if the custom action *ipconfig* is redirected to a text file named *ipconfig\_output.txt*, then that output is not available for piping to the next action.

## Create a Custom Action

You can create customized action types by defining substitution parameters. Your custom action types are added to the Action Types drop-down list with the predefined action types.

### Follow these steps:

1. In the Explore pane, select the Data Center node.
2. Click Resources, Policy, and then click the Custom Action Types tab.  
The Custom Action Types page appears.
3. Click + (Add).  
The Custom Action Types: Add New section appears.

4. Complete the following fields to define a new action type and a substitution parameter, and then click Save:

### Action Type Name

Defines the name of the new action type.

### Command

Defines the command line structure for the action type. You can define substitution parameters such as %SERVER%, \$MYKEY\$, and so on, for replacement as part of a command. Substitution keys can only be used once per command. For example, the %SERVER% substitution key can only be used once in a command.

### Substitution Key

Defines a unique string for the substitution key. The substitution key name must match what is defined in the command. When defining multiple substitution keys, define each substitution key individually.

### Prompt

Defines the argument name associated with the substitution parameter to input when creating actions.

### Default Value

Defines the default substitution key value.

The new parameter appears in the substitution parameter list.

5. Select Save from the Actions drop-down list.

The custom action type is saved.

## Define an Action Sequence

You can define action sequences for your rules. If the conditions for a rule evaluate to true, the action sequence that you defined runs. You can also create custom conditions and build them into your sequence.

**Note:** The action sequence can also be scheduled as a job or can be run using the `dmpolicy runaction` CLI command.

### Follow these steps:

1. In the Explore pane, select the Data Center node.
2. Click Resources, Policy, and then click the Action tab.  
The Actions page appears.
3. Click + (Add new action).  
The Action Definition: New section appears.



4. Type a meaningful name for the action sequence, and then select Run Action Sequence from the Type drop-down menu.

The Condition Logic section appears.

5. Leave the Restart if Interrupted check box selected to restart the sequence after an abnormal termination. The sequence restarts the last action that was executed and continues. Clear the check box to prevent the sequence from continuing after an abnormal termination.
6. Click + (Add action) in the Action Sequence pane to add actions to the action sequence. Add Action adds a new action at the end of the action sequence. If you want to insert an action in the middle of the sequence, remove all actions after the desired position. Insert the new action, and then redefine the actions that you removed.
7. Select a condition to build your condition logic for the action sequence. New condition logic can only be added to the end of the condition logic sequence. If you want to insert new condition logic in the middle of the sequence, remove all condition logic after the desired insertion point. Insert the new condition logic, and then redefine the condition logic that you removed.
8. Select the type of condition logic evaluation for each additional condition logic sequence. Output Types include the following:

#### **ReturnCode**

Evaluates the action return code.

**Note:** Valid comparison operators for Return Code evaluation are: ==, !=, >, <, >=, <=

#### **STDOUT**

Searches the standard output for a specific string.

#### **STDERR**

Searches the standard error for a specific string.

**Note:** Valid comparison operators for STDOUT and STDERR are "Contains" and "Does Not Contain".

**Note:** You can use the Logic OP field (AND/OR) to link conditions. Logic OP is set to NOOP automatically for the final condition.

The new condition logic is added to the sequence.

9. When you complete your conditions, click Save Condition.

The condition is saved.

10. Click Save in the Action Sequence pane.

The action is saved.

For testing purposes, you can run the action from the Actions page by selecting the action and clicking the Run action icon.

## Define a Schedule

You can schedule actions to run at predefined times. For example, you can use the default Windows scheduler to schedule actions that must be performed every day, or actions that are performed periodically, such as maintenance tasks.

### Follow these steps:

1. In the Explore pane, select the Data Center node.
2. Click Resources, Policy, and then click the Scheduled Actions tab.

The Scheduled Actions page appears.

3. Complete the following fields:

#### **Name**

Defines a name for the scheduled action.

#### **Pre Notification**

Specifies whether to generate an event before the scheduled action runs. The event appears in the dashboard.

#### **Post Notification**

Specifies whether to generate an event after the scheduled action runs. The event appears in the dashboard.

#### **Frequency**

Specifies how often the scheduled action runs: once, daily, weekly, monthly (day), or monthly (day of week).

#### **Date**

Defines a date on which to start the scheduled action.

#### **Time**

Defines a time of day to run the scheduled action.

**Note:** You do not need to enter seconds because they are not used for scheduling jobs.

#### **Type**

Specifies the action type used for the action you are scheduling.

**Note:** The scheduler does not support an action that contains substitution parameters (the only exceptions are %AutoIncrement(0)% and %AutoDecrement(0)%). You can run the actions that contain substitution parameters only through Policy rule evaluation.

**Action**

Lists the actions that have already been created for each action type.

**Note:** The list does not include actions that specify a help desk approval requirement.

4. Select Save from the drop-down list.

A message confirms that your action is scheduled. The scheduled action appears in the list of Scheduled Jobs in the Scheduled Actions list.

**Note:** Actions that specify a help desk approval requirement cannot be used for action scheduling. If you need the same action for a scheduled action, create a second action that does not include the help desk approval requirement.

## Use Cases for Policies

The following scenarios demonstrate some use cases for implementing policies.

**More information:**

[Use Case: Adding a Server to a Service](#) (see page 139)

[Use Case: Adding a New Rule to a Service](#) (see page 140)

[Use Case: Defining an Action](#) (see page 140)

### Use Case: Adding a Server to a Service

This use case illustrates the process for adding a server to a previously created service.

1. Verify the prerequisites for adding the server to the service:
  - The service exists.
  - The server exists.
  - The service already has a priority assigned.
  - The user has access to modify a service.
2. Add the server to the service.
3. Verify the results of adding the server to the service:
  - The server is now a member of the service.
  - The server is now included in the utilization of the service.
  - Inclusion of this service now affects any service rules for utilization.

## Use Case: Adding a New Rule to a Service

This use case illustrates the process for adding a new rule to a service.

1. Verify the prerequisites for adding the rule to the service:
  - The service exists.
  - The user has access to create rules.
  - The servers are in the service.
2. Create the rule definition for this service.
3. Verify the results of adding the rule to the service:
  - The new rule has been created.
  - The new rule is being evaluated for all services that are valid for the conditions of the rule.

## Use Case: Defining an Action

This use case illustrates the process for defining an action for use in scheduling jobs or policy rules.

1. Verify the prerequisites for defining the action:
  - The user has access to define an action.
  - The resources required for the intended action definition have been discovered.
2. Define the attributes of the action and the name of the action in the CA Virtual Assurance user interface.
3. Verify the results of adding the server to the service:
  - The action has been created with the description provided by the user.
  - The action is now available for rules.
  - The action is now available for job scheduling.

**Note:** Actions that specify a help desk approval requirement cannot be used for action scheduling. If you need the same action for a scheduled action, create a second action that does not include the help desk approval requirement.

## Configuring Data Collection

You can control how the following data is collected in your data center:

- The time intervals for metrics collection to occur.
- The systems from which you want to collect metrics (filtering).
- The metrics that you want to collect for each server.
- Data aging and the expiration of data (how long to retain data).

**More information:**

[Key Points About Metrics Collection](#) (see page 141)

[Configure Data Collection for a Data Center](#) (see page 144)

[Configure Data Collection for a Server](#) (see page 145)

[Configure Data Collection for a Virtual Resource](#) (see page 147)

[Configure Performance Thresholds](#) (see page 149)

[Configure the Metric Filter](#) (see page 150)

## Key Points About Metrics Collection

To make informed decisions when you select metrics, review these points to understand CA Virtual Assurance performance and application metrics collection:

- How does CA Virtual Assurance collect metrics data? CA Virtual Assurance communicates with the CA Performance Agent or with the SystemEDGE agent on the remote computer to collect the specified system metrics. The CA Performance Agent is a lightweight version of the NSM Performance Agent. If you have installed the NSM Performance Agent, CA Virtual Assurance can also poll that agent.

**Note:** The CA Performance Agent works differently from the NSM Performance Agent, so the metrics available for collection on Linux can be different, depending on which agent you use.

The CA Performance Agent or the SystemEDGE agent must be installed on any server from which you want to collect the base system metrics, unless you are already using NSM performance agents or SystemEDGE agents. If NSM performance agents or SystemEDGE agents are present, then the CA Performance Agent is not required. If necessary, you can install the SystemEDGE agent using the product user interface. All performance metrics are stored in the Performance DB.

- How is overall utilization calculated? Overall utilization is an aggregate calculation of all the metrics that are currently being collected for servers managed by CA Virtual Assurance. The calculation is based on the value of the metrics and the user-defined thresholds that define the parameters for normal operation. Any new metric that you select for collection is not used in the overall utilization calculation unless you select Include for Overall Calculation in the Policy, Metrics, Thresholds section of the user interface. When you do this, CA Virtual Assurance does not provide false results when evaluating the state of the servers.
- How is overall utilization impacted by metric evaluations? The metric details provided in the tables help you understand how CA Virtual Assurance evaluates the different metrics. Each metric has a method property set to either *exact* or *complement*. A higher exact value is a worse scenario than a lower exact value because it indicates an increase in overall utilization. A higher complement value is a positive scenario because it indicates a decrease in overall utilization. Generally, a high exact value negatively impacts overall utilization and a low exact value positively affects overall utilization. By contrast, a high complement value positively impacts overall utilization, and a low complement value negatively affects overall utilization. For example, if the value of Memory: Percentage Committed Bytes In Use increases, overall utilization of the system increases. If the value of Memory: Available MB increases, overall utilization decreases.

- What are the default metrics? The default metric definitions are located in the metric list in the Filter section for all supported platforms. You can find the default metrics indicator on the metric list with the value Yes in the Default column. CA Virtual Assurance uses this list to obtain the metric definitions when you add a new server. You can configure platforms, types, subtypes, instances, and the type of data to collect in the Filter section. The metric filter and definitions for each server are stored in the Performance DB.
- Is performance data currently available for my systems? By default, if data cannot be collected, CA Virtual Assurance does not negatively affect server state because lack of data does not reflect server criticality. By reviewing the Events list or selecting a specific system, you can determine whether metric data is being collected. However, if a more immediate means of determining this is needed, or if performance data is critical, CA Virtual Assurance can be configured to change the state of a system automatically to Warning or Critical if performance data cannot be collected. To enable easy identification of systems where performance data is not available, modify the `caaipconf.cfg` file located in the CA Virtual Assurance `install_path\conf` directory. Open the file with a text editor and locate the health state property as follows:

```
<property name="CONFIG_KEY_DEFAULT_HEALTH_STATE">
  <!-- Valid values: 0 (Unknown); 5 (OK); 10 (Warning); 15 (Minor Failure); 20 (Major Failure); 25
(CriticalFailure) -->
  <!-- Changes the value of HealthState for the CA_CollectionState object associated to the
CA_ComputerSystem -->
  <!-- If set to 30, CE will not set the HealthState. -->
  <value>5</value>
  <displayName>Default node health state when problem encountered in metric or data
collection</displayName>
</property>
```

By modifying the value surrounded by the value XML elements to one of the other supported values such as 5 or 10, for OK or Warning respectively, CA Virtual Assurance reflects the desired state when performance data cannot be collected. For example:

```
<property name="CONFIG_KEY_DEFAULT_HEALTH_STATE">
  <!-- Valid values: 0 (Unknown); 5 (OK); 10 (Warning); 15 (Minor Failure); 20 (Major Failure); 25
(CriticalFailure) -->
  <!-- Changes the value of HealthState for the CA_CollectionState object associated to the
CA_ComputerSystem -->
  <!-- If set to 30, CE will not set the HealthState. -->
  <value>10</value>
  <displayName>Default node health state when problem encountered in metric or data
collection</displayName>
</property>
```

Because `<value>` was changed to "10", systems that do not have performance data available are displayed in a warning state in the CA Virtual Assurance user interface.

**Note:** For a list of performance metrics and descriptions, see the *Reference Guide*.

## Configure Data Collection for a Data Center

You can configure data collection at the Data Center level. The Data Center level policy takes effect immediately.

**Follow these steps:**

1. Click Resources, and select the Data Center folder in the Explore pane.
2. Right-click, and select Policy, Configure Collection Settings.

The Settings dialog appears.

3. Complete the following fields in the Collection Setting section:

**Data recording interval (seconds)**

Defines how often the data is collected and stored in the Performance DB.

**Default:** 300 seconds

**Polled data retention (days)**

Defines how long to store the polled data in the Performance DB. Consider the number of managed systems, services, and metrics collected when defining this number. The stored polled data objects accumulate over time and can impact performance. If performance issues arise, decrease the number of retention days.

**Default:** 14 days

**Daily rollup data retention (days)**

Defines how long to store the average of the daily data in the Performance DB.

**Default:** 365 days

4. Enter the threshold limits in the Thresholds section and then click Save.

Your settings are saved.



## Configure Data Collection for a Server

You can configure data collection for individual servers. Use this procedure to configure specific servers to collect data for the data center. You can also select metrics to monitor, set threshold values for individual metrics, and include metrics in overall utilization.

**Follow these steps:**

1. Open the Explore pane.  
Available groups, services, and systems appear.
2. Open the Data Center folder, and select the service to which the server belongs.
3. Right-click and select Policy.  
The Policy submenu appears.
4. Click Metrics.  
The Metrics wizard opens.
5. Select the server for which you want to configure data collection.
6. Complete the following fields in the Set Interval dialog:

**Use Default**

Specifies the data center level as the default when selected. If you leave the check box cleared, the values that you specify are used instead.

### **Data recording interval (seconds)**

Defines how often the data is collected and stored in the Performance DB.

**Default:** 300 seconds

### **Daily rollup data retention (days)**

Defines how long to store the average of the daily data in the Performance DB.

**Default:** 365 days

### **Polled data retention (days)**

Defines how long to store the polled data in the Performance DB. Consider the number of managed systems, services, and metrics collected when defining this number. The stored polled data objects accumulate over time and can impact performance. If performance issues arise, decrease the number of retention days.

**Default:** 14 days

7. Select the metrics to monitor from the Available Metrics section and then click the down-arrow.

The selected metrics are moved to the Selected Metrics to Collect section.

**Note:** If you disable the default metrics (CPU and memory) and enable others, you will not see an overall utilization until you modify the thresholds of the newly selected metrics.

8. You can configure which performance metrics to monitor for each server and set threshold boundaries for each metric. Select the metric for which you want to set thresholds and complete the following fields:

#### **Upper Threshold**

Defines the upper limit of utilization for the selected metric group.

**Default:** 80%

#### **Lower Threshold**

Defines the lower limit of utilization for the selected metric group.

**Default:** 20%

#### **Include for Overall**

Specifies that you want the selected metrics to be included in the overall utilization calculation and evaluated by CA Server Automation.

9. Click Finish to save your settings.

## Configure Data Collection for a Virtual Resource

You can configure data collection for virtual platforms and the virtual resources created and managed on those platforms. Use this procedure when you want to configure specific virtual machines or other resources to collect data at an interval that differs from the default for the data center. You can also select metrics to monitor, set threshold values for individual metrics and include metrics in overall utilization.

You can configure data collection for the following virtual platform objects:

- vCenter Server
- vCenter Data Center
- vCenter ESX Server
- vCenter Virtual Machine
- Hyper-V
- Microsoft Clusters
- Microsoft Cluster Nodes
- IBM AIX LPAR Server
- IBM Logical Partition
- Solaris Zones Server
- Solaris Zone

### To configure data collection for a virtual resource

1. Click Resources, and open the Explore pane.
2. Expand the Data Center or MS Cluster Service folder, then any subfolder, and select the object that you want to configure.

Subtabs for that object appear in the right pane.

**Note:** If you select a top-level folder (such as VMware vCenter Server) or an object for which no data is collected (such as a vCenter cluster), you must select the specific object contained within the folder or object for which to configure data collection.

**Note:** If you select MS Cluster Service as the top-level folder, then you see clusters and their nodes.

3. Right-click and select Policy, Configure Server Metrics Collection.

**Note:** If you select the top-level folder for Solaris Zones, the Hardware Class column in the System section always shows the value Other.

4. Select the metrics that you want to monitor from the Available Metrics section and then click the down arrow.

The metrics you select move to the Selected Metrics to Collect section.

**Note:** If you disable the default metrics (CPU and memory) and enable others, you will not see an overall utilization until you modify the thresholds of the newly selected metrics.

5. Click Save to apply the selected metrics.
6. Right-click the resource, and select Policy, Configure Collection Settings.
7. Complete the following fields in the Collection Setting section:

**Use Default**

Specifies the data center level as the default when selected. If you leave the check box cleared, the values that you specify are used instead.

**Data recording interval (seconds)**

Defines how often the data is collected and stored in the Performance DB.

**Default:** 300 seconds

**Daily rollup data retention (days)**

Defines how long to store the average of the daily data in the Performance DB.

**Default:** 365 days

**Polled data retention (days)**

Defines how long to store the polled data in the Performance DB. Consider the number of managed systems, services, and metrics collected when defining this number. The stored polled data objects accumulate over time and can impact performance. If performance issues arise, decrease the number of retention days.

**Default:** 14 days

8. Click Save to save your settings.

**Note:** The default thresholds are used. If you want to modify thresholds, you must do this separately.

## Configure Performance Thresholds

You can configure which performance metrics to monitor for each server and set threshold boundaries for each metric.

**Follow these steps:**

1. Open the Explore pane.  
Available groups, services, and systems appear.
2. Expand the Data Center folder and any subfolder, then select the server that you want to configure. Navigate to a virtual server to select a specific virtual resource, such as a virtual machine or logical partition.
3. Right-click and select Policy.  
The Policy submenu appears.
4. Click Configure Threshold Settings.  
The Configure Threshold Settings appears.
5. Select the metric for which you want to set thresholds and complete the following fields:

**Upper Threshold (%)**

Defines the upper limit of utilization for the selected metric group.

**Default:** 80%

**Lower Threshold (%)**

Defines the lower limit of utilization for the selected metric group.

**Default:** 20%

**Include for Overall Utilization Calculation**

Specifies that you want the selected metrics to be included in the overall utilization calculation and evaluated by CA Virtual Assurance.

6. Click Modify to save your settings.

## Configure the Metric Filter

You may want to add or delete metrics to or from the metric filter for the Data Center, depending on which performance metrics you want to monitor.

**Note:** For information about the performance metrics available for NSM Performance Agent users, see the *NSM Systems Performance Metrics Help* in the NSM documentation set.

### To configure the metric filter

1. Select the Data Center folder in the Explore pane.
2. Right-click and select Policy, Configure Collection Criteria.  
The Collection Criteria dialog appears.
3. Do *one* of the following:
  - Select the check box for an existing metric to modify an existing entry. The information for the selected metric populates the fields of the Details section. Make any changes and click Update.
  - Select an OS, and complete the fields in the Details section to add a new metric, then click Add.

The metric is saved.

The Details section contains the following fields:

#### OS

Defines the operating system for the metric being monitored.

#### Type

Defines the type of metric being monitored.

##### Example:

Type: CA Disk Group

Sub Type: Writes per second (average)

#### Sub Type

Defines which aspect of the metric is being monitored.

##### Example:

Type: CA Disk Group

Sub Type: Reads per second (average)

**Instance**

Defines the instance of the managed object in the MIB hierarchy.

**Example:**

Type: vmvcaim.StatClusterEffectiveCPU

Sub Type: 1.3.6.1.4.1.546.16.52.2.7.2.1.14

Instance: %3 [%2]

%<n> where <n> is the numeric value listed under Instance matched to any value corresponding to the n<sup>th</sup> column in the respective AIM MIB table. For example, vmvcAimStatClusterTable for all row entries (instances for the same managed object). This is useful when collecting metrics for the managed object instantaneously for all instances when they are available with no user input.

**Upper Threshold (%)**

Defines the upper limit of utilization for the selected metric group.

**Default:** 80%

**Lower Threshold (%)**

Defines the lower limit of utilization for the selected metric group.

**Default:** 20%

**Lag**

Defines how many consecutive times the threshold breach occurs before a threshold event is generated. Configure this option to avoid flooding events for threshold evaluation. You can define an action to log threshold breach events and set up rules for threshold monitoring.

**Method**

Specifies whether the collection method is complementary, complementary delta, exact, or exact delta. The complementary method includes metrics that are not already included in a subset of that set. The exact method collects the exact metric specified.

**Category**

Specifies whether the monitored metric is a system, application, or SNMP metric.

**Default Selected Metric(s) for Collection**

Specifies whether the collection engine collects the metrics specified by the filter by default. Unless a metrics filter is set as default, the collection engine will not automatically collect the metrics specified.

**Include for Overall Utilization Calculation**

Specifies that you want the selected metrics to be included in the overall utilization calculation and evaluated by CA Virtual Assurance.

### Activate for Collection

Specifies that the metric filter is effective for use when evaluating what metrics are available for collection.

4. Select the check box for any metrics that you want to delete, then click Delete.  
The selected entries are deleted.

## UCS Action Types for Policies

CA Virtual Assurance provides Cisco UCS action types for blade power operations and service profile operations.

### More information:

[Create a Blade Power Action](#) (see page 152)

[Service Profiles](#) (see page 99)

[How to Create or Update a Service Profile](#) (see page 100)

[Associate Service Profiles with Blades](#) (see page 155)

## Create a Blade Power Action

You can define actions to conduct blade power operations.

### To create a blade power action

1. Click Resources, and open the Explore pane.
2. Select the Data Center node, and click Policy.
3. Click Actions.  
The Actions page appears.
4. Click + (Add new action).  
The Action Definition: New page appears.
5. Click Configure Power on the Type drop-down list.
6. Click Cisco UCS on the Environment drop-down list.  
The Configure Power form displays.
7. Specify the blade that you want to have the power operation performed on by selecting the UCS Manager and chassis which contains the blade.
8. Select a power operation, and click Save on the Actions drop-down menu.  
The blade power action is created.



## Service Profiles

A *service profile* contains configuration information about Cisco UCS hardware, including interfaces, fabric connectivity, and network and server identity. You can create a service profile for a specific UCS manager or centrally in the CA Virtual Assurance Management DB. Service profiles on a UCS manager can be *imported* into the Management DB, from which they can be *exported* to other UCS managers.

Service profiles allow supporting Cisco UCS hardware to be abstracted from the operating system. By addition or removal of a service profile, services can be brought online or taken offline. By reassignment of a service profile from one blade to another, a set of services (operating system and applications) can be moved to different servers.

Service profile information can include:

- Blade servers by device UUID or virtual UUID pool
- Storage in any configuration of local storage, RAID, or SAN (HBA, WWNN, and WWNN pools)
- Networking (MAC, vNIC 0, vNIC 1, and MAC pools)
- Server boot order or other policies
- Server assignment type (assign later, provision a slot in advance, provision an existing server, select server from pool)

## How to Create or Update a Service Profile

CA Virtual Assurance provides a wizard that administrators can use to create service profiles, and you can update service profiles with predefined policy options. System, network, and storage administrators can collaborate to create service profiles with unique identity characteristics, and required connectivity characteristics.

The service profile wizard provides a subset of the Cisco UCS Manager interface to take advantage of knowledge and experience in the Cisco environment.

### Example: Create a Service Profile

Using the CA Virtual Assurance user interface, select the create-service-profile option and specify the server profile name to create and choose whether the server profile is to be:

- Hardware-based
- Simple server profile with default networking and storage connectivity
- Based on an existing service profile template
- Custom service profile which must be created explicitly

Based on the option selected, the wizard leads you through an interview process to obtain the required identity and connectivity information. You can take the defaults or specify identity (UUID), network (MAC/VLAN), storage (WWN/vHBA), and boot policy information explicitly.

## Associate Service Profiles with Blades

Services profiles can be associated with blades, unassociated, or set to apply at failover.

### To adjust a UCS service profile

1. Right-click and select Policy.  
The Policy submenu appears.
2. Right-click and select Policy, Actions & Rules.  
The Actions & Rules page appears.
3. Click Actions.  
The Actions page appears.
4. Click + (Add new action).  
The Action Definition: New page appears.
5. Click the action type Configure Service Profile on the Type drop-down list.  
The Configure Service Profile form displays.
6. Specify the UCS resource details to which you want the service profile to apply.  
Select the profile operation.  
**Note:** If help desk approval is required, enter information as needed.
7. Click Save on the Actions drop-down.  
The service profile relationship is modified.

## Policy Configuration

You can configure managed agents and apply the configuration to multiple systems in one operation using centralized remote deployment from the CA Virtual Assurance user interface. Policy configuration lets you configure SystemEDGE and the SRM AIM in a centralized location and distribute the policy across the enterprise in a consistent, reliable, and secure manner.

Remote policy configuration using CA Virtual Assurance provides the following benefits:

- The ability to create platform independent monitoring policies to use across monitoring platforms
- The ability to apply configuration policies to single servers or groups of servers
- The ability to create monitoring templates that you can combine into one policy

- An audit trail of configuration events and actions
- The ability to track policy compliance across the enterprise through events and reports
- Integration with the deployment solution, and, similar to deployment, only a minimal footprint on the target system
- Scalability to thousands of concurrent configurations
- Support for multiple agent configuration sources (CA Virtual Assurance, NSM, SystemEDGE, and so on), and the ability to accept or reject changes through CA Virtual Assurance
- The ability to remotely control the AIMs loaded by SystemEDGE
- The ability to import existing SystemEDGE configurations for use in future policy configuration
- Pick lists during configuration for many monitor definitions, eliminating the requirement of entering individual OID numbers
- Automatic monitor index assignment that eliminates the need to manually define indexes and avoids conflicts

### More Information

[How to Create SystemEDGE Policy](#) (see page 156)

[How to Create SRM Policy](#) (see page 166)

[Apply Policy to Machines](#) (see page 169)

[Review Policy Application Progress](#) (see page 170)

[Configure and View Applied Policies](#) (see page 171)

[Agent Policy Dashboard Views](#) (see page 173)

[Example: How to Monitor User-specific Metrics \(MIB Extensions\)](#) (see page 173)

[Example: How to Monitor a Specific Windows Performance Registry Metric](#) (see page 175)

## How to Create SystemEDGE Policy

You create SystemEDGE policy to define a set of monitors, AIMs to load, configuration preferences, and other settings that control how the agent runs and what it monitors. Once you create a policy, you can apply it to any number of systems running SystemEDGE agents in managed mode. Policy lets you perform all configuration operations that you can manually configure locally with the benefit of a consolidated interface, pick lists, and dynamic deployment to remote systems.

The following process describes how to create SystemEDGE policy:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.  
The SystemEDGE pane appears.
2. Click + (New) on the Available Policies toolbar.  
The New SystemEDGE Policy dialog appears.
3. Enter a name and description for the policy and whether to base it on an existing policy and click OK.  
The policy is created, and a configuration screen appears in the right pane.
4. [Define monitors to include](#) (see page 157).
5. [Define control settings](#) (see page 158).
6. [Define SNMP settings](#) (see page 163).
7. [Define MIB extensions](#) (see page 166).
8. Click Save Policy.  
The policy is saved.

## Monitor Definition

When configuring policy, click the Monitors tab to view a summary of all monitors currently defined. In the Monitors tab, subtabs exist for defining the following monitors:

- Self monitors
- Process monitors
- Log file monitors
- Windows event monitors
- History monitors
- Process group monitors

From each of these subtabs, you can create a new monitor, delete or modify existing monitors, or copy a monitor. Monitor definition using configuration policy provides the following functional benefits that are not available through local file manipulation:

- Intuitive field names
- Automatic index creation
- A platform designation that lets you limit a monitor to a specific platform

- Pick lists for object class and attribute for self monitors to eliminate the need to remember and enter individual OID numbers
- Pick lists for key properties in other monitors
- A monitor type designation for log file monitoring that lets you specify whether you are monitoring a log file or directory and adjusts the available fields dynamically
- Individual check boxes for setting flags to eliminate the need to provide a hexadecimal string for flag settings

**Note:** See the `sysedge.cf` configuration file or the *SystemEDGE User Guide* for descriptions of these hexadecimal flags.

- Maintenance window settings integrated into the monitor

For more information about defining individual monitors for configuration policy, see the *Online Help*.

## Define SystemEDGE Policy Control Settings

You can control the following agent behavior using the SystemEDGE policy control settings:

- Security settings
- SNMP settings
- MIB table population
- UNIX settings
- Performance monitoring settings

You can apply the control settings defined in the policy to all machines.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Control Settings.

The Controls page appears.

4. (Optional) Click Use Defaults.

The default selections pane appears. You can change the default settings.

5. Configure the following control settings:

**SNMP**

Lets you define the following basic SNMP properties:

**Bind Address**

Specifies an interface to which the agent binds and listens for incoming SNMP requests. Valid addresses are IPv4 or IPv6 address.

**Note:** The corresponding default \_port is specified during installation.

**Bind Port**

Specifies the trap port the agent binds to for sending SNMP traps. If no bind\_address is specified the agent binds to all available UDP addresses.

**Default:** Port selected by the system

**IP Family**

Specifies the agent communication method: IPv4 only, IPv6 only, or both. By default, the agent tries using IPv4 and then IPv6.

**FIPS Mode**

Specifies the agent to use FIPS-compliant encryption. Select Non-FIPS Mode to enable the CA eTrust Public Key Infrastructure libraries, and if this method fails, fall back to the internal minimum security solution. Select FIPS Co-existence Mode to enable FIPS-compliant encryption, and if this method fails, fall back to the CA eTrust Public Key Infrastructure Libraries. Select FIPS Only Mode to enable the RSA BSAFE Crypto-C Micro Edition FIPS-compliant libraries and perform no encryption if they fail.

**Default:** Non-FIPS Mode

**Trap Source**

Specifies the source address used to send traps. Valid addresses are IPv4, IPv6 address, or a host name.

**Default:** Host name of the agent

**Security Settings**

Lets you define the following security preferences:

**Authentication Traps**

Sends an authentication failure trap when the agent receives an SNMP message with a community name that the agent cannot recognize.

**Default:** Disabled

**Process Sets**

Permits access to processes and other software running on agent systems in the Process table and Running Software table. Allowing SNMP Sets on these tables can cause security issues.

### Remote Shell Group

Permits management systems to remotely instruct the agent to run shell scripts and programs on the agent system through the Remote Shell group. The disclosure of this type of information can post a potential security risk.

### Execution Action

Enables the execution of action commands with the monitoring tables when a threshold breach occurs. The capability to run action commands and scripts can be a security issue.

### MIB Table Population

Populates the following tables in the Systems Management MIB:

- Process Table
- User Group Table
- Who Table
- Trap Community Table
- Monitor Mirror Table
- Aggregate Mirror Table
- Top Processes Table

Each table either contains sensitive information that you can expose in a MIB or non-essential information that you can disable to save disk space. The default settings enable population of all tables except for the process table.

### Miscellaneous

Lets you define the following miscellaneous settings:

#### Allow agent to be Updated using SNMP

Permits agent updates using SNMP Sets (for example, removes write communities). If you permit SNMP Sets on the agent, any updates through this method cause a notification of an SNMP Set change and also an exception when viewing policy details for the system.



**Notify Manager of Configuration Updates**

Enables the agent to send a notification to the manager for any SNMP Set request that the agent processes.

**Warm Start Discovery**

Enables an agent rediscovery of all devices after every warm start configuration update. If you manage a system with many devices, a discovery after every warm start can consume too much time and too many resources.

**Use Perl Compatible Regular Expressions**

Perl Compatible Regular Expressions (PCRE) enables you to specify i18n compatible regular expressions while defining monitors that support regular expressions. The examples of regular expressions are log file, process, process group, Windows services and Windows events. You can also use this option to create more complex regular expressions. This option is provided in SystemEDGE agent 5.1.0 and above versions.

**Automatically Resolve Index Conflicts**

Enables you to resolve Index conflicts. When you apply the layered templates to all machines, indexes are assigned to the monitors added in the template. If the assigned indexes conflict with existing indexes either within the base policy or another template, this option reassigns unique index values.

**Note:** Indexes contained within the base policy are always maintained in the delivered configuration. If this option is disabled, you cannot resolve conflicting indexes. However, when you apply layered templates to the machines, the conflicting indexes are displayed as errors on the layered templates that caused the conflicting indexes.

**Historical Performance Monitoring**

Lets you define the following settings for the Performance Cube AIM, which collects history information into Systems Performance cubes for historical performance management:

**Collection Interval**

Specifies how often to collect information from the History table into performance cubes.

#### **Index Range Start**

Specifies the beginning of the reserved range of indices, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

#### **Index Range End**

Specifies the end of the reserved range of indices, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

#### **UNIX Control Settings**

Lets you define the following settings for agents running on UNIX systems:

##### **Sub-program Group**

Specifies a group name other than root under which to run subprograms.

##### **Sub-program User**

Specifies a user name other than root under which to run subprograms.

##### **Linux Freemem Include**

Specifies whether to include system buffers, disk cached memory, or both in free memory calculation.

#### **Query System Devices**

Lets you enable querying of the following system device metrics:

- Serial device status
- Floppy disk status
- Disk size, capacity, description, and other properties (Probe Disks)
- NFS file system status
- HP-UX graphics status

Querying these metrics can cause issues with potential agent blocking. The default settings enable querying of only serial device status and NFS file system status.

6. Click Plugins.

The Plugins pane appears. This pane controls which AIMs to load with the agent.

7. Do one of the following:

- Select 'Load all available plugins' to load all AIMs available on the agent system.
- Select 'Load plugins selected in the table'.
- Click + (New) on the External Plugins toolbar to add an AIM to the External Plugins table.

**Note:** For more information about available AIMs, see the *SystemEDGE User Guide*.

AIM loading is configured.

8. Click Aggregate Monitors.

Configure aggregate monitors as described in [Configure Object Aggregation](#) (see page 108).

The control settings are defined.

9. Click Save Policy.

The policy is saved.

**More Information:**

[Configure Object Aggregation](#) (see page 108)

## Define Traps and Communities

SNMP settings define the communities that the agent uses and the destinations to which it sends traps.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Traps and Communities tab.

The Communities page appears.

4. Select one of the following and click Action, select Apply:

- Include only Server specific SNMP settings
- Include Server specific SNMP settings and all Default settings
- Select 'Include Server specific SNMP settings and selected Default settings

The SNMP settings are updated and the community page in the Communities table displays the following:

**Name**

Specifies the name of the community string.

**Port**

Specifies the port of SNMP.

**SNMP Version**

Specifies the SNMP version that the community uses.

**Access Rights**

Specifies whether the community should have read write or read-only permissions.

**Note:** Add at least one read-only and one read write community.

**Community/User**

Specifies the community name.

**Authentication Protocol**

Specifies the protocol to authenticate SNMPv3 data.

**Privacy Protocol**

Specifies the protocol to authenticate SNMPv3 data.

**Access Control List**

Specifies a space separated list of IP addresses to restrict community usage to those addresses only. If you leave the list blank, the agent grants access to any system that uses the associated community name. Access lists are only for communities that use SNMPv1.

**Note:** For information about defining SNMPv2c and SNMPv3 access lists, see the *SystemEDGE User Guide*.

5. (Optional) Add, update, or delete other communities as necessary.

6. Click Save Policy.

The policy is saved.

7. Click Trap Destinations.

The Trap Destinations page appears.

8. Define a trap destination using the following controls and click Add:

**Trap Type**

Specifies the type of trap to send, depending on the SNMP version.

**Destination**

Specifies the IPv4 or IPv6 address to which to send traps.

**Port**

Specifies the UDP port to which to send traps.

**Community**

Specifies the community name sent with the traps.

**Encoding**

(Optional) Specifies how to include the source address you defined in the Trap Source field of the Control Settings pane in traps. This parameter is important if the trap source translates to an IPv6 address. Enter the encoding parameter in a three digit format XYZ, assuming leading zeros.

**Default:** 000

**X**

Controls extending the four byte IPv4 source address field (SNMPv1 traps only). Enter 0 to not extend the source address field to include the 16 byte IPv6 address, and enter 1 to extend the source address field.

**Y,Z**

Controls the inclusion of source information into the trap's varbind (Y) or UDP packet (Z; SNMPv1 traps only). Enter one of the following for these digits:

**0:** Do not modify the trap's varbind or the outer UDP packet.

**1:** Include the trap\_source parameter as is in the varbind or packet (IPv4/IPv6 address or host name).

**2:** Include the trap\_source parameter preferably as an IPv4 address (then IPv6 address, then host name).

**3:** Include the trap\_source parameter preferably as an IPv6 address (then IPv4 address, then host name).

**4:** Include the trap\_source parameter preferably as a host name (then IPv4, then IPv6).

**5:** Follow the preference for 2 and include the host name.

**6:** Follow the preference for 3 and include the host name.

**7:** Follow the preference for 1 and include the host name (if trap\_source is an IPv6 address).

### Trap Source

(Optional) Specifies the IPv4 or IPv6 address or the host name to use as trap source.

**Default:** Global Trap

The trap destination appears in the Defined Trap Destinations table.

9. (Optional) Add, update, or delete other trap destinations as necessary.

10. Click Save Policy.

The policy is saved.

**Note:** For more information, see the *SystemEDGE User Guide*.

## Define MIB Extensions

Defining MIB extensions provide functional benefits that are not available in local file manipulation. The policy configuration feature provides field names and the list of key properties such as object type.

When you configure a policy or a monitoring template, click the MIB Extensions tab to add the following objects:

- MIB Extensions
- Windows Performance
- Windows Registry

**Note:** MIB Extensions within templates are supported for the purposes of applying the MIB Extensions directly to monitored systems. MIB Extensions for use within policies should be created directly in the Policy itself.

## How to Create SRM Policy

You create SRM policy to define tests to perform, thresholds to monitor, configuration preferences, and other settings that control how the agent runs and what it monitors. Once you create a policy, you can apply it to any number of systems running SystemEDGE agents with the SRM AIM in managed mode. Policy lets you perform all configuration operations that you can manually configure locally with the benefit of a consolidated interface, pick lists, and dynamic deployment to remote systems.

The following process describes how to create SRM policy:

1. Click the Resources tab, open the Configure pane, expand Policies, then click Service Response.

The Service Response pane appears.

2. Click + (New) on the Available Policies toolbar.

The New Service Response Monitoring Policy dialog appears.

3. Enter a name and description for the policy and whether to base it on an existing policy and click OK.

The policy is created, and a configuration screen appears in the right pane.

4. [Define tests to include](#) (see page 167).

5. [Define test thresholds](#) (see page 167).

6. [Define control settings](#) (see page 168).

7. Click Save Policy.

The policy is saved.

## Test Definition

When configuring SRM policy click the Tests tab to view a summary of all tests currently defined. From the Tests pane you can define new tests and modify, delete, or copy existing tests. Defining tests using configuration policy provides the following functional benefits that are not available through local file manipulation:

- Intuitive field names
- Pick lists for key properties
- A pick list to specify the test type, after which the pane dynamically updates with the properties relevant to the test you specify
- A table for adding threshold definitions to a test

For more information about defining test for configuration policy, see the *Online Help*.

## Threshold Definition

You can define threshold monitors for any attribute collected by an SRM test and add the monitors to any existing test. These monitors create SystemEDGE self monitor entries and let you track status information for SRM tests. Defining monitor thresholds using configuration policy provides the following benefits over manual file manipulation:

- Intuitive field names
- A pick list for choosing the attribute to monitor
- Fields for warning, minor, major, critical, and fatal values that provide the flexibility to define as many severities as necessary
- The ability to dynamically associate a threshold with any SRM test in the Test tab

For more information about defining monitor thresholds for tests, see the *Online Help*.

## Define SRM Control Settings

SRM control settings define various aspects of AIM behavior that you typically control in the svcrsp.cf file, including the following:

- Security settings
- Log level
- Index reservations

Control settings defined in SRM policy are applied to all machines to which you apply the policy.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click Service Response.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Control Settings tab on the policy page.

The Control pane appears.

4. Configure the following control settings:

#### Maximum Number of Threads

Specifies the number of threads the jcollector should use to perform tests

#### Log Level

Specifies the log level of the SRM AIM. Default is Warning.

#### Allow External Scripts

Specifies if execution of external scripts is allowed.

#### Allow Execution of File I/O Tests

Specifies if execution of file I/O tests is allowed.

#### Allow Untrusted SSL Certificates

Specifies if SSL tests with sites that do not have trusted SSL certificates is allowed.

#### Java bin Location

Defines the location of the Java executable.

**Note:** Specify the complete path and binary on AIX.



**Override CLASSPATH in Environment**

Defines extra classes to load. Overrides CLASSPATH in environment if defined.

**No Collector**

Specifies if SystemEDGE should start jcollector.

**Bypass JRE Internal Cache**

Specifies whether to bypass JRE internal cache.

**No TOS for IPv4 (HP-UX)**

Specifies whether to disable TOS.

**Shared Memory Name**

Defines the ID for the shared memory.

**Reserved Test Indices**

Defines reserved range of test indexes.

The control settings are defined.

5. Click Save Policy.

The policy is saved.

## Apply Policy to Machines

After you create configuration policy, apply it to machines across the enterprise. When you apply configuration policy, CA Virtual Assurance pushes a compiled configuration file containing all policy settings to all specified agent machines, and the new policy is implemented after an automatic agent warm start.

You can also reapply the policy to machines if one of the following occurs:

- You updated the policy
- You received a notification that the configuration on an agent machine has changed

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Select the policy you want to apply.

Policy details appear in the right pane.

4. Click Action and select Apply.

Tabs appear for selecting machines on which to apply the policy. The 'Update machines running this policy' tab lets you apply the policy to machines that are already running the policy. The 'Apply to Machines not running this Policy' tab lets you apply policy to machines without any policy or using a different policy.

5. (Optional) Do one of the following from the 'Update machines running this policy' tab:

- Select 'Update all machines using this policy' to deploy the policy on all machines currently running it. This option is useful if you have made configuration policy changes that you want to apply globally.
- Select 'Update selected groups of machines' to update only machines that meet any of the following criteria:

- Running an out of date version of the policy
- Policy exceptions have occurred
- Running a current version of the policy

Select any of these options. Policy exceptions occur when a user applies a point configuration change to an agent that is not represented in the applied policy.

- Select Advanced and manually the machines in the Select Machines pane to which the policy is to be reapplied.

6. (Optional) Select machines from the 'Apply to Machines not running this Policy' tab to which to apply the policy.

7. Click Apply Policy.

The policy application is initiated.

## Review Policy Application Progress

You can review the progress of policy application operations in detail for each individual policy.

### Follow these steps:

1. Select the Resources tab, open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Managed Machines tab.

The Managed Machines page appears with a list of machines currently running the policy that lets you view the configuration status.

4. (Optional) Click View Exceptions.

The Policy Exceptions pane appears and lets you view SNMP sets that have been applied to the system since the policy was last applied.

**Note:** This screen is only available for SystemEDGE Policies.

5. (Optional) Click View Configuration.

The Policy Configuration page appears and lets you view the configuration file delivered for the agent.

6. (Optional) Click View Errors.

The Policy Errors pane appears. If the policy could not be successfully applied, you can view a list of errors returned by the agent when the policy was rejected.

## Configure and View Applied Policies

The Policy feature lets you manage the policies and templates applied to an individual server, a server group, or a service. You can perform the following operations:

- Update Policies and Templates
- View exceptions. Exceptions are SNMP sets that have been applied to the system since the policy or template that was last applied.
- View policy configuration
- View policy errors
- Bulk update Policies
- Delete Templates

### Follow these steps:

1. Open the Explore pane.

Available groups, services, and systems appear.

2. Select a system or a service. Click the Resources page, and then Monitoring Software.

The Machine Details page appears.

3. Click Policies.

The Policies page displays a list of Policies of SystemEDGE and SRM, and SystemEDGE Templates.

**Note:** The Filter can display a list of layered templates that are Pending, Delivered (successful), Configured, and Failed.

4. You can bulk update policies and templates. In the Policies and Templates table, select the policy or template you want to bulk update, click Actions, and select one of the following options:

- Bulk update SystemEDGE Policies.
- Bulk update Service Response Policies.
- Bulk update SystemEDGE Templates.

**Note:** If the policy is being applied to a single server, you are prompted for the policy name.

**Bulk update for Policies:**

If the selected policy is applied to a service group, a dialog provides an option to select the machine on which you want to apply the policy.

**Bulk update for Templates:**

A dialog provides an option to select the templates from the Available Templates. After you select the templates, click one of the following options:

- Apply the selected templates as the set for all machines - Removes the existing templates that are applied to all machines and applies the selected templates to all machines.
- Apply the selected templates to the existing set already applied to all machines - Adds the selected templates to all machines. If the templates selected include the existing templates that are applied to machines, the existing templates are reapplied.

5. Click Apply Policy to apply the policy or template to the machines.

On the Policies page, you can view the progress of the policy or template being applied to the machines.

6. (Optional) Click View Configuration Icon.

The Policy Configuration page appears. For a machine with a template, it displays the Policies and Templates, and SystemEDGE Configuration file. For a machine with a Service Response Monitor, it additionally displays the Service Response Monitor Configuration file.

7. Click Save Policy.

The policy is saved.

## Agent Policy Dashboard Views

The following views are available on the Dashboard for tracking agent policy assignments:

### Policy Status Summary

Displays a pie chart and list showing the number of policies. A system can be in five different states:

#### Unconfigured

SystemEDGE agent is installed but no policy is configured.

#### Agent Installed

SystemEDGE agent is installed.

#### Configured

SystemEDGE agent is installed and a policy is configured.

#### Configuration Error

SystemEDGE is installed and a policy is configured but the last configuration failed.

#### Installed but not managed

SystemEDGE is installed but running in a mode that cannot be managed by policy configuration.

### Policy Breakdown

Displays a pie chart and list showing all policies and how many systems contain each policy.

### Machines with nonstandard Policies

Displays a table listing systems that contain nonstandard changes to an applied policy.

## Example: How to Monitor User-specific Metrics (MIB Extensions)

This step-by-step example describes how to monitor a user-specific metric.

### How to monitor user-specific metrics (MIB extensions)

1. Create a program that returns the data required. For example, a simple DOS batch script on the agent system to return some fixed data.

```
@echo off  
echo 99
```

2. Open a text editor and store these two lines in data.bat on the C: drive.

3. Create an MIB extension that references this batch file.
  - a. From the user interface, click *Policy*, open *Configuration* in the navigation pane, expand the Policy tree, and open a SystemEDGE policy.  
The policy details appear in the right pane.
  - b. Click the MIB Extensions tab.  
The MIB Extensions pane opens.
  - c. Add the following data into the fields:  
Index: 1 (if it is the first MIB extension)  
Type: integer  
Extension Command: C:\data.bat  
Access Rights: Read Only
  - d. Click *Add*.  
The MIB Extension is added to the Policy.
  - e. Click *Save Policy*.  
The policy is saved.
4. Create a threshold monitor to check the value of the new monitor.
  - a. Click *Monitors* and then *Thresholds*.  
The Threshold Monitor Details Edit pane appears.
  - b. Add the following data into the fields:  
Index: (automatically added)  
Platform: OS Independent  
Object Class: extensionGroup [Extended mib from adding new scalar variables]  
Object Attribute: 1  
Object Instance Name: MyData  
Interval: 60  
Severity: Major Alarm  
Operator: greater than or equal to  
Value: 50  
Scale: 1  
Sample Type: absolute value

- c. Click *Save*.

The policy is saved. A 'major' alarm with threshold '50' is added. This threshold will be breached immediately as the script created previously always returns the value '99'.

- d. Click *Action* and then *Apply*, to apply the policy to a computer.

The Selected Machines pane appears.

- e. Verify that the selected machines are correct and click *Apply*.

The policy with the MIB extension is applied to the selected computers.

Click *Return to Policy*.

The policy details pane appears.

Once the agent is configured, you can view the state of this threshold monitor from the Resources tab. You can see that the "major" threshold has been breached.

## Example: How to Monitor a Specific Windows Performance Registry Metric

This step-by-step example describes how to monitor a user-specific metric.

### How to monitor user-specific metrics (MIB extensions)

1. Create a MIB extension for a Windows Performance Registry metric.
  - a. From the user interface, click *Policy*, open *Configuration* in the navigation pane, expand the Policy tree, and open a SystemEDGE policy.

The policy details appear in the right pane.
  - b. Click the *MIB Extensions* tab.

The MIB Extensions pane opens.
  - c. Click *Windows Performance*.

The Windows Performance Defined Extensions pane appears.
  - d. Add the following data into the fields:

Index: 1 (If the extension is the first one)

Type: integer

Object: System

Counter: Processes (Provides the total number of running processes).

The System metrics have no 'instance' so this field is left blank.

- e. Click *Add*.  
The MIB Extension is added to the Policy.
  - f. Click *Save Policy*.  
The policy is saved.
2. Create a threshold monitor to check the value of the new monitor.
  - a. Click *Monitors* and then *Thresholds*.  
The Threshold Monitor Details Edit pane appears.
  - b. Add the following data into the fields:  
Index: (automatically added)  
Platform: OS Independent  
Object Class: ntRegPerf (NT Registry and Performance Group)  
Object Attribute: 1  
Object Instance Name: NumProcesses  
Interval: 60  
Severity: Warning Alarm  
Operator: greater than or equal to  
Value: 50  
Scale: 1  
Sample Type: absolute value
  - c. Click *Save*.  
The policy is saved. A 'warning' alarm with threshold '50' is added.
  - d. Click *Action* and then *Apply*, to apply the policy to a computer.  
The Selected Machines pane appears.
  - e. Verify that the selected machines are correct and click *Apply*.  
The policy with the MIB extension is applied to the selected computers.  
Click *Return to Policy*.  
The policy details pane appears.

Once the agent is configured, you can view the state of this threshold monitor from the Resources tab.



# Chapter 7: Provisioning Resources

---

This section contains the following topics:

[Imaging Services](#) (see page 177)

[VMware vCenter Provisioning](#) (see page 178)

[Add a Virtual Machine \(Hyper-V Server\)](#) (see page 185)

[Add a Solaris Zone](#) (see page 186)

[Add a Logical Partition for an IBM AIX Computer](#) (see page 187)

[Additional Provisioning Information](#) (see page 189)

[Remote Deployment](#) (see page 190)

[Infrastructure Deployment Process](#) (see page 215)

[Deployment to Windows Vista, Windows 2008 and Windows XP Computers Running Firewall Software](#) (see page 223)

## Imaging Services

CA Virtual Assurance can provision new virtual machines, logical partitions, or Solaris Zones, and reimage existing resources where appropriate. Provisioning functions let you clone, migrate, configure and change the properties of VMs, and create and manage LPARs and Solaris Zones.

The imaging service component uses and integrates with the following technologies to perform provisioning operations:

- Solaris Zones servers for zones provisioning
- VMware vCenter Server integration for VM provisioning. VM provisioning supports Standard Switches and Distributed Virtual Switches.
- Hyper-V integration for VM provisioning based on local templates on Hyper-V servers out of the box and optional integration with Microsoft System Center Virtual Machine Manager (SCVMM) for reuse of existing SCVMM image libraries.

Imaging services generate events for the following actions:

- An imaging job has been submitted to the imaging server.
- Imaging job status changes occurred.
- The target computer is discovered after the imaging job succeeds.

## VMware vCenter Provisioning

The imaging service communicates with the VC PMM, which integrates with VMware vCenter Server to provide VM provisioning for CA Virtual Assurance. This integration permits the imaging service to perform operations on VMs such as powering on, powering off, setting shares for CPU or memory, migrating a VM, or creating and provisioning a new VM.

This documentation assumes that you are familiar with VMware vCenter Server. All requirements and restrictions imposed by VMware on vCenter Server are valid for CA Virtual Assurance.

**Note:** For documentation about VMware vCenter Server and vSphere, see the website <http://pubs.vmware.com>.

## VMware vCenter and vSphere User Permissions

All VMware vCenter operations that CA Virtual Assurance supports require that you configure user permissions in the VMware product. The minimum requirements are described here. Some users may require additional permissions depending on their role. These permissions apply to VMware vCenter versions 2.5, 4.0, and 4.1.

Permissions are set in the Edit Role dialog under All Privileges, Virtual Machine.

Users are located on the Administration, Configuration, Provisioning: vCenter Server page. After configuring permissions, perform the following actions to ensure that VM provisioning does not fail:

- Add the users to the hosts and clusters level in VMware. See <http://kb.vmware.com/kb/1008192> for details.

**Note:** For more information about setting permissions, see the VMware documentation.

### Power Operation

To let users perform power operations (Adjust VC Power Action), select the following options under Interaction:

- Power on
- Power off
- Suspend
- Reset

### Resource Adjustment

Resource adjustment permissions let users perform the following operations:

- Set CPU Limit
- Set CPU Reservation
- Set Memory Limit
- Set Memory Reservation

To configure resource adjustment permissions, select Change Resource under Configuration.

### Share Adjustment

Share adjustment permissions let users perform the following operations:

- Set Memory
- Add Memory
- Subtract Memory
- Set CPU
- Add CPU
- Subtract CPU

To configure share adjustment permissions, select the following settings under Configuration:

- Change CPU Count
- Memory
- Settings

### Delete VM

Delete VM permissions let users delete virtual machines (Delete VC Virtual Machine).

To configure delete VM permissions, select Remove under Inventory.

### VM and Template Conversion

VM and template conversion permissions let users convert templates to VMs and VMs to templates.

To configure VM and template conversion permissions, select the following settings:

#### **under Provisioning**

- Mark As Template
- Mark as Virtual Machine

#### **under Resource**

- Assign Virtual Machine to Resource Pool

### Cloning

Cloning permissions let users perform cloning operations including VM to template cloning or from VM to VM.

To configure vCenter cloning permissions, select the following settings:

#### **under Inventory**

- Create
- Remove
- Move

#### **under Interaction**

- Power On
- Power Off
- Suspend
- Reset
- Device connection
- Tools Install

**under Configuration**

- Add Existing Disk
- Add New Disk
- Change CPU Count
- Memory
- Modify Device Settings
- Settings
- Change Resource

**under Provisioning**

- Customize
- Clone
- Deploy Template
- Mark as Template
- Mark as Virtual Machine
- Read Customization Specifications
- Modify Customization Specification

**under Resource**

Assign Virtual Machine to Resource Pool

**VMotion**

VMotion permissions let users migrate VMs from one physical server to another.

To configure VMotion permissions, select the following settings under Resource:

- Migrate
- Relocate
- Query VMotion

### VM Snapshot

VM snapshot permissions let users take snapshots.

To configure VM snapshot permissions, select the following settings under State:

- Create Snapshot
- Revert To Snapshot
- Remove Snapshot
- Rename Snapshot

### More Information

[User-scoped Authentication for vCenter Server](#) (see page 37)

## Add a Virtual Machine (vCenter Server)

You can use one of two methods to add a VM:

- Clone a predefined template
- Clone an existing VM and a customization specification. The customization specification defines the characteristics of the Guest OS.

VM provisioning supports Standard Switches and Distributed Virtual Switches. When provisioning a VM that is attached to a Distributed Virtual Switch, you can specify the appropriate discovered dvPort Group in the user interface. dvPort Groups define how a connection is made to the network through the Distributed Virtual Switch.

### To add a VM

1. Right-click VMware vCenter Server in the Explore pane and select Provisioning, Provision VMware VM.

VMware vCenter Provisioning dialog appears.

2. Select options from the drop-down lists to specify the settings.

**Note:** The virtual machines listed for cloning are limited to virtual machines that are monitored by CA Virtual Assurance. Access to VMs is restricted to ensure security. If you want to clone a system that is unavailable, discover that system as you would any other system to make it available in the drop-down list.

3. Enter your user name, password, and the host name to use. Otherwise the name indicated in the specification is used by default.

**Note:** The user name and password for Windows and Linux must match those defined in the customization specification file.

4. Select *one* of the following options and click Next:

- VC Virtual Machine to use an existing VM
- VC Template to use a template to create a new VM
- VC Specification to select a customization specification from the available list

The Virtual Machine Memory page appears.

5. (Optional) Adjust the memory for the VM and click Next.

#### **Memory**

Populates the field with the memory value defined in the VM template or VM.

**Default:** 4 MB minimum and 16 GB maximum

**Note:** Configure these values in the caimgconf.cfg file.

The Virtual Machine CPU page appears.

6. (Optional) Adjust the CPU for the VM and click Next.

#### **Virtual Processors**

Populates the field with the number of virtual processors defined in the VM template or VM.

**Default:** 1 CPU minimum and 4 CPU maximum

**Note:** Configure these values in the caimgconf.cfg file.

The Disk page appears with the fields populated with the default values from the selected VM or template that you selected.

7. (Optional) Set the drive size and click Add Drive to add drives, configure which data store to associate the hard disk with, and which SCSI controller to use from the drop-down lists and click Next.

#### **Datastore**

Identifies the data store name of the VMware ESX host where the VM will be created.

#### **Drive size**

Lets you specify a drive size and add more hard disks to the VM.

**Limits:** The minimum drive size is 1 MB, but cannot exceed the drive size for the data store you selected.

#### **SCSI controller**

Specifies which SCSI controller to use as the virtual adapter.

The Network page appears and the table is populated with the default values from the selected template.

8. (Optional) Click inside the cells in the Network Management table to activate drop-down lists, change any settings desired.

If your custom specification specifies the use of DHCP, you will only be able to edit the network connection cell in the table. Network connections now support both networks for standard and distributed virtual switches. You can distinguish the names of Standard Switches and Distributed Virtual Switches based on the following naming convention:

- For Standard Switches, the name is the network name.
- For Distributed Virtual Switches, the name is a concatenation of the dvPort group name followed by the Distributed Virtual Switch name enclosed in parentheses: dvPortGroupName (dvSwitchName)

If your custom specification specifies the use of a static IP address, you will be able to edit all cells except the NIC cell. CA Virtual Assurance does not support the custom specification network setting "Prompt User." Custom Specifications that use this setting will be filtered out and unavailable.

Click Next.

9. Click Add Computer.

A confirmation message appears at the top of the pane.

**Note:** Imaging takes time, so you should expect a delay during operating system installation. For more efficient discovery, you can adjust the discovery retry time or the interval in the caimgconf.cfg file located at: *install\_path\CA\productname\conf*.

10. Click Refresh to see the new VM in the left pane.

Your data center has a new cloned VM. You can view the events of the imaging process in the dashboard and you can generate an imaging job report.



## Add a Virtual Machine (Hyper-V Server)

You can create a VM to your data center. You must use a predefined template to create a VM.

**Note:** The value for Hyper-V "Total Storage" includes the total space required to create the VM from the template. This value is a combination of several factors that include all virtual disks, RAM size for the VM, snapshots, and a buffer. Use this information as guidance for the maximum amount of storage required to create a VM based on the template selected.

### To create a VM

1. Select Resources, Explore.  
The Explore pane appears.
2. Right-click a Hyper-V resource, and select Provisioning, Provision Hyper-V VM.
3. Specify the following details and click Next.
  - SCVMM server and the Hyper-V server.
  - Template name that you want to use to create a VM.
  - Destination path where you want to create the VM.
  - Name of the VM that you want to create.
  - Specify whether to start the VM after it is created.

The Virtual Machine Memory page appears.

4. Specify the VM memory details and click Next.  
The Guest OS Customization page appears.
5. Specify the guest operating system details and click Next.  
The Network Management page appears.
6. Specify the network details of the VM and click Next.

**Note:** If your custom specification specifies the use of DHCP, you will only be able to edit the network connection cell in the table. If your custom specification specifies the use of a static IP address, you will be able to edit all cells except the NIC cell. CA Virtual Assurance does not support the custom specification network setting "Prompt User." Custom Specifications that use this setting will be filtered out and unavailable.

7. Click Add Computer.

A confirmation message appears at the top of the pane.

**Note:** Imaging takes time, so you should expect a delay during operating system installation. For more efficient discovery, you can adjust the discovery retry time or the interval in the `caimgconf.cfg` file located at: `install_path\CA\productname\conf..`

## Add a Solaris Zone

Solaris Zones servers use Zones to run applications in isolated environments to make it appear as if they are running on physically separate systems. Each Zone on a server takes its resources from a resource pool and includes virtual network interfaces, file systems, memory, and other dedicated units. When you create a Zone, you must supply all of this information. The Zone installs automatically after creation.

### To add a Solaris Zone

1. Select the Resources tab, right-click the Zone Host in the Explore pane, and select Provisioning, Provision Zone.

The Solaris Zone Provisioning wizard appears.

2. Complete the following fields on the Zone Identity and Type page and click Next:

#### Host

Defines the host on which to create the Zone.

#### Name

Defines the name of the Zone.

#### Description

(Optional) Defines a description of the Zone.

#### Type

Defines whether the Zone is Native or Branded. A Branded Zone is based on an existing Zone template.

#### Template Name

(Optional) Defines the template from which to create the Zone when you set Type to Branded.

#### Install Archive Path

Defines the directory path of the installation archive on the Zone. This field is only required if you set Type to Branded.

The CPU, Memory, and Additional page appears.

3. Complete the following fields and click Finish:

#### Type

Defines the scheduler type. Select FSS to use the Fair Share Scheduling class to control CPU allocation based on the number of CPU shares assigned to tasks.

#### Capacity

Defines the amount of physical memory capacity to allocate to the Zone, in megabytes.

**Swap Memory**

Defines the amount of swap memory to allocate to the Zone, in megabytes. The swap memory must be at least 50 MB.

**Lock Memory**

Defines the amount of lock memory to allocate to the Zone, in megabytes. The lock memory must be less than the physical memory.

**Zone Path**

Defines the root directory path of the Zone.

**NIC Type**

(Optional) Defines the NIC type. Select a type from the drop-down list. If you do not select a NIC, the Zone is not assigned a NIC card or IP address.

**IP Address**

(Optional) Defines the IP address of the Zone.

**Resource Pool**

Defines the resource pool to use with the Zone. Select a pool from the drop-down list. If you want to use a new resource pool with the Zone, create the pool first. If you do not select a pool, the default is used.

**Auto Reboot**

Defines whether to reboot the Zone automatically when the global Zone is rebooted.

## Add a Logical Partition for an IBM AIX Computer

You can use the Provisioning wizard to create logical partitions for an IBM AIX system.

**To add a logical partition for an IBM AIX computer**

1. Click Resources.
2. Right-click IBM PowerVM Server in the Explore pane, and select Provisioning, Provision LPAR.

The Provisioning wizard appears with the Partition and Memory page.

3. Select the HMC/IVM server and managed system name. Specify the partition name and, if using an HMC server, the profile name. Specify the minimum, desired, and maximum memory for the partition. Click Next.

The Processor page appears.

4. Specify the minimum, desired, and maximum number of processors, and indicate whether to allocate a shared or dedicated processor. Advanced settings are available for shared modes and virtual processors. Click Next.

The I/O Components page appears.

5. Select the I/O devices to associate with the partition, and click Next.

**Note:** For each I/O device, you can specify that the I/O device is required or optional for logical partition activation. If the I/O device is required, the partition cannot be activated if the I/O device is unavailable or in use by another logical partition. If the I/O device is optional, and if the desired I/O device is available when the partition is activated, the managed system commits the I/O device to the partition. If the optional I/O device is not available, the managed system skips the I/O device.

The I/O Pools page appears.

6. (Optional) To create a new I/O pool, click the + (Add) on the I/O Pools table, enter a numerical value, and click Save.

**Note:** When you add an I/O device to a partition, and the I/O device belongs to an I/O pool. When this partition is activated, the managed system automatically adds the I/O pools defined for the partition to the logical partition.

7. Click Next.

If an HMC server was selected, the Virtual Serial page appears.

8. (Optional) Specify the maximum virtual adapters for the partition. To create a new virtual serial adapters, click + (Add) and specify the Adapter ID, Remote Partition, and Remote Slot Number. You can require that the virtual adapter must be allocated and the managed system must have enough memory to run the required virtual adapters for the partition profile, or the logical partition does not activate.

9. Click Next.

The Virtual Ethernet page appears.

10. Specify the maximum virtual ethernet adapters for the partition. (Optional) You can add new virtual ethernet adapters by clicking + (Add) and selecting an Adapter ID, Virtual LAN ID, Access External Network, Trunk Priority, IEEE 802.1 Q Compatibility, additional Virtual LAN IDs, and whether the Ethernet adapter is required.

11. Click Next.

The Virtual SCSI page appears.

12. Specify the virtual SCSI devices for the partition. (Optional) To add a new virtual SCSI adapter, click + (Add) on the Virtual SCSI Adapters table.

Select an Adapter ID, specify whether the SCSI adapter is Required, and pick a Device name from the SCSI Devices table. If the desired device is on the SCSI Devices list, click OK, click Next in the Virtual SCSI panel, and skip to the last step. To add a new SCSI backing device, click + (New Backing Device) on the SCSI devices table.

**Note:** If the selected device has a slot number, that is the slot number of the virtual SCSI server adapter defined to the Virtual I/O server partition.

If the selected device doesn't have a slot number, that means it isn't associated with a virtual SCSI server adapter yet. When the job to create the partition takes place, the virtual SCSI server adapter will be created and assigned to the device. Click Next.

The Summary page appears.

13. Verify the Summary and click Add Computer.

The logical partition is created.

## Additional Provisioning Information

You can find more information about application and operating system deployment on the CA Support Online website. The following information provides you with advanced CA Desktop and Server Management topics and system requirement information.

- *CA Desktop and Server Management Advanced Topics* guide:

[https://support.ca.com/phpdocs/0/common/impcd/r11/troubleshooting/doc/DSM\\_Adv\\_topics\\_r11.pdf](https://support.ca.com/phpdocs/0/common/impcd/r11/troubleshooting/doc/DSM_Adv_topics_r11.pdf)

- *CA Desktop and Server Management Green Books*:

[https://support.ca.com/phpdocs/7/common/greenbooks/CA\\_Unicenter\\_Desktop\\_Server\\_Management\\_Green\\_Book\\_ENU.pdf](https://support.ca.com/phpdocs/7/common/greenbooks/CA_Unicenter_Desktop_Server_Management_Green_Book_ENU.pdf)

- The following information identifies the operating systems and versions that are supported for provisioning:

**Windows and Linux:**

CA Software Delivery compatibility matrix:

[https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/234/234\\_comp\\_matrix.html](https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/234/234_comp_matrix.html)

CA IT Client Manager compatibility matrix:

[https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/8176/8176\\_compmatrix.html](https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/8176/8176_compmatrix.html)

**Note:** For specific information about supported operating environments and hardware requirements, see the *Release Notes*.

## Remote Deployment

CA Virtual Assurance provides a comprehensive solution for remotely deploying SystemEDGE and other agents to all managed systems. You can create deployment templates based on provided packages that contain customized installation parameters and simultaneously deploy these templates to numerous managed systems. This automated deployment solution provides one location from which to deploy and configure the agents throughout your enterprise.

Remote deployment provides the following features:

**Deployment configuration**

Allows creation, editing, and deletion of configurations which define how software packages are installed on target systems. These are referred to as Package Wrappers.

**Deployment job management**

Allows creation, start, and cancelation of deployment jobs, allowing the concurrent deployment of packages to multiple targets using multiple distribution servers.

**Deployment job reporting**


Provides the ability to query the status of deployment jobs.

**Deployment events**

Provide a source of deployment-related events that tracks the state of the managed nodes.

## Remote Deployment Architecture

The overall architecture of the deployment solution is driven by the Domain Server and Distribution Server components. The following diagram represents an overview of the deployment-related components:



## Deployment Components

This section lists and briefly describes the deployment key components:

### Domain Server

The domain server is the repository for all configuration and control data. The server is responsible for managing configuration and software package data required for deployment operations and it manages all configuration operations. Detailed event data is passed between the domain and distribution servers during the deployment process. A single Domain Server is responsible for maintaining the status of all Distribution Server jobs.

### Distribution Server

The distribution server controls the Infrastructure Deployment Manager (IDManager) server located on the same machine. The architecture allows for multiple distribution servers offering deployment services.

### Infrastructure Deployment

Infrastructure Deployment initiates and manages deployment jobs. During the deployment process, the Infrastructure Deployment Manager (IDManager) provides access to remote systems and the Infrastructure Deployment Primer (ID Primer) provides the mechanism to remotely install agent software packages. The IDPrimer is used to transfer the deployment package data to the target computer and run the installation. All subsequent deployments to the same target computer can use the existing IDPrimer installation. The IDManager controls all the deployment operations and handles job status.

## Using Multiple Distribution Servers

Although the remote deployment solution lets you have a single central server (manager machine) to manage all your deployments, it is recommended that you install a remote distribution server that points to the central domain server if you have any of the following requirements:

- You have two or more geographically remote locations where you need the agent software deployed to but need these managed centrally using a single manager machine.

In this case we recommend that each location has at least one distribution server with the distribution server connects to the central domain server.

- You have a single location but have several hundreds of machines that you have the need to deploy to.

In this case you can choose to install many distribution servers split logically across subnets and these distribution servers will connect to the central domain server.



## Change the Domain Server a Distribution Server Connects To

In the situation where the network address of the domain server machine changes after the original installation, it is necessary to reconfigure the distribution server to connect to a new network address.

Prior to making the configuration change shown below it is important to verify that the new network address is connectable from the distribution server. If the distribution cannot make a connection to the domain server using the new address then deployment functionality will not work correctly.

### To change the domain server a distribution server connects to

1. From the Start menu, open Administrative Tools, Services.

The Services user interface appears, listing the installed services.

2. Right-click CA SM Distribution Server and select Properties.

The Properties dialog appears.

3. Click Stop to stop the service.

4. Enter the following parameter into the Start parameters field:

`-m domainserver`

The *domainserver* parameter specifies the IP address or DNS name of the domain server.

5. Click Start

The distribution server will now attempt to connect to the domainserver address entered.

## Scalability

The deployment system provides a degree of scalability using multiple distribution servers as a scalability layer. Each distribution server communicates with one IDManager instance. The IDManager can manage multiple component deployments to multiple target computers. CA Virtual Assurance supports many simultaneous deployments because of this federated model.

### More Information

[Remote Deployment and Policy Configuration Recommendations](#) (see page 281)

[Scalability Use Cases](#) (see page 283)

## Audit Trail

Jobs and tasks are the two fundamental concepts of the deployment system. A deployment job specifies one or more configurations to be installed on one or more target systems. Each individual installation of a software package on a target system is represented by a task. Deployment job reporting provides the ability to query the status of deployment jobs.

You can create, control and inquire the state of deployment jobs. When the job is started, its individual deployment tasks are delegated to available distribution servers which perform the actual deployment. You can track the progress of the job as it occurs, to verify that the deployment is going well and to quickly identify and correct any problems.

Remote deployment is able to provide the following information:

- Which deployment jobs are currently:
  - Inactive (not yet started)
  - Active
  - Completed, and of those, which were:
    - Successful
    - Partially successful
    - Unsuccessful
- Which deployment jobs are:
  - Associated with a specific target machine
  - Associated with a specific package/package group
- What packages have been deployed to a specific target machine
- Which user created/started deployment of a specific package
- Which machines are targeted in a specific deployment job
- Which machines are targeted by active deployment jobs

**Note:** Remote Deployment supports deploying software to UNIX/Linux systems with /tmp file system mounted with noexec flag.

The Jobs pane in the UI displays a table of all created deployment jobs that lists the job ID and description, the packages included, job status, and so on. From this table, you can drill down to view more details about a specific job, including why a job failed. Click the Jobs folder in the UI to access this view. For more information, see the *Online Help*.

## Deployment Restrictions

Consider the following before performing a deployment:

- If you want to install an agent on the CA Virtual Assurance Manager system, you must perform a manual standalone agent installation. Deployment of agents on the CA Virtual Assurance Manager system is not supported.
- The deployment process is dependent on the availability of existing host operating system services in order to gain remote access to target systems. When these services are not available on the target nodes it will be necessary to install the IDPrimer client package and a corresponding key on a target system. For more information, see the section *Manual Installation of the Remote Deployment Primer Software*.
- Deployment is supported to most, but not all, supported agent platforms. For more information about deployment support, see the *Release Notes*.

## Deployment Credential Restrictions

The UI limits the entries for both username and password fields to 64 characters.

## Remote Deployment to UNIX/Linux Using Non Privileged User Account

If you want to use a nonprivileged user account, consider the following requirements about the sudo configuration:

- Sudo must not enforce that the executed program has a valid pseudo terminal attached to it. To disable such validation for a particular user (if it is globally enabled), add the line “Defaults:\$username !requiretty” to the /etc/sudoers file. Replace \$username by the actual username used for Remote Deployment.  
  
The standard way to edit the file is using the visudo command. The visudo command invokes \$EDITOR. When editing is finished, it verifies the syntax of the file. If the result is not valid, visudo blocks saving the file.
- Sudo must not ask the user for a password before running the elevated program. To achieve this behavior, the NOPASSWD: keyword must be present on the line giving privileges to the user.
- Sudo must be allowed to run the necessary commands or all. Configuration entries (lines in /etc/sudoers) satisfying the previous requirements are, for example:

```
$username ALL=(ALL) NOPASSWD: ALL
```

or

```
$username ALL = NOPASSWD: /usr/bin/id,/bin/sh /tmp/idprimer/PifInst -r /tmp/idprimer/install.rsp
```

**Note:** Replace \$username by the actual username used for Remote Deployment.

On Solaris, consider the following requirements for pfexec:

- Any local user can be given profile “Primary Administrator” with the following command

```
usermod -P “Primary Administrator” {user}
```

- Any nonlocal user can be given profile “Primary Administrator” by manually adding an entry in file /etc/user\_attr:

```
user:::type=normal;profiles=Primary Administrator
```

## Deployment Sizing Key Factors

A number of key factors having a considerable impact on the infrastructure sizing, and system performance, including the following:

- Size of software packages to deliver.
- Number of software packages to deliver.
- Frequency of software package delivery.
- Network latency between deployment components and target computers.
- Network bandwidth management.

The initial deployment to a target installs IDPrimer, a small installation agent. Once IDPrimer is installed, subsequent deployments to the same target should take less time.

Under ideal conditions, a single distribution server can install about 50-100 SystemEDGE packages per hour. However, if many tasks fail due to any of the deployment prerequisites not being met, the throughput of a distribution server will be reduced.

Recommendations:

- Verify that target servers typically meet the requirements for deploying software remotely.
- A single distribution server can handle 500 target servers comfortably.
- Install additional distribution servers local to the target location.
- Deploy using distribution servers that are local to targets if possible.
- Schedule deployments to start during periods of low network traffic if possible.

## Deploying/Installing SystemEDGE Agents Using Custom Ports

Deployment of SystemEDGE agents to a non-standard port requires a number of settings to be configured. To ensure the manager can discover and manage the system once it has been deployed, perform the following actions:

1. Update the package wrapper:
  - If you are using the remote deployment solution, you must first configure the package wrapper to specify the port to be used. Navigate to Provisioning, Deployment in the user interface and change the Port field. The Write Community string can also be updated here.
2. Update SNMP Community strings in CA Virtual Assurance:
  - For the Manager to successfully monitor and manage a machine using a non-standard port, it must know the appropriate Port / Write Community string combination to use for monitoring and management. This can be done either by creating a global SNMP entry that can be used to monitor and manage multiple systems, or by creating a server-specific SNMP entry:
    - To update global SNMP settings: Navigate to Administration, SNMP in the user interface and add a new entry with the appropriate SNMP community string / port combination.
    - To update server-specific SNMP settings: Navigate to Policy, Explore, *Machine\_Name*, Metrics, SNMP Settings and add a new entry for the required port / write community string.

Once these settings have been updated, the agent can be deployed / installed in the usual way. The SystemEDGE Platform Management Module will then use the custom port / write community string combination to discover, monitor and manage the server.

### Reconfigure Ports for SystemEDGE Agents

You can reconfigure the port for SystemEDGE agent by reinstalling the agent. After reinstalling the agent, the settings remain unchanged with a provision to edit the details of the port to be reconfigured.

## Deployment Package Library

The package library contains a configurable set of installable software packages where you can control which products, versions and platforms are available for deployment. You can control the way these products are installed by creating standard package configurations that define the parameters required for an unattended installation of a configured software package.

Each package must have an associated package configuration file. The configuration file provides information describing both the package details and how the package installation can be configured. For more information, see the [Deployment Package Configuration File](#) (see page 201) section.

The package library is located in the following directory:

%AllUsersProfile%\Application Data\CA\SM\domainserver\Deployment\Packages

The directory tree layout is defined by the requirements of the IDManager component. The package library itself consists of a top-level packages directory which contains two sub-directories, Public and Private. The Public directory contains all the deployable software packages.

```
..
├── SM
│   ├── CA_LiteAgent
│   ├── CA_ProcProbe_Utility
│   ├── CA_SystemEDGE_AdvancedEncryption
│   ├── CA_SystemEDGE_Core
│   │   ├── 5.6.0
│   │   │   ├── ENU
│   │   │   │   ├── AIX_aix
│   │   │   │   ├── HPUX_hp
│   │   │   │   ├── HPUX_ia64
│   │   │   │   ├── Linux_ia64
│   │   │   │   ├── Linux_x86
│   │   │   │   ├── Solaris_sparc
│   │   │   │   ├── Solaris_x86
│   │   │   │   ├── Windows_ia64
│   │   │   │   ├── Windows_x64
│   │   │   │   └── Windows_x86
│   │   ├── CA_SystemEDGE_ESAD
│   │   ├── CA_SystemEDGE_HyperV
│   │   ├── CA_SystemEDGE_LPAR
│   │   ├── CA_SystemEDGE_MSCS
│   │   ├── CA_SystemEDGE_RM
│   │   ├── CA_SystemEDGE_SRM
│   │   ├── CA_SystemEDGE_UCS
│   │   ├── CA_SystemEDGE_Zone
│   │   └── CA_VMVCAIM
```

The top-level Public directory has five sub-directories:

**Component Name**

Must be the IDManager instance name, which for CA Virtual Assurance is SM.

**Software Package**

Contains all versions, localizations and architectures of a single deployable package, for example CA\_SystemEDGE\_Core, CA\_SystemEDGE\_SRM

**Version**

The version of the software packages contained within, for example 5.0.1

**Language**

The installation package language, for example. ENU

**Architecture**

The architecture specific installation materials, for example Windows\_ia64, Solaris\_x86. The architecture directory name must be one of the platforms supported by IDManager.

When run within the distribution server machine, the IDManager component uses directories under the distribution server. This contains a temporary cache of encrypted packages for its internal use. These packages should be removed upon job completion.

The private IDPrimer installation materials are contained in a different directory. By default these are stored under the installation directory of the IDManager component itself, in the following directory:

<CA Shared Components>/IDMgrApi/packages/private/idprimer

This directory contains the IDPrimer installation materials for all platforms supported by the infrastructure deployment component.

## Deployment Package Library Maintenance

You may want to perform maintenance of the deployment packages, for example:

- Add new deployment packages
- Copy existing packages to a new package to create a patched version of the same base package (master image patch).
- Delete packages which are no longer required.
- Copy CA Virtual Assurance compatible packages from installation media, or another specified directory into the package library.
- Copy an entire registered package to a new name and generate a new UUID for that package.
- Perform consistency checking of the package library and the associated deployment configurations data.

## Package Filter

If upgrades have been applied to the server, Remote Deployment can show an increasing number of package versions. The default behavior of this release is to show only the latest packages available. As a consequence, the data in the Packages - Details tab is also filtered to show only the latest versions of each package. Selecting a wrapper from this panel, expands the tree at the selected wrapper position.

If you want to see all packages, you can override the default filtering behavior by the check box “Display Latest Package Versions Only” in the Package Information Panel. When enabled (default), it filters out any older package versions from the left tree and the Package Details Tab.

### To change the filtering behavior

1. Select/Unselect the check box for “Display Latest Package Versions Only” in the Package Information Panel.
2. Refresh the Deployment view in the user interface.

The latest package versions appear.

**Note:** All other locations within the UI where package versions are displayed are not affected.



## Deployment Package Configuration File

In addition to the software package installation materials, each deployable package must be referenced by an additional package configuration file, `pkginfo_PLATFORM.xml`. The package configuration file describes what packages to install and how the installation is configured. The configuration files provide the following:

- A localizable description of the installation package
- Text for any End User License Agreement (EULA) that must be accepted in the UI
- A mechanism by which package dependencies may be encoded in a machine-readable format
- Documenting the publicly accessible installation parameter types
- Additional context to the parameter types so a level of validation may be performed within the UI
- A mapping between parameter names and the tokens used to represent those parameters in the packages installation program, in a platform independent manner
- Specifies how the installation materials should be executed on a target machine
- Mapping between the installer exit codes and those understood by the deployment system

Localized elements of the `pkginfo.xml` file may optionally be provided using either side by side locale-specific files or embedded within a single file. At load time any file whose name matches filename `pkginfo_PLATFORM.xml` may be loaded to obtain localized message data.

The deployment system requires the package configuration files to be located parallel to the platform-specific sub-directory in the packaging tree. A single package description file can describe the installation for multiple platforms. See the following directory:

```
%AllUsersProfile%\Application  
Data\CA\SM\domainserver\Deployment\Packages\SM\CA_SystemEDGE_Core\5.6.0\EN  
U
```

```
pkginfo_AIX.xml  
pkginfo_HPUX.xml  
pkginfo_Linux.xml  
pkginfo_Solaris_sparc.xml  
pkginfo_Solaris_x86.xml  
pkginfo_Windows.xml
```

## Deployment Packages

Deployment packages provide the materials necessary to deploy monitoring software to systems across your enterprise. The following deployment packages are available:

### **CCA Agent**

Provides the CCA Agent.

### **CA LiteAgent**

Provides the CA Performance Lite Agent.

### **SystemEDGE**

Provides the core SystemEDGE agent.

### **SystemEDGE Advanced Encryption**

Provides a FIPS 140 compliant encryption package for SystemEDGE.

### **SystemEDGE AIX LPAR**

Provides the LPAR AIM.

### **SystemEDGE ADES**

Provides the Active Directory and Exchange Server (ADES) AIM.

### **SystemEDGE Hyper-V**

Provides the Hyper-V AIM.

### **SystemEDGE MSCS**

Provides the Microsoft Cluster Support (MSCS) AIM.

### **SystemEDGE RM**

Provides the Remote Monitoring AIM.

### **SystemEDGE Solaris Zone**

Provides the Solaris Zone AIM.

### **SystemEDGE SRM**

Provides the SRM AIM.

### **SystemEDGE UCS**

Provides the Cisco UCS AIM.

### **SystemEDGE VC**

Provides the VMware vCenter AIM.

The Solaris Zone, Remote Monitoring, vCenter, and SRM AIMs all depend on both SystemEDGE 5.6 and Advanced Encryption packages. Therefore, if you want to deploy any of these packages, SystemEDGE and Advanced Encryption must either already exist on the system, or you must include these in the deployment job.

Deployment packages are broken into platform-specific variants, package wrappers are available for all platforms that support deployment.

## Default Package Wrappers

Default package wrappers are provided out of the box for the software packages that can be deployed using Remote Deployment. These package wrappers contain installer parameters with default values where possible for the chosen software package. Where a package needs a mandatory parameter which is not safe to default for security reasons (for example, SNMP write community string, user name, password), or cannot be defaulted because it is user environment specific (vCenter host name), the default package wrappers will not contain a pre-filled default. This means the default package wrappers containing mandatory parameters need to be edited first with the mandatory parameter values filled in and saved before they can be used in a deployment job.

The following package wrappers contain mandatory parameters that need to be filled in prior to deployment:

### **SystemEDGE RM**

Specify a user name and a password to deploy this package.

### **SystemEDGE MSCS**

Specify an MSCS hostname, MSCS user name, and a password to deploy this package.

### **SystemEDGE VC**

Specify a vCenter host name, user name and password to deploy this package.

### **SystemEDGE LPAR AIM**

Specify a host name, user name, and password to deploy this package.

### **SystemEDGE Solaris Zone**

Specify a host name, user name, and password to deploy this package.

### **SystemEDGE UCS**

Specify a host name, user name, password, protocol, and port to deploy this package.

### **SystemEDGE Hyper-V**

Specify a host name, user name, and password to deploy this package.

The mandatory values for the default package wrapper will need editing only once before the first deployment. These need not be edited again unless there is a need to modify the installer parameter values for a package. If you proceed to deploy a package with a default package wrapper that does not have the mandatory parameters filled in, the deployment wizard will indicate that the package wrapper is not in a deployable state.

For more information, see the *Online Help*.

## **Agent Configuration Without Write Community**

Although It is not mandatory to provide a write community for SystemEDGE package wrappers, consider the following:

- The SystemEDGE agent can be discovered by the SystemEDGE PMM even if the agent is configured with only SNMP read community and no write community. However, point configuration changes cannot be made to the agent without the agent configured with SNMP write community.
- Full vCenter and Remote Monitoring functionality is only supported if the agent is configured with write community. AIM configuration and administration from the CA Virtual Assurance UI is not possible without the agent configured with SNMP write community.
- An agent without write community can be configured post-installation from the CA Virtual Assurance UI using Policy Configuration. Policy Configuration also allows you to configure the agent to use SNMP v3, which is more secure than SNMP v1/2.

## Create a New Package Wrapper

Package wrappers provide platform-specific instructions for the deployment mechanism to follow when deploying a specific package. Each package contains a default package wrapper for all platforms that support remote deployment. You can create new package wrappers if certain systems require different settings than the default.

**Follow these steps:**

1. Select Resources, Deploy.

The Deployment pane displays the Packages, Templates, and Jobs folders.

2. Expand Packages.

The list of available packages appears in the Deployment pane.

3. Right-click a package name in the Deployment pane and select Create New Wrapper. You can also click + (New) on the Available Wrappers toolbar.

The New Wrapper dialog appears.

4. Enter a name and an optional description for the wrapper, specify the platform the wrapper should support, and click OK.

The wrapper is created, and details appear in the right pane.

**Note:** If you create a SystemEDGE package wrapper, consider the dependency between the Trap Port, Trap Destination, and Trap Community fields. The behavior is that either none of these fields or all must be set. In case of a partial setting, the installer displays an error message.

## Modify a Package Wrapper

Package wrappers define a set of platform-specific installation settings for a deployment package, such as installation path, port, trap communities, and so on. You can edit a user created or default package wrapper to change this set of installation settings. The available properties vary by the package type.

**To modify a package wrapper**

1. Select Resources, Deployment.

The Deployment pane displays the Packages, Templates, and Jobs folders.

2. Expand Packages, the specific package type, and the wrapper platform, and select the wrapper to modify.

The wrapper details appear in the right pane.

3. Modify the package properties as necessary and click Save. The options that appear in the Properties pane depend on the package type that you select.

### **Package types:**

#### **CA AIX LPAR AIM**

##### **Hostname**

Specifies the hostname to use for connecting to IBM LPAR server. You must specify the name of the IBM LPAR host to deploy this package.

##### **Username**

Specifies the username to use for connecting to IBM LPAR server. You must specify the name of the IBM LPAR user to deploy this package.

##### **Password**

Specifies the password to use for connecting to IBM LPAR server. You must specify a IBM LPAR password in order to deploy this package.

##### **Serial Number**

Specifies the serial number to use for connecting to IBM LPAR server. You must specify a serial number in order to deploy this package.

#### **CA LiteAgent**

##### **Shared Path**

Defines the root installation directory to use for CA Shared Components if not already set.

##### **Install Path**

Defines the root installation directory for the package.

##### **Suppress Reboot check box**

Specifies whether to suppress any automatic reboot that might otherwise occur at end of the installation.

**CA Solaris Zone AIM****Zones Host**

Specifies the hostname to use for connecting to Solaris Zone server. You must specify a Solaris Zone hostname to deploy this package.

**Username**

Specifies the username to use for connecting to Solaris Zone server. You must specify a Solaris Zone username to deploy this package.

**Password**

Specifies the password to use for connecting to Solaris Zone server. You must specify a Solaris Zone password to deploy this package.

**SystemEDGE****Install Path**

Defines the root installation directory for the package.

**Data Path**

Defines the data directory for the package.

**Shared Path**

Defines the root installation directory to use for CA Shared Components if not already set.

**Port**

Defines the SystemEDGE port number.

**Default:** 161

**Description**

Defines the SNMP system description.

**Location**

Defines the SNMP system location.

**Contact**

Defines the SNMP System contact.

**Read Community**

Defines the SNMP read-only community string.

**Default:** public

**Read-Write Community**

Defines the SNMP read-write community string

**Trap Community**

Defines the SNMP trap community string.

**Trap Destination**

Defines the SNMP trap destination host name.

**Trap Port**

Defines the SNMP trap port.

**Default:** 162

**Privilege Separation User (UNIX/Linux)**

Specifies the user name under which credentials the agent run during SNMP communication.

This entry instructs the agent to run SNMP communication under another user account. The agent also uses this user's default group as an effective group.

**Default:** The agent operates using root account.

**Start Agent check box**

Specifies whether to automatically start SystemEDGE at the end of the installation.

**Suppress Reboot check box**

Specifies whether to suppress any automatic reboot that might otherwise occur at end of the installation.

**Disable Native Agent check box**

Specifies whether to replace the native SNMP agent.

**Use native settings check box**

Specifies whether to use native SNMP agent settings (if replacing a native SNMP agent).



**Run in Managed Mode check box**

Specifies whether to run SystemEDGE in managed mode.

**Managed Policy Name drop-down list**

Specifies a list of available SystemEDGE policies.

**Note:** When you upgrade SystemEDGE from Version 4.3 or Version 4.2 patch level 3, the installer uses the following parameters only:

CASE\_PUBDATADIR  
CASE\_MANAGER\_HOSTNAME  
CASE\_MANAGER\_POLICY\_NAME  
CASE\_START\_AFTER\_INSTALL  
CASE\_LEGACY\_MODE  
CASE\_SNMP\_PORT  
CASE\_INSTALL\_DOCS  
CASE\_SNMP\_TRAP\_COMMUNITY <sup>(1)</sup>  
CASE\_SNMP\_TRAP\_DESTINATION <sup>(1)</sup>  
CASE\_SNMP\_TRAP\_PORT <sup>(1)</sup>  
CASE\_SNMP\_READ\_COMMUNITY <sup>(1)</sup>  
CASE\_SNMP\_WRITE\_COMMUNITY <sup>(1)</sup>  
CASE\_SNMP\_READ\_ALLOWED\_MANAGERS <sup>(1)</sup>  
CASE\_SNMP\_WRITE\_ALLOWED\_MANAGERS <sup>(1)</sup>

Other parameters are ignored.

(1) These parameters are special. Their settings are appended to the existing SystemEDGE 4.x settings allowing both the SystemEDGE 4.x manager and SystemEDGE 5.x manager to function.

**Note:** For more information about the above parameters, see the *Installation and Deployment* chapter in the *SystemEDGE User Guide*.

**SystemEDGE Advanced Encryption****Suppress Reboot check box**

Specifies whether to suppress automatic reboot at end of the installation.

**SystemEDGE Hyper-V****Suppress Reboot check box**

Specifies whether to suppress automatic reboot at end of the installation.

### **SystemEDGE RM**

#### **Default WMI username**

Defines the default username to use for connecting to remote machines. You must specify a username in order to deploy this package.

#### **Default WMI password**

Defines the default password to use for connecting to remote machines. You must specify a password in order to deploy this package.

#### **Suppress Reboot check box**

Specifies whether to suppress automatic reboot at end of the installation.

### **SystemEDGE SRM**

#### **Allow scripts check box**

Specifies whether to allow running scripts as tests.

#### **Allow File I/O tests check box**

Specifies whether to allow running file I/O as tests.

#### **Disable user TOS check box**

Specifies whether to disable applications from setting type of service bits in outgoing IP packets.

#### **Allow untrusted SSL check box**

Specifies whether to allow accessing SSL site with unverified certificates.

#### **Suppress Reboot check box**

Specifies whether to suppress automatic reboot at end of the installation.

**SystemEDGE UCS****UCS hostname**

Specifies the hostname to use for connecting to UCS. You must specify a UCS host name to deploy this package.

**UCS username**

Specifies the username to use for connecting to UCS. You must specify a UCS user name to deploy this package.

**UCS password**

Specifies the password to use for connecting to UCS. You must specify a UCS password to deploy this package.

**UCS Protocol**

Specifies what protocol to use, HTTP or HTTPS.

**Port**

Defines the UCS port number.

**Default:** 80 for HTTP or 443 for HTTPS.

**SystemEDGE VC****VirtualCenter hostname**

Specifies the hostname to use for connecting to VirtualCenter. You must specify a VirtualCenter host name in order to deploy this package.

**VirtualCenter username**

Specifies the username to use for connecting to VirtualCenter. You must specify a VirtualCenter user name in order to deploy this package.

**VirtualCenter password**

Specifies the password to use for connecting to VirtualCenter. You must specify a VirtualCenter password in order to deploy this package.

**VirtualCenter Port**

Defines the VirtualCenter port number.

**Default:** 443

**VirtualCenter Protocol**

Specifies what protocol to use, HTTP or HTTPS.

**Suppress Reboot check box**

Specifies whether to suppress automatic reboot at end of the installation.

## Deployment Jobs

To deploy agents to target systems, you must create a deployment job. Deployment jobs contain the details that are required for CA Virtual Assurance to deliver the deployment packages to the appropriate systems at the appropriate time. When you create a deployment job, you specify the following information:

### **Job information**

Includes the job name and whether to base the job on an existing template.

### **Deployment package**

Includes the platform, the packages to deploy, and the specific wrappers for each package.

### **Machine information**

Includes the systems to which to deploy the packages and system credentials required to establish a connection.

### **Deployment time**

Specifies when to perform the deployment: immediately, staggered over a specific time period, or scheduled for a specific time in the future.

You can also save the job as a template after you create it. A template saves the package and machine selections so that you can easily reuse them for subsequent jobs.

For more information about creating a deployment job and running a deployment, see the *Online Help*.

## Track Deployment Job Status

Once a job to deploy a set of agent packages to a set of computers has been started, you can track its progress and status. The Jobs folder displays a table of all created deployment jobs that lists the job name, packages included, job status, and so on. From this table, you can drill down to view more details about a specific job, including why a job failed.

### **To track a deployment job status**

1. Select Resources, Deploy pane.  
The Deployment pane displays the Packages, Templates, and Jobs folders.
2. Click the Jobs folder.xxxxxx  
The Job Status pane appears.
3. Click the job you want to view.  
The Job Information page appears.

4. In the Task Status pane, filter out particular job tasks by using any of the filters available. Alternatively, use the paging interface to identify the desired task.
5. Click Status Extended to view extended information about a job.

The Extended Status information dialog appears with details about the job:

**Information**

Displays general information about the job

**Message**

Displays a message about the job, for example, Package delivery failed

**Reason**

Displays the reason for the failure

**Action**

Displays what action to take to correct the problem

Common causes for deployment failure include the following:

- Lack of machine availability
- Invalid system credentials
- Inability to resolve the system host name
- Unfulfilled package dependency (for example, SystemEDGE is required for a successful SRM AIM deployment).

## View Deployment History

Deployment history information is available from the following places:

**Deployment Pane**

Displays a count of completed, active, pending, and failed deployment tasks and a summary of successful deployment. Click the top-level Deployment folder to access this view.

**Jobs pane**

Displays a table of all created deployment jobs that list the job name, packages included, and job status. From this table, you can drill down to view more details about a specific job, including why a job failed. Common causes for deployment failure include the following:

- Invalid system credentials
- Inability to resolve the system host name
- Unfulfilled package dependency (for example, SystemEDGE existence is required for a successful SRM AIM deployment)

You can resubmit a job from this pane to correct the reason for failure and redeploy.

Click the Jobs folder to access this view.

## Deployment Dashboard Views

The following views are available on the Dashboard for tracking deployment metrics:

### **Deployment Task Summary**

Displays a pie chart and a list showing the number of completed, active, pending, and failed deployment tasks.

### **Unresolved Deployment Tasks**

Displays a list of deployments that did not complete successfully. You can click the job ID to view details about why the task is unresolved.

### **Active Deployment Tasks**

Displays a list of deployment tasks that are currently active. Details include the associated deployment job, target, package, and current state. You can click the task ID for details about the current status.

### **Deployment Package Summary**

Displays a bar chart showing the number of deployment for each deployment package type.

### **Completed Deployment Jobs**

Displays a list of deployments that completed successfully. You can click the job ID to view details about the job.

# Infrastructure Deployment Process

When executing a deployment, the primary steps of the process are as follows:

1. From the administrator computer, the infrastructure deployment client component issues a request to the IDManager to install an agent on a list of one or more target computers. The deployment manager may be running on a computer that is remote from the client. The list of targets can consist of explicit machine names or IPv4 addresses.

**Note:** Only discovered resources can be deployed to.

For deployment to succeed to each target computer, it is important to verify that its name, whether entered explicitly or obtained from a container, is suitable for resolving to the address of the target as seen on the deployment manager computer. If, for example, the list of targets retrieved from a directory is not fully qualified with network domain names, deployment may not be able to proceed in certain network configurations.

2. A check is made to see if the IDPrimer is already installed on the target computer. If not, IDPrimer will be installed first on the target computer. The IDManager tries to deliver the IDPrimer installation package. The delivery method used depends on the target operating environment and the security that has been enabled on it. After the IDPrimer image is copied across to the target computer, its installation is initiated.

As some operating systems do not have a method for remote invocation of the IDPrimer installation, in which circumstances the IDPrimer installation may have to be performed manually.

3. The IDPrimer installer installs itself and the CA Messaging (CAM) component on the target computer. Once the IDPrimer is installed and IDManager has received the 'installation complete' signal from the target computer, package deployment can be initiated. An IDManager that has previously installed an IDPrimer and has authenticated with it can deploy packages without needing to resupply user names or passwords. On subsequent deployments, IDPrimer uses asymmetric cryptographic keys to authenticate and limit access to those managers from which we have already gained access.

For details on using the deployment manager through the CA Virtual Assurance UI, see the *Online Help*.

### More Information

[Prerequisites for Automatically Deploying CA Virtual Assurance Infrastructure](#) (see page 216)

[Notes on Infrastructure Deployment Using IPv6 Addresses](#) (see page 218)

[Protocols for Transferring Packages Employed by IDManager](#) (see page 219)

[Manual Installation of the Infrastructure Deployment Primer Software](#) (see page 219)

[Deployment Primer Installation on Windows](#) (see page 219)

[Deployment Primer Installation on Linux or UNIX](#) (see page 220)

[Provide the Deployment Management Certificate to a Primer Installation](#) (see page 220)

[Deployment Management Certificate on Windows](#) (see page 220)

[Deployment Management Certificate on Linux or UNIX](#) (see page 221)

[Compatibility Libraries for Linux](#) (see page 221)

## Prerequisites for Automatically Deploying CA Virtual Assurance Infrastructure

The Infrastructure Deployment component lets you remotely install agent software to target computers. The installation can only be done using the facilities offered by the underlying operating systems on source and target computers, and is subject to any restrictions imposed by an enterprise network configuration.

The initial step when deploying software is to install a small primer application remotely, the IDPrimer, onto the target computer. This IDPrimer software is responsible for subsequent transfer of software component installation images, and the invocation of their installation. When delivering the IDPrimer to the target computers, the deployment manager must supply user credentials that are valid on the target.

The IDPrimer is transferred to the target system using one of the following mechanisms. If the target computer's operating system is known to the deployment manager, an appropriate transfer mechanism is selected. If the target operating system cannot be determined, each of the following mechanisms is attempted in turn.

- Opening a network share

The deployment manager tries to connect to a Windows network share on the target system. By default, the share name used is ADMIN\$, but this is controlled by an IDManager configuration option. This mechanism is available only from deployment managers running on a Windows-based environment and will only succeed on some Windows targets. Windows variants such as Windows XP Home do not support this deployment mechanism.



- Opening a network connection to the target computer using the SSH protocol, and transferring the primer installation package using SFTP

This mechanism works on any computer running an SSH server, however, it is useful when targeting Linux or UNIX computers.

Note: When deploying to Solaris systems, we recommend that you use either SunSSH v1.1 (or higher) or the latest version of OpenSSH. Refer to the following website for additional details about patches applicable for Solaris platforms and versions: <http://opensolaris.org/os/community/security/projects/SSH>.

If you are running a firewall on the target computer, verify that the SSH port (22) is enabled to permit connection from the deployment manager. You should also check that the SSH server on the target computer is configured to use an RSA key with the 3DES cipher for encryption and the HMAC-SHA1 message authentication code (MAC). Most SSH servers support this configuration by default, but if they do not, consult your SSH server documentation for further instructions.

To successfully deploy to a UNIX or Linux agent, configure the `/etc/ssh/sshd_config` configuration file of your recent SSH implementation as follows:

- Set PasswordAuthentication to Yes
- Set PermitRootLogin to Yes or configure sudo/pfexec as described in section [Remote Deployment to UNIX/Linux Using Non Privileged User Account](#) (see page 195)
- Verify that SFTP subsystem is enabled

Remote Deployment supports deploying software to systems with `/tmp` file system mounted with `noexec` flag.

When deploying to some IBM AIX systems that are running both an IPv4 and IPv6 stack, using an IPv6 address, the target computer SSH server may be listening only on port 22 for IPv4. This would cause the deployment to fail. To correct this, edit the `sshd_config` configuration file and set the `ListenAddress` to `:::`.

**Note:** If you want the SSH communication between the deployment manager and the target computer to be FIPS-compliant, you must verify that the SSH server running on the target is also using FIPS-compliant cryptographic module, apart from setting FIPS-only mode on the deployment manager.

**Important!** Some modern operating systems do not encourage, and sometimes actively prohibit, the remote installation of software. If you try to deploy software to these systems, you will usually see the deployment fail with a status of No Primer Transport. In such cases, installation of software components may be performed in other ways, for example, installation off physical distribution media such as DVD.

Alternatively, you can pre-install or provision machines with the IDPrimer software. This will allow deployment without having to rely on facilities offered by the underlying operating systems. In cases where no authentication has been carried out, valid credentials would need to be supplied before deployments being authorized.

To determine whether automatic deployment is possible in your environment, you can perform some simple checks by running the following standard operating system operations:

- For delivery of the IDPrimer image using Windows shares, you must be able to map a share (default: ADMIN\$) from your deployment manager host computer to each deployment target computer using the target user credentials supplied in the deployment request.
- For delivery of the IDPrimer image using SSH, you must be able to connect using SSH from the deployment manager to the deployment target computers.

### More Information

[Remote Deployment to UNIX/Linux Using Non Privileged User Account](#) (see page 195)

## Notes on Infrastructure Deployment Using IPv6 Addresses

If you are going to use CA Virtual Assurance deployment services in an IPV6 environment, you should be aware of the following prerequisites:

1. The following registry key needs to be set to 1 on the Manager machine (and each deployment (distribution) server):
  - HKLM\System\CurrentControlSet\Services\smb\Parameters\IPv6EnableOutboundGlobal (REG\_DWORD)
2. The three hot fix updates listed below must be applied to Windows 2003 Manager machines:
  - <http://support.microsoft.com/kb/947369/en-us>
  - <http://support.microsoft.com/kb/950092/en-us>
  - <http://support.microsoft.com/kb/974927/en-us>
3. The host name of the target machine must resolve to a global IPv6 address, and the reverse lookup of the IPv6 address must resolve to the same host name.
4. The Infrastructure deployment configuration policy option, usehostnames must have the value 1 on each manager machine. This file is located in the following directory by default:  
C:\Program Files\CA\SC\IDMgrApi\config\SM\idconfig.xml

## Protocols for Transferring Packages Employed by IDManager

IDManager uses the following protocols to transfer packages to target computers when you deploy using the distribution server:

### **Windows Network Share**

Uses this mechanism if the distribution server and the target computer are on Windows.

### **SSH/SFTP**

Uses this mechanism if either the distribution server or the target machine is on Linux or Unix.

For more information about these transfer mechanisms, see [Prerequisites for Automatically Deploying CA Virtual Assurance Infrastructure](#) (see page 216).

## Manual Installation of the Infrastructure Deployment Primer Software

Even if automatic deployment to target computers is not possible for any reason, you can still deploy software if you manually install the primer software on the target computer. This can be done by physically installing the primer package or running the installation through login scripts.

As well as installing the primer software itself, you must install a security key that is generated by the deployment manager you want to use to deploy to your target computers.

## Deployment Primer Installation on Windows

The installation of the deployment primer on a target computer running Windows requires the following actions:

- Make the CA Virtual Assurance installation media (DVD) available on the target computer, or manually copy the primer setup file to the target computer. The primer setup file is stored on the installation media in the following directory:  
`%PROGRAMFILES%\SC\IDMgrApi\packages\private\idprimer\Windows_x86`
- Run IDPrimer\_Setup.exe on the target computer to install the primer.

## Deployment Primer Installation on Linux or UNIX

The installation of the deployment primer on a Linux or UNIX target computer requires the following actions:

- Make the CA Virtual Assurance installation media (DVD) available on the target computer, or manually copy the primer installation image to the target computer. The primer installation image is stored on the installation media in the following directory:

`%PROGRAMFILES%\SC\IDMgrApi\packages\private\idprimer\Linux_x86`

- Change to the directory containing the primer installation image on the target computer and run the following installation command to install the primer:  
`# sh installidp`

## Provide the Deployment Management Certificate to a Primer Installation

The deployment manager generates a certificate that needs to be transferred to the target computer before the primer on the target computer will accept deployment packages. The deployment certificate file is named `dmkeydat.cer`

The location of the certificate is configurable at installation time. You may configure a different file location if you want to store the certificate in a more secure area or in a location shared between two managers providing a failover solution. In the latter case, sharing the certificate enables deployment managers to communicate with IDPrimer components delivered from either manager without the need to resupply authentication credentials.

## Deployment Management Certificate on Windows

On Windows, the deployment certificate is located in the following directory:

`C:\Program Files\CA\SC\IDMgrApi\config\SM`

The certificate file (with the suffix.PMR for example, `MANAGER1 SM.PMR`) must be copied to the primer installation folder on the target computer, which by default is the following:

`\Program Files\CA\SC\IDPrimer`

## Deployment Management Certificate on Linux or UNIX

On Linux and UNIX, the deployment certificate must be copied to the primer installation folder on the target computer, which by default is the following:

```
/opt/CA/SharedComponents/ID/primer/bin
```

## Compatibility Libraries for Linux

The IDPrimer installer assumes that certain dependent libraries are present. If these libraries are not present, installed components may not work properly.

The following table details the software library prerequisites. These libraries must be present on Linux hosts prior to installing IDPrimer.

| Linux Distribution/Version              | Required RPM Packages                                                                                                                                                                                                                                        |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Red Hat 5.0 Enterprise Linux            | glibc 2.3.3-84<br>compat-libstdc++ 33-3.2.3-61<br>For 64-bit versions of the OS:<br>ncurses 5.4-1.3 (i386) or ncurses-devel 5.4-13 (i386)<br>zlib 1.1.4-8.1 (i386) or zlib-devel 1.2.1.2-1.2 (i386)<br>libstdc++.so.6                                        |
| Red Hat 4.0 Enterprise Linux (AS/ES/WS) | glibc 2.3.3-84<br>compat-libstdc++ 296-2.96-132.7.2<br>compat-libstdc++ 33-3.2.3-47.3<br>For 64-bit versions of the OS:<br>ncurses 5.4-1.3 (i386) or ncurses-devel 5.4-13 (i386)<br>zlib 1.1.4-8.1 (i386) or zlib-devel 1.2.1.2-1.2 (i386)<br>libstdc++.so.6 |
| SuSE Linux Enterprise Server 10         | glibc 2.3.3-84<br>compat 2004.4.2-3<br>libstdc++ 3.2.2-38<br>libstdc++.so.6                                                                                                                                                                                  |
| SuSE Linux Enterprise Server 9 (SP3)    | glibc 2.3.3-84<br>compat 2003.5.12-56<br>libstdc++ 3.2.2-38<br>libstdc++.so.6                                                                                                                                                                                |

|                                      |                                                                               |
|--------------------------------------|-------------------------------------------------------------------------------|
| SuSE Linux Enterprise Server 9 (x64) | glibc 2.3.3-84<br>compat 2003.5.12-56<br>libstdc++ 3.2.2-38<br>libstdc++.so.6 |
| SuSE Linux 10.1                      | glibc 2.3.3-84<br>compat 2004.4.2-3<br>libstdc++ 3.2.2-38<br>libstdc++.so.6   |
| SuSE Linux Professional 10           | glibc 2.3.3-84<br>compat 2004.4.2-3<br>libstdc++ 3.2.2-38<br>libstdc++.so.6   |
| SuSE Linux Professional 9.3          | glibc 2.3.3-84<br>compat 2003.5.12-56<br>libstdc++ 3.2.2-38<br>libstdc++.so.6 |
| SuSE Linux Professional 9.2          | glibc 2.3.3-84<br>compat 2003.5.12-56<br>libstdc++ 3.2.2-38<br>libstdc++.so.6 |
| SuSE Linux Professional 9.1          | glibc 2.3.3-84<br>compat 2003.5.12-56<br>libstdc++ 3.2.2-38<br>libstdc++.so.6 |
| SuSE Linux Desktop 9                 | glibc 2.3.3-84<br>compat 2003.5.12-56<br>libstdc++ 3.2.2-38<br>libstdc++.so.6 |

**Note:** To see the latest information about required compatibility libraries and additional system packages, visit the support web site of your Linux supplier.

## Deployment to Windows Vista, Windows 2008 and Windows XP Computers Running Firewall Software

To enable deployment of agents to computers running firewall software, consider the following:

- If the firewall of a target computer running Windows Vista or Windows 2008 operating system is *off* (disabled) and deployment to the computer fails, create or set the following registry variable so that it is a DWORD type with a value 0x1:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\LocalAccountTokenFilterPolicy
```

This is required because User Account Control (UAC) in Windows Vista or Windows 2008 does not automatically grant administrative rights to local users. This occurs even though the local users are members of the Administrators group.

**Note:** Setting this value will result in remote UAC access token filtering being disabled.

Setting this value is only worthwhile if the user has a local administrator account on the computer running Windows Vista or Windows 2008. Domain administrators will not benefit from this change.

- If the firewall of a target computer running Windows Vista or Windows 2008 is *on* (enabled), the following ports should be opened in addition to file sharing ports, to enable deployment to that computer:

### UDP ports

CAM: 4104

File and printer sharing, and so on: 137, 138

### TCP ports

IDManager: 135

File and printer sharing, and so on: 139, 445

- If deployment still fails, the following Outbound Rules in the firewall for Windows Vista or Windows 2008 should be fully enabled:
  - Remote Assistance
  - Network Discovery
  - File and Printer Sharing
  - Core Networking

- To enable agent deployment to Windows XP computers that run firewall software you must perform the following actions manually:
  1. Change Security Policy Network Access: Sharing and security model for local accounts from 'Guest only - local users authenticate as Guest' to 'Classic - local users authenticate as themselves'.

The Classic model allows fine control over access to resources and prevents network logons that use local accounts from being mapped to the Guest account, which typically has Read Only access to a given resource.
  2. Configure the following firewall settings:
    - Allow File and Printer Sharing
    - Open UDP Port 4104
    - Open TCP Port 135



# Chapter 8: Managing Traps

---

This section contains the following topics:

[SNMP Trap Receiver](#) (see page 225)

[SNMP Trap Receiver Configuration File](#) (see page 226)

[SystemEDGE Trap Forwarding](#) (see page 229)

## SNMP Trap Receiver

The SNMP Trap Receiver component is the integration point for CA Virtual Assurance to receive incoming SNMP traps. The SNMP Trap Receiver component listens on a User Datagram Protocol (UDP) port (default 162) that catches the following types of traps:

- Incoming SNMP v1 traps
- Incoming SNMP v2c traps
- Incoming SNMP v3 traps
- v2c notification traps
- v3 notification traps

SNMP Trap Receiver handles SNMP v1, v2c, and v3 traps in the same manner and in some cases, there is no noticeable difference.

All traps that are not viewed as duplicates by the trap storm window are forwarded to ActiveMQ, where other components listen for specific trap IDs.

For SNMPv3 traps, you can configure SNMPv3 trap user authorization using the `sysedgev3.cf` file located in the `Install_Path\CA\productname\conf` directory.

**Note:** For more information about configuring for SNMPv3 in `sysedgeV3.cf`, see the *SystemEDGE User Guide*.

## SNMP Trap Receiver Configuration File

The trapreceiver.conf configuration file is an XML-like file that contains all of the SNMP Trap Receiver settings. This file is provided so you can change these settings as needed. You can configure settings to enable trap blocking, format traps, send traps to CA Virtual Assurance Event Manager, or convert traps into CIM\_SNMPTripIndications. These traps can be sent to the Product Management Modules (PMMs). You can configure the file dynamically without restarting Apache for all settings except for SNMPPort which requires you to restart Apache. All other changes take effect 30 seconds after you save them.

Locate the trapreceiver.conf file in the Install\_Path\CA\productname\config directory and open with a text editor.

Configuration parameters include the following:

### **SNMPPort**

Defines the UDP port to which Trap Receiver connects and listens for incoming SNMP traps. If another application is connected to the port specified by this parameter, Trap Receiver does not start. Instead, it logs a message to the Event Manager and attempts to reconnect to that port every 30 seconds. If this parameter is modified with a new port number, Trap Receiver attempts to connect to the new port.

**Default:** 162

**Note:** If you change the SNMPPort setting after Trap Receiver is connected to the SNMP port, restart Apache for the new port to take effect.

### **TrapStormTimer**

Prevents duplicate trap handling. A configurable value prevents identical traps from coming in more often than the TrapStormTimer value setting. If the value is set to 0, then trap storm is disabled.

**Default:** 30

**PreFilter**

Filters all incoming SNMP traps before being evaluated by the Converter or Subscriber components in the configuration file.

**Default:** All Traps

**FilterType**

Applies only to the PreFilter. Options include the following:

**positive**

Determines which traps are permitted to pass for handling.

**negative**

Determines which traps are blocked.

**Default:** positive if `<FilterType>negative</FilterType>` or `<FilterType>positive</FilterType>` is not included.

**Note:** If you specify the asterisk \* for all fields in negative prefilters, then all prefilters are disabled and traps are blocked.

**ConvertToDCAEvent (Including Filter and Format)**

Specifies a descriptive name for the converted trap and can be duplicated when specifying other trap names.

**Example:** `<Converter Name="eHealthTrap25">`

**Format**

Specifies a string containing %s to represent the values that are defined by the list of varbind values below the format string.

**Varbind**

Specifies the variables that are used to replace the %s in the format string.

Valid values include the following:

SourceIP

EnterpriseOID

GenericTrapID

SpecificTrapID

**Note:** The SpecificTrapID filter entry is only evaluated if generic type from the incoming trap is 6 (enterpriseSpecific). If GenericTrapID is \* and the generic type from the incoming trap is 0-5, then both GenericTrapID and SpecificTrapID are matches.

Varbind1, Varbind2 - up to and including the number of expected varbinds in the trap. The number of variables below the Format must match the number of %s values in the Format.

### EventSource

(Optional) Substitutes the SourceIP of the trap to another value, such as a specific varbind of the incoming trap. Some incoming traps originate from the manager and are not from the node where the trap originated.

**Default:** SourceIP

### EventTarget

(Optional) Substitutes the localhost node name with a value from the incoming trap, such as a specific varbind.

**Default:** localhost

### SubscribedIndicationHandlers (Including Filter and Topic)

Specifies a descriptive name for the subscriber and can be duplicated when specifying other trap names.

**Example:** <Subscriber Name="SUNZoneAim\_PMM">

### Topic

Specifies the ActiveMQ topic name where messages are sent for CA Virtual Assurance.

### Filter

These settings apply to Filter and PreFilter. They allow incoming traps to be evaluated to determine if each Converter or Subscriber uses them. If the PreFilter blocks traps, then the other Filters do not evaluate them.

### SourceIP

Specifies any regular expression string variable that is a subset of the range of IP addresses.

#### Examples:

111.\*.\*.\*

112.223.\*.\*

112.223.334.5

\*

**Limits:** IPv4 addresses only.

### EnterpriseOID

Specifies any regular expression string variable that is a subset of the range of enterprise OIDs. Use backslash (\) before any period (.) that appears in the regular expressions. The period (.) is a special character in regular expression syntax.

**GenericTrapID**

Specifies an integer-based regular expression that is used to compare the trap Protocol Data Unit (PDU) TrapType field. Valid values are \*, a single integer 0 through 6, or a range (0-1) or (0-6).

Valid SNMP TrapType values include the following:

coldStart(0)

warmStart(1)

linkDown(2)

linkup(3)

authenticationFailure(4)

egpNeighborloss(5)

enterpriseSpecific(6)

**SpecificTrapID**

Specifies an integer-based regular expression that is used to compare the trap PDU's TrapType field.

Valid values can be *one* of the following:

\*

A single integer

A range from (1200-2225) or (165201-165299)

**Comments**

Specifies comments which can be on the same line or separate lines.

<!--

Defines the beginning of a comment.

-->

Defines the end of a comment.

**Example:**

<!-- If the port is already connected, it does not change dynamically. Restart Apache. -->

## SystemEDGE Trap Forwarding

The SNMP Trap Receiver can receive any trap from SystemEDGE when configured to forward traps to the CA Virtual Assurance manager. The incoming SNMP traps are evaluated using a filter and can then be sent to the PMMs or converted to CA Virtual Assurance events.

**Note:** For more information about forwarding SNMP traps using the `sysedge.cf` and `sysedgeV3.cf` files, see the *SystemEDGE User Guide* in this bookshelf.

# Chapter 9: Reporting

---

This section contains the following topics:

[Report Availability](#) (see page 231)

[Run Reports](#) (see page 231)

## Report Availability

Some reports in the CA Virtual Assurance user interface require functionality that is not provided by a standalone installation of CA Virtual Assurance and therefore may not be available. The following reports require CA Virtual Assurance to be installed with CA Server Automation:

- Package Delivery Failure/Success
- Packaging Job History
- All Reservation Manager Reports

All other reports function correctly with a CA Virtual Assurance standalone installation.

## Run Reports

You can run reports based on the metrics that CA Virtual Assurance collects about the servers and services in your data center.

**Follow these steps:**

1. Click Reporting.

The Reports page appears. At the Reports level, a menu provides frequently requested selections.

2. Expand the folder for the type of report you are interested in.

The specific reports that are available for the category appear.

3. Select the report that you want to run.

A progress indicator appears in the right pane and an empty pane opens with the report description.

4. Select start and end times from the drop-down list and any other prompts requested, then click Create Report.

A progress indicator appears in the right pane and the charted data appears in this pane.



# Chapter 10: Remote Monitoring

---

This section contains the following topics:

[Overview](#) (see page 233)

[Advantages of Remote Monitoring](#) (see page 233)

[Features and Benefits](#) (see page 234)

[Architecture](#) (see page 236)

[Use Case Scenario](#) (see page 238)

[Configuration Prerequisites](#) (see page 239)

[Configuring Remote Monitor Systems](#) (see page 240)

[Managing Systems Using Remote Monitoring](#) (see page 243)

## Overview

Remote Monitoring (RM) lets you monitor the health state of agent-less systems. RM provides the flexibility of monitoring systems without the need to install the monitoring agents (such as SystemEDGE) on the remote systems.

RM employs a mid-level manager named RM AIM to monitor the remote systems. RM AIM collects the metrics information using WMI queries on the remote Windows systems.

## Advantages of Remote Monitoring

Remote Monitoring involves agent-less rather than agent-based technology and there are advantages to both strategies. Use this information when deciding whether to use RM or deployed agents.

RM offers the following benefits:

- Costs less to set up, configure, and deploy
- Simplifies software upgrades and maintenance
- Deploys quickly and is less intrusive on the monitored environment
- Utilizes fewer resources on the managed server

A deployed agent offers the following benefits:

- Provides more detailed data and higher levels of functionality for the monitored servers and applications
- Requires less network bandwidth to operate
- Provides a higher degree of scalability, scaling to thousands of servers
- Continues to monitor server health and conduct data gathering when network connections are unavailable (as agent can work autonomously)
- Provides stronger command and control functions over the managed servers

## Features and Benefits

Remote Monitoring provides *seamless* integration of monitoring from an end-user perspective (that is, equal look-and-feel of management interfaces for the monitored systems whether by agent or RM).

RM includes features that let you manage systems by monitoring health states and key performance indicator (KPI) metrics. RM provides reports on system status and utilization metrics. RM includes benefits such as resilience, scalability, integration, and automation. The primary features and benefits are described in the sections that follow.

## Agent-less Monitored Systems

Remote Monitoring enables seamless health monitoring for systems managed with agent-based and agent-less technologies.

The RM manager component (RM PMM) creates CIM system objects representing the RM systems and their health state.

This information is presented in the Dashboard and the Resources Panel.

### **Follow these steps:**

1. Open Resources, Explore and expand the Remote Monitoring folder.  
The discovered systems appear in the components tree.
2. Select a system.  
The page of that system appears in the right pane.
3. Open the Remote Monitoring tab.  
The agent-less gathered data appears.

## Key Performance Indicator Metrics

Remote Monitoring collects and provide Windows metric information by performing WMI queries on the monitored RM systems. A rich set of information is available in various Win32 CIM classes, made available through the RM AIM.

## Visualization

The RM UI lets you configure the following information:

- What systems are remotely monitored
- What metrics are collected for those systems
- If and how those metrics are monitored (including severity and threshold)

## Configuration

Remote Monitoring monitors KPIs out-of-the-box on a remote system when it is selected for monitoring without requiring configuration of the monitoring being performed. You can adjust the out-of-the-box monitoring thresholds to suit your needs.

You can also define and store configuration settings in a configuration set, which can then be assigned to one or more RM systems.

## Access Control

When a user logs in to the UI as admin or as a nonadmin user, security mechanisms provide authentication and authorization functionality. Remote Monitoring allows or disallows certain actions (such as configuring an RM system) based on whether the user is an admin or nonadmin user.

The RM AIM accesses the RM system through a WMI connection to the root\CIMV2 namespace (using DCOM). The local RM system administrator user and password credentials are required for access. These credentials (provided by the user when an RM system is to be monitored) are stored in a file using password encryption.

## Resilience

The RM AIM is a separate process from SystemEDGE; an error in the RM AIM does not cause SystemEDGE to crash. If the RM AIM crashes or no longer responds to SystemEDGE requests, the *RM AIM alive* check in SystemEDGE restarts it.

### Scalability

There is one RM AIM per SystemEDGE and each RM AIM can monitor approximately 200 RM systems. There is a single RM PMM per manager and each RM PMM can manage approximately 20 RM AIMs. The default configuration set contains ten monitored metrics with two monitors for each metric.

In terms of SystemEDGE scalability, this results in the following:

- $10 * 2 * 200 = 4000$  monitorTable entries
- $10 * 200 = 2000$  aggregateTable entries

### Integration

RM monitor information is exposed in an SNMP MIB to enable easy access for Spectrum, eHealth, and NSM managers.

### Automation

The RM AIM includes a command line utility (rmonwatch), which allows remote configuration of RM systems and their credentials using a script.


## Architecture

The following diagram provides an overview of the main RM components.

One or more RM AIMs perform monitoring for Windows servers through WMI over DCOM/RPC. Within a particular site or subnet, direct TCP connectivity from the AIM to the monitored Windows servers is required. The AIM is deployed through the deployment component.

A Platform Management Module (RM PMM) provides the interface to the manager infrastructure and creates managed objects in the CIM object model. The PMM communicates with the RM AIM using SNMP.

### Interaction Between Remote Monitoring Components




## Use Case Scenario

Consider the following use case scenario for Remote Monitoring. An enterprise offers a web book store consisting of services to order books and services to browse for books.

- The order services are available from two servers in Munich and one server in London
- The browse services are available from two servers in London and one server in Munich

The servers are GERSERV1 and GERSERV2 in Munich and UKSERV1 and UKSERV2 in London. They are configured for load balancing and failover.



Monitoring of the services depends whether it is an order service or a browse service. In this example, two configuration sets (one for each service type) are defined. They comprise queries and monitors for the following information:

- CPU

The percentage of the total CPU idle time.

- FSys

The free space of the logical disk important for the respective service type (C: for an order service and D: for a browse service).

- Proc

The working set size of the order process (single process order.exe) or the sum of the working set size of all browse processes (group of processes browse).

For each monitored system, depending on the service role (order, browse, or both), the following configuration sets are assigned:

| SystemName | ConfigSet |
|------------|-----------|
| GERSERV1   | order     |
|            | browse    |
| GERSERV2   | order     |
| UKSERV1    | order     |
|            | browse    |
| UKSERV2    | browse    |

## Configuration Prerequisites

Before you configure remote agents, verify that the following prerequisites are met.

- Firewall and port requirements for RM systems

The RM AIM system accesses the RM system through a WMI connection. WMI uses DCOM communication which uses an End Point Mapper (EPMAP) Port TCP (135) and a DCOM TCP port which EPMAP dynamically identifies.

To simplify configuration, the RM AIM must be located within the same firewall boundaries as the RM systems.

**Note:** For more information about using a fixed port search for the article "Setting Up a Fixed Port for WMI" on the Microsoft MSDN website.

- Firewall and port requirements for the manager system

The RM AIM utilizes the SNMP infrastructure provided by SystemEDGE; it does not require additional ports.

RM configurations are performed using SNMP. Because the configuration data includes passwords, RM uses password encryption.

- SystemEDGE ports and SNMP

The manager system accesses the SystemEDGE system through SNMP, which requires that the SNMP port (UDP 161 incoming) is open on the SystemEDGE system. The SystemEDGE system sends SNMP traps (UDP 162 outgoing).

- SystemEDGE ports and policy-based configuration

The manager system accesses the SystemEDGE system through CAM, which requires that the UDP (4104) or TCP (4105) port are opened on the SystemEDGE AIM system. The SystemEDGE AIM system uses CAM to send messages to the manager system.

- WMI access best practices

The RM AIM connects to RM Systems using WMI and requires credentials. As a best practice the RM systems must be a member of an AD Domain (for example, RIVER). This membership lets you use a domain account and avoids the need to define local user accounts on each RM System. Create a CARMuser domain account that is a member of the Domain Admins group of the AD Domain.

When user credential settings are prompted for during RM installation, provide the domain account (for example, RIVER\CARMuser) with the password. For any system member of this domain, no additional configuration is required.

**Note:** If necessary, you can restrict the CARMuser access rights so the user is not a member of the Domain Admins group. In this case, configure WMI Namespace access and DCOM access. For more information about defining WMI Namespace access and DCOM access, see the Microsoft website.

## Configuring Remote Monitor Systems

A configuration set is the entity assigned to an RM System; it defines what metrics (WQL queries) are collected and how those metrics are monitored.

A configuration set consists of several configuration items. The configuration items consist of a metric definition (WQL query) and a monitoring definition (threshold, severity, and so on).



RM provides the following configuration sets with out-of-the box metric and monitoring definitions:

- default
- extended
- metricDisk
- metricFS
- metricNet

If different metrics must be monitored for an RM system with different threshold and severity settings, clone the out-of-the-box configuration set and adjust the cloned set to your system-specific monitoring needs.

The following table lists the RM metrics and the config sets to which they belong.

| <b>Metric</b>       | <b>Config Set</b> |
|---------------------|-------------------|
| CPU_PercentIdle     | default           |
| Disk_PercentIdle    | default           |
| Event_SystemErrors  | default           |
| FSys_FreeMB         | default           |
| FSys_FreeMBDecrease | default           |
| Mem_PercentUsed     | default           |
| Net_MACAddress      | default           |
| Net_MACIndex        | default           |
| Net_QueueLength     | default           |
| Proc_PercentCPU     | default           |
| Proc_PercentMemory  | default           |
| Srvc_StoppedAuto    | default           |
| Sys_LastBootTime    | default           |
| Sys_LastLocalTime   | default           |
| Sys_OSInfo          | default           |
| Sys_PhysMemKB       | default           |
| Disk_ReadPerSec     | extended          |
| Disk_WritePerSec    | extended          |
| Disk_QueueLength    | extended          |

| Metric                              | Config Set |
|-------------------------------------|------------|
| Mem_FreeMB                          | extended   |
| Mem_FreePages                       | extended   |
| Mem_NonPagedMB_3GB                  | extended   |
| Mem_PagedMB                         | extended   |
| Mem_PagedMB_3GB                     | extended   |
| Mem_PagingPerSec                    | extended   |
| Mem_NonPagedMB                      | extended   |
| Net_PercentBusy                     | extended   |
| Sys_Is64bit                         | extended   |
| Sys_Has3GBSwitch                    | extended   |
| Sys_OSType                          | extended   |
| BIOS_Version                        | extended   |
| BIOS_SerialNumber                   | extended   |
| Disk_AvgDiskBytesPerRead            | metricDisk |
| Disk_AvgDiskBytesPerWrite           | metricDisk |
| Disk_AvgDiskReadQueueLength         | metricDisk |
| Disk_AvgDiskWriteQueueLength        | metricDisk |
| Disk_DiskWritesPersec               | metricDisk |
| Disk_PercentDiskReadTime            | metricDisk |
| Disk_PercentDiskWriteTime           | metricDisk |
| Disk_SplitIOPerSec                  | metricDisk |
| Net_PacketsOutboundErrors           | metricNet  |
| Net_PacketsReceivedErrors           | metricNet  |
| Net_PacketsReceivedDiscarded        | metricNet  |
| Net_PacketsReceivedNonUnicastPersec | metricNet  |
| Net_PacketsReceivedUnicastPersec    | metricNet  |
| Net_PacketsSentNonUnicastPersec     | metricNet  |
| Net_PacketsSentUnicastPersec        | metricNet  |
| FSys_PercentFreeSpace               | metricFS   |

**Note:** For more information about the RM metrics, see the *Reference Guide*.

## Collection Engine and Report Support for Remote Monitoring Metrics

The Collection Engine and Reports support a fixed set of RM metrics in the default config set.

As a result, assign the default config set (or a config set or group of sets containing those metrics) to all systems for which you want to use reports.

The supported default config set metrics are as follows:

- CPU\_PercentIdle
- Disk\_PercentIdle
- Event\_SystemErrors
- Mem\_PercentUsed
- FSys\_FreeMB
- Fsys\_FreeMBDecrease
- Net\_QueueLength
- Proc\_PercentCPU
- Proc\_PercentMemory
- Svc\_StoppedAuto

## Managing Systems Using Remote Monitoring

Access the RM information and settings necessary to manage your systems by highlighting a managed resource in the Resource pane and clicking Remote Monitoring. The Remote Monitoring pages let you perform the following actions:

- Add remote systems for monitoring
- Manage queries
- Manage credential settings
- Create configuration sets
- Manage configuration entries

**Note:** For more information about using the UI RM sub-tabs, see the *User Help*.

For the Dashboard, the following RM modules are available:

- CA SystemEDGE Machines Status
- CA SystemEDGE Objects Status



# Chapter 11: Active Directory and Exchange Server AIM

---

This section contains the following topics:

[Overview](#) (see page 245)

[Architecture](#) (see page 246)

[Features](#) (see page 247)

[Advantages](#) (see page 248)

[Performance](#) (see page 248)

[Install ADES AIM](#) (see page 248)

[Configure the Environment to Enable ADES AIM Monitoring](#) (see page 250)

[Configure ADES AIM with NodeCfgUtil in Command Mode](#) (see page 251)

[Configure ADES AIM with NodeCfgUtil in Dialog Mode](#) (see page 252)

[Uninstall the ADES AIM](#) (see page 254)


## Overview

Active Directory and Exchange Server AIM (ADES AIM) lets you monitor Active Directory and Exchange Server environments. This AIM enables configuration of Active Directory and Exchange Server environments, and monitoring of the key performance indicators.

## Architecture

Active Directory and Exchange Server AIM addresses scalability and performance issues by its distributed and reusable architecture.

The following diagram provides an overview of the AIM:



**Note:** The ADES AIM is supported for use with SystemEDGE. CA Virtual Assurance for Infrastructure Managers provides Remote Deployment support and local configuration support based on the NodeCfgUtil utility. For details to install the AIM remotely, see the Remote Deployment section in the CA Virtual Assurance Administration Guide.

## Features

ADES AIM includes features that let you monitor the health states and key performance indicator (KPI) metrics. The primary features and functionality of ADES AIM are as follows:

- Monitors message records manager, logical disk usage, logical disk read/write, and memory and CPU usage of the mailbox server.
- Monitors network latencies, queue, and mail delivery metrics of the hub transport server.
- Monitors active directory performance and replication.
- Supports Remote Deployment through CA Virtual Assurance.

Active Directory and Exchange Server AIM has the following components:

### **Domain Data Collector**

Collects configuration data from domains identified for monitoring.

### **Exchange Server Data Collector**

Collects configuration data from Exchange Servers identified for monitoring.

### **Domain Performance Data Collector**

Collects Performance data from Exchange Servers identified for monitoring.

### **Exchange Performance Data Collector**

Collects performance data from domains identified for monitoring.

## Advantages

Active Directory and Exchange Server AIM lets you monitor Active Directory and Exchange Server environments.

Active Directory and Exchange Server AIM offers the following benefits:

- Quick and seamless deployment.
- Less intrusive on the monitored environment.
- Low resource utilization on the managed server.
- Less bandwidth-intensive during operation.

## Performance

To obtain optimal performance from Active Directory and Exchange Server AIM, do the following:

- Deploy ADES AIM in geographical proximity to the servers it is intended to monitor. It avoids end-to-end metric poll performance degradation due to network latencies in geographically distributed Active Directory and Exchange Server environment.
- Configure the poll intervals proportional to the number of servers to obtain optimal performance.

## Install ADES AIM

Use Remote Deployment to install ADES AIM. Specify the following parameters for the package wrapper:

**User**

Specifies the name of an administrative user.

**Password**

Specifies the password of that user.

**Domainname**

Specifies the name of the domain that is monitored through the ADES AIM.



**Management Entity**

Specifies the managed entity.

**0**

Monitors Active Directory only.

**1**

Monitors Exchange Server only.

**2**

Monitors both the Active Directory and Exchange Server.

**Management Mode**

Specifies the option for providing management.

**0**

Monitors all the hosts in the current domain.

**Note:** Hosts of Child domains are not monitored.

**1**

Monitors a specific host of the domain.

## Configure the Environment to Enable ADES AIM Monitoring

Apply the following settings to enable Microsoft Active Directory and Exchange Server monitoring:

### Prerequisites:

Verify that the following are installed on the same host where ADES AIM is installed:

- SystemEDGE
- Advanced Encryption
- .Net 3.5 or higher version
- Windows PowerShell 2.0
- Exchange Management Tools

**Note:** Exchange Management Tools SP3 is required for monitoring Exchange Server 2007.

### Configure PowerShell on Exchange Servers

As an administrator, apply the PowerShell configuration settings on Exchange Server systems to enable remote connectivity and Exchange Server monitoring.

### Follow these steps:

1. Navigate to Start menu, Programs, Accessories, and launch Windows PowerShell.  
PowerShell command prompt appears.
2. Enter the following command:  
`Enable-PSRemoting`  
WinRM Quick configuration message appears.
3. Specify Y to run the Set-WSManQuickConfig command.  
WinRM setup initiates remote management and creates a WinRM listener to accept WS-Man requests. A confirmation to register session configuration appears.
4. Confirm Registering session configuration by specifying Y.  
Registering session configuration starts.
5. Enter the following command to set trustedhosts:  
`set-item wsman:localhost\client\trustedhosts -value *`  
Adds all the computers into the trustedhosts list and establishes connectivity.
6. Enter the following PowerShell command to restart the WinRM service:  
`restart-Service winrm`  
Enables the new TrustedHosts settings and Exchange Server monitoring.

## Configure ADES AIM with NodeCfgUtil in Command Mode

NodeCfgUtil.exe lets you modify the AIM configurations for Active Directory and Exchange Server. When you use the utility in command mode, you can only add managed nodes to AIM configuration.

This command has the following format:

```
nodecfgutil esad -u user -p passwd -d domainname -e entity -o option
```

**esad**

Specifies the Active Directory and Exchange Server AIM.

**-u *user***

Specifies the name of a user with domain admin & exchange admin privileges. In the format <user>@<fqdn>.

**-p *password***

Specifies the user password.

**-d *domainname***

Specifies the name of the domain that is monitored through the ADES AIM.

**-e *entity***

Specifies the managed entity.

**0**

Specifies that the AIM monitors only Active Directory servers in the domain.

**1**

Specifies that the AIM monitors only Exchange servers in the domain.

**2**

Specifies that the AIM monitors both Active Directory and Exchange servers in the domain.

### **-o option**

Specifies the option for providing management.

**0**

Specifies that the ADES AIM discovers and monitors all servers in the domain.

**Note:** Hosts of Child domains are not monitored.

**1**

Specifies that the ADES AIM discovers all servers in the domain but monitors only those that are configured through the manager.

### **Follow these steps:**

1. Open a command prompt on the system on which the AIM resides.

The command prompt appears.

2. Enter the following command:

```
nodecfgutil esad -u user@fqdn -p passwd -d fqdomainname -e {0|1|2} -o {0|1}
```

Authenticates and stores the input credentials for Active Directory and Exchange Server.

The utility writes a configuration file for ADES (esad.cfg) to the Sysedge\_Install\_Path\plugins\AIPCommon directory.

## Configure ADES AIM with NodeCfgUtil in Dialog Mode

NodeCfgUtil lets you modify the AIM configurations for Active Directory and Exchange Server. When you use the utility in dialog mode, you can configure which nodes the ADES AIM manages.

### **Follow these steps:**

1. Open Windows Explorer on the system on which the AIM resides, change to the *SystemEDGE\_InstallPath*\plugins\AIPCommon directory, and start NodeCfgUtil.exe. NodeCfgUtil discovers the installed ADES AIM and lists only those in subsequent dialogs.
2. Enter **1** to add a new managed node.

NodeCfgUtil can only discover installed AIM.

3. Follow the on-screen instructions to complete the configuration. Each node requires a valid user name and password for authentication.
4. After a successful configuration, enter 0 to return to previous menus, or to exit the utility.

NodeCfgUtil writes a configuration file for ADES (esad.cfg) to the *SystemEDGE\_InstallPath\plugins\AIPCommon* directory. You can also use the NodeCfgUtil utility to edit or remove existing entries. The corresponding dialogs are self-explaining.

The ADES AIM is a multi-instance AIM. So you can repeat this procedure and add more entities that you want to manage with this AIM.

**Note:** When removing a server entry, NodeCfgUtil marks this server for deletion, but it does not remove the server from the user interface or database. To remove the server from the user interface and database, rediscover the Windows server on which the AIM runs. The new discovery process detects the deletion mark and removes the corresponding server from the user interface and database.

### Examples

The following example shows the Install Managed Node dialog for mydomain that has been successfully added to the configuration of the ADES AIM. Management Entity is set to Active Directory. Management Mode is set to Entire Domain. For details, see the NodeCfgUtil command mode. Managed Mode corresponds to the option parameter in the NodeCfgUtil command mode.

\*\*\*\* Choose Managed Node \*\*\*\*

1. Microsoft Exchange Server and Active Directory
0. Go Back to Previous Menu

\*\*\*\*\*

Enter choice: 2

Enter following information for the Microsoft Exchange Server and Active Directory Node...

(At any point to go back to previous menu, Enter 'CTRL Q')

1. Domain Name: mydomain
2. User Name: administrator
3. Password: \*\*\*\*\*
4. Management Entity: 0
5. Management Mode: 0

CAAC1016 Authenticating, please wait...

CAAC1018 Credential authentication SUCCESSFUL.

Press any key to continue . . .

## Uninstall the ADES AIM

You can uninstall the SystemEDGE agent and the ADES AIM from the Windows Add or Remove Programs or from the command line. The uninstaller removes ADES AIM from your system. You can specify whether to remove the configuration data from the data directory. To uninstall SystemEDGE or ADES AIM from the Windows Add or Remove Programs Window

Follow these steps:

1. Select Start, Settings, Control Panel, Add or Remove Programs.

The Add or Remove Programs window appears and lists the following components:  
CA\_SystemEDGE\_ESAD

2. According to the uninstallation sequence, right-click the appropriate component and select Uninstall.

On selection of SystemEDGE, a dialog prompts you to specify whether to preserve the configuration files.

3. Click Yes or No.

A dialog charts the uninstallation process. When the uninstallation process completes, the dialog closes.

Uninstallation is complete.

# Chapter 12: Scenarios and Best Practices

---

This section contains the following topics:

[How to Configure the vCenter Server Management Components](#) (see page 256)

[How to Use Policy Actions to Identify Performance Issues](#) (see page 271)


[Scalability Best Practices](#) (see page 274)

[How to Apply Policy and Layered Templates to Servers](#) (see page 289)

## How to Configure the vCenter Server Management Components

The following diagram provides an overview about the required actions. The diagram includes corresponding troubleshooting strategies in case of connection problems.

**How to Configure the vCenter Server Management Components**





**Follow these steps:**

[Review Requirements](#) (see page 257)

[Review Interactions Between vCenter Server Management Components](#) (see page 258)

[Add a New vCenter Server Connection to the Manager](#) (see page 260)

[Add the AIM Instance for the vCenter Server](#) (see page 264)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 270)

[Troubleshoot the vCenter Server Connection](#) (see page 261)

[Troubleshoot the vCenter AIM Instance Connection](#) (see page 265)

## Review Requirements

Review the following requirements before you start configuring the vCenter Server management components of CA Virtual Assurance for Infrastructure Managers:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You have a basic understanding of CA Virtual Assurance, CA SystemEDGE, and VMware vSphere.
- You can access a CA Virtual Assurance manager installation that includes the vCenter Platform Management Module (PMM), vCenter Application Insight Module (AIM), and Monitoring Agent (CA SystemEDGE).
- You can access the CA Virtual Assurance user interface.
- You have valid credentials available (user name and password) to access the vCenter Server of the new vSphere environment that you want to manage.
- You have found out which protocol (HTTP or HTTPS) and port to use for accessing the vCenter Server of the vSphere environment through web services. Default: HTTPS, Port 443
- You have verified that the new vSphere environment and its vCenter Server are running properly.
- If the vCenter PMM and vCenter AIM are installed on different systems, you have verified that the SNMP settings on these systems are consistent. Read and write community strings and SNMP port number must be identical.
- You have verified that the CA Virtual Assurance manager has discovered any remote vCenter AIM Servers that you want to use.

### More information:

[Review Interactions Between vCenter Server Management Components](#) (see page 258)

[Add a New vCenter Server Connection to the Manager](#) (see page 260)

[Add the AIM Instance for the vCenter Server](#) (see page 264)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 270)

## Review Interactions Between vCenter Server Management Components

As a System Administrator, you want to manage a new VMware vSphere environment with CA Virtual Assurance. CA Virtual Assurance allows you to manage the physical and virtual resources of one or more vSphere environments dynamically.

vSphere consists of one vCenter Server, physical ESXi hosts, and a virtual infrastructure that runs on the ESXi hosts. A vCenter Server is the central point of control of a vSphere environment with its entire virtual infrastructure. This infrastructure can consist of datacenters, clusters, resource pools, vApps, VMs, virtual devices, and virtual switches. To manage vSphere, CA Virtual Assurance requires network connections between its vCenter Platform Management Module (PMM), vCenter Application Insight Module (AIM), and VMware vCenter Servers. To establish these network connections, configure the CA Virtual Assurance vCenter Server management components, that is, vCenter PMM and vCenter AIM.


The vCenter AIMs is a SystemEDGE agent plug-in which extends the functional scope of SystemEDGE. The vCenter AIM enables SystemEDGE to monitor the performance of multiple vSphere environments and to evaluate the states of monitored vSphere resources. Typical monitored resources are virtual CPUs, virtual memory, virtual switches, virtual disks, resource pools, vApps, and so on. Based on thresholds, SystemEDGE and the vCenter AIM determine the status of a monitored resource and propagates this information to the CA Virtual Assurance manager using SNMP.

The vCenter PMM is a component of the CA Virtual Assurance manager. The PMM is responsible for providing connection and support for all VMware vCenter operations using web services. The PMM manages connections with vCenter Servers, performs vSphere-related operations, retrieves data from the vCenter AIM, and populates the CA Virtual Assurance Management Database. Typical operations include but are not limited to: Creating, starting, stopping, or cloning a VM, adding, or removing CPU shares, adding memory to the VM while the VM is running, and so on.

Because the vCenter PMM and the AIM interact with each other, CA Virtual Assurance can dynamically manage multiple vSphere environments. CA Virtual Assurance can run operations automatically controlled by thresholds, status, and values gathered by the AIM. For example, CA Virtual Assurance can add or remove CPU shares dynamically according to the workload of a VM.

The following diagram shows the interaction of the affected components in an example environment of four vSphere environments represented by four vCenter Servers. In general, the vCenter PMM and each vCenter AIM with its multi-instance support can connect to multiple vCenter Servers. The number of connections shown in the diagram do not specify any limitations. The required network connections are based on TCP/IP, SNMP, and web services.

### Interaction Between vCenter Server Management Components



When you have configured the CA Virtual Assurance components successfully, CA Virtual Assurance discovers the new vSphere environment. After a successful discovery, the vCenter Server of the vSphere environment and its virtual infrastructure appear in the resources tree of the CA Virtual Assurance Explore pane. You can then manage the new vSphere environment.

**More information:**

[Add a New vCenter Server Connection to the Manager](#) (see page 260)

[Add the AIM Instance for the vCenter Server](#) (see page 264)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 270)

## Add a New vCenter Server Connection to the Manager

You can add a vCenter Server connection using the Administration tab of the CA Virtual Assurance user interface.

**Follow these steps:**

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select vCenter Server from the Provisioning section in the left pane.

The right pane refreshes and displays the managed vCenter Servers, associated vCenter AIM Servers, and the AIM Instance for the vCenter Server.

3. Click  (Add) on the vCenter Servers pane toolbar.

The New vCenter Server dialog appears.

4. Enter the required connection data (server name, user, password, protocol, port), enable Managed Status (checkbox), and click OK.

If the network connection has been established successfully, the vCenter Server is added to the top right vCenter Servers pane with a green status icon. CA Virtual Assurance discovers the vCenter Server automatically.

If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the vCenter Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added. For troubleshooting the connection, see Troubleshoot the vCenter Server Connection (see page 261).

**More information:**

[Add the AIM Instance for the vCenter Server](#) (see page 264)


[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 270)

[Troubleshoot the vCenter Server Connection](#) (see page 261)

## Troubleshoot the vCenter Server Connection

The vCenter Server connection has failed. Follow the troubleshooting information indicated in the following diagram:

### How to Troubleshoot the vCenter Server Connection



**Follow these steps:**

[Troubleshooting: vCenter Server Connection Failed](#) (see page 262)

[Add the AIM Instance for the vCenter Server](#) (see page 264)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 270)

## Troubleshooting: vCenter Server Connection Failed

**Symptom:**

After I have added a new vCenter Server connection under Administration, Configuration, the validation of the connection to the vCenter Server failed.

**Solution:**

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used vCenter Server connection data (server name, user, password, protocol, port) is still valid. If necessary, update the connection data.
- Verify, if the vCenter Server system is running and accessible.
- Verify, if the VMware Management Service on the vCenter Server system is running properly.

**To update the vCenter Server connection data**

1. Click  (Add) or  (Edit) that is associated with the failed connection.

The New or Edit vCenter Server dialog appears.

2. Add the valid server name, user, password, protocol, and port. Enable Managed Status and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the vCenter Server cannot be established, continue with the next procedure.

**To verify, if the vCenter Server system is running and accessible**

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

```
nslookup <vCenter Server Name>
ping <IP Address of vCenter Server>
```

2. Verify the output of the commands to find out whether the vCenter Server has a valid DNS entry and IP address.

If the vCenter Server is not in the DNS, add the vCenter Server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

If the vCenter Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress <vCenter Server Name>
```

Enter the correct IP address and vCenter Server name. For example:

```
192.168.50.50 myvCenter
```

4. Click  (Validate) in the upper-right corner.

If the vCenter Server credentials and connection data are correct and you can ping the vCenter Server, the connection can still fail. In this case, it is possible that the vCenter Server causes the problem. If the connection to the vCenter Server cannot be established, continue with the next procedure.

**To verify, if the VMware Management Service on the vCenter Server system is running properly**

1. Contact the vSphere Administrator to access the vCenter Server system.
2. Log in to the vCenter Server system and open Administrative Tools, Services from the Start menu.

The Services window opens.

3. Select the service *VMware VirtualCenter Server*. Start or restart the service.
4. Change to the CA Virtual Assurance user interface, vCenter Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Virtual Assurance validates the vCenter Server connection.

If the connection to the vCenter Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

Work with the vSphere administrator or VMware support to fix the vCenter Server connection problem.

## Add the AIM Instance for the vCenter Server

After adding a new vCenter Server connection to the CA Virtual Assurance manager, add a vCenter AIM instance to manage the new vCenter Server. CA Virtual Assurance then discovers the entire vSphere environment with all its physical and virtual components, such as vCenter Server, ESX Servers, VMs, and so on.

### Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select vCenter Server from the Provisioning section in the left pane.

The right pane refreshes and displays the managed vCenter Servers, associated vCenter AIM Servers, and the AIM Instances for managed vCenter Servers.

3. Click  (Add) on the vCenter AIM Servers pane toolbar.

The New vCenter AIM Server dialog appears.

4. Open the vCenter AIM Server drop-down list.

The list of discovered vCenter AIM Servers appears. If you have installed the vCenter AIM on the local system, the name of the local system appears in the list too.

5. Select a vCenter AIM Server from the drop-down list.

CA Virtual Assurance populates the vCenter Server drop-down list with the vCenter Servers listed in the vCenter Servers pane. That is, you can only manage those vCenter Servers for which your CA Virtual Assurance manager has a valid connection established.

6. Select the vCenter Server you want to manage and click OK.

A new AIM instance for the selected vCenter Server is added. If the instance is not in an error or stopped state, CA Virtual Assurance starts to discover the associated vSphere environment. When the discovery process is complete, you can start managing the virtual and physical resources of vSphere.

### More information:

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 270)

[Troubleshoot the vCenter AIM Instance Connection](#) (see page 265)




## Troubleshoot the vCenter AIM Instance Connection

The vCenter AIM Connection is in not-ready status. One of the following status icons appears:

-  Discovery in progress - Wait until the platform manager synchronizes all data.
-  Multiple instances - This instance is managed by multiple AIMs. Remove the duplicate instance.
-  Error - Unable to connect to the AIM. Check the network configuration.
-  Warning - Some instances on the AIM are still in error state.
-  Disabled - This instance is not managed.

Follow the troubleshooting information indicated in the following diagram:

### How to Troubleshoot the vCenter AIM Instance Connection



#### More information:

[Troubleshooting: vCenter AIM Instance Status Icon Shows Discovery in Progress](#) (see page 267)

[Troubleshooting: vCenter AIM Instance Status Icon Shows Multiple Instances](#) (see page 267)

[Troubleshooting: vCenter AIM Instance Status Icon Shows Error](#) (see page 268)

[Troubleshooting: vCenter AIM Instance Status Icon Shows Warning](#) (see page 269)

[Troubleshooting: vCenter AIM Instance Status Icon Shows Disabled](#) (see page 270)

## Troubleshooting: vCenter AIM Instance Status Icon Shows Discovery in Progress

### Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Discovery in Progress).

### Solution:

Wait until the discovery process of the vSphere environment has completed. The discovery duration depends on the amount of managed objects related to virtual and physical resources in vSphere. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job has completed, CA Virtual Assurance adds a vCenter Server folder to the resources tree. Then you can start managing vSphere and its entire virtual infrastructure.

## Troubleshooting: vCenter AIM Instance Status Icon Shows Multiple Instances

### Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Multiple AIMs manage this instance).

### Solution:

Verify that your CA Virtual Assurance manager manages each vCenter Server with one vCenter AIM instance only. If a CA Virtual Assurance manager manages a vCenter Server through multiple AIM instances, management problems would occur. CA Virtual Assurance stops monitoring the associated vCenter Server.

Decide which AIM instance you want to use to manage the vCenter Server and remove the other instances from the vCenter AIM Servers pane.

### Follow these steps:

1. Select the AIM instance you want to delete and click  (Delete).

The Delete Item dialog appears.

2. Click Yes.

Repeat these steps with other multiple instances until you have unique relationships between manager and AIM instance established.

## Troubleshooting: vCenter AIM Instance Status Icon Shows Error

### Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

### Solution:

The following procedures address the most common issues which can cause a connection failure to the vCenter AIM:

- Verify, if the vCenter AIM Server is accessible.
- Verify, if SystemEDGE is running. Start or restart SystemEDGE if necessary.

### To verify, if the vCenter AIM server system is accessible

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

```
ping servename
```

2. Verify the output of the commands to find out whether the vCenter AIM server has a valid DNS entry and IP address.

If the vCenter AIM server is not in the DNS, add the vCenter AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

If the vCenter Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress servename
```

Enter the correct IP address and vCenter AIM server name. For example:

```
192.168.50.51 myvCenterAIM
```

4. Click  (Validate) in the upper-right corner of the vCenter AIM Server pane.

If the error status remains unchanged, continue with the next procedure.

**To verify, if SystemEDGE is running**

1. Log in to the vCenter AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.

The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.

2. Start or restart SystemEDGE.

Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.

3. Change to the CA Virtual Assurance user interface, vCenter AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Virtual Assurance validates the vCenter AIM Server connection.

If the error status remains unchanged, verify whether the data you gathered according to the requirements for this scenario is still valid.

## Troubleshooting: vCenter AIM Instance Status Icon Shows Warning

**Symptom:**

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Warning). Other instances on the AIM are still in error state.

**Solution:**

No specific actions are required for the associated instance. The warning informs you that other instances on the same AIM are in error state. The warning disappears when the problems of the other instances are resolved. If one or more instances of an AIM are in error state, then all other instances in the AIM show warning state.

## Troubleshooting: vCenter AIM Instance Status Icon Shows Disabled

### Symptom:

After CA Virtual Assurance has discovered vCenter AIM instances in the network, the status icons of several instances show  (Disabled). This vCenter AIM instance is not managed.

This status appears, if CA Virtual Assurance has discovered a vCenter AIM with the following relationships:

- The vCenter AIM is configured for a vCenter Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a vCenter Server that has not been configured in the vCenter Servers pane.

### Solution:

To change the status of the AIM instance to ready, do one of the following:

- Add the missing vCenter Server connection to the CA Virtual Assurance manager.
- Edit the existing vCenter Server connection and change its managed status to enabled.

## Verify the vCenter Server Folder Appearance in the Resources Tree

After a successful configuration and discovery, the new vCenter Server is listed in the Resources Explore pane under the VMware vCenter Server folder.

### Follow these steps:

1. Click Resources, Explore.

The resources tree appears.

2. Expand VMware vCenter Server.

The managed vCenter Servers appear.

3. Expand the new vCenter Server entry.

The managed vSphere infrastructure appears: VMware Datacenters, ESX Servers, Resource Pools, VMs, ...

CA Virtual Assurance is now ready to manage the added vSphere environment with its virtual infrastructure.


## How to Use Policy Actions to Identify Performance Issues

This scenario provides information about how a system administrator can identify and dynamically address performance issues. This information is meant for System Administrator to optimize the allocation of resource shares of their managed vCenter environments.

The policy actions identify VM resources and dynamically adjust the allocation of CPU shares. *Shares* determine which VM gets resources when there is competition for resources among VM's. Using shares allows dynamic allocation of CPU resources. Each VM is allocated a specified number of shares. The amount of resources the VM receives depends on the proportion of its share against the total number of available shares. The allocation is dynamically changed based on the current usage of CPU resources on the ESX Server host.

If CPU usage of any VM is over 50 percent, allocation of CPU shares increases dynamically. If CPU usage is less than 30 percent, the CPU shares allocation decreases dynamically. The policy component not only identifies the problematic virtual machines but helps ensure dynamic actions which sustain business continuity. Using policy actions helps ensure that resources are allocated where to virtual machines that are in need and deallocate when the need is gone.

### How to use policy actions to identify performance issues



To identify and address performance issues using policy actions, follow these steps:

1. [Create an action for CPU metric.](#) (see page 272)
2. If CPU usage is more than 50 percent, [create a rule for CPU metric to increase allocation.](#) (see page 273)
3. If CPU usage is less than 30 percent, [create a rule for CPU metric to decrease allocation.](#) (see page 273)

## Create an Action for CPU Metric

Policy provides the creation of rules and actions that can be used to create policies for the automated management of systems. Custom actions can be created for actions not included in the default library.

### Follow these steps:

1. Click Resources.  
The Resources page appears.
2. Open the Explore pane.  
Available groups, services, and systems appear.
3. Click Policy, then click Actions.  
The Actions page appears.
4. Click '+' on upper right side bar to add new action.
5. Enter name of the Action.
6. Select Resource Configuration from Category drop-down list.
7. Select Configure Shares from Type drop-down list.
8. In VC Server field, leave the entry as "%VCServer%" to apply this action on any VM across any VC Server.
9. In VC Data Center field, leave the entry as "%DATACENTER%".
10. In the Target VM Machine field, leave the entry as "%VMNAME%".
11. Select "Set CPU" from Operations drop-down list and enter Values as "10000".  
The number is arbitrary and the share values are set to normal.  
**Note:** Use higher or lower numbers to increase and decrease the share allocations accordingly.
12. If the changes require approval, enable Help Desk Approval.  
A Message will appear in the Event Console after the Action is created.  
CAAP4521 Policy: Action <action name> was created.



## Create a Rule for CPU Metric to Increase Allocation

Creating a rule for CPU metric to increase CPU allocation helps ensure dynamic resource allocation when the usage exceeds the threshold.

**Follow these steps:**

1. Click on Resources tab, Policy, Rules.
2. Click '+' on upper right side bar to add new rule.
3. Enter Name of rule and click Next.
4. Select the action from the "Action Selection" list for the rule and click Next.
5. Enter the metric-based rule where CPU usage is greater than 50 percent to increase the CPU shares of the VM.

## Create a Rule for CPU Metric to Decrease Allocation

Create a Rule for CPU metric to decrease allocation

**Follow these steps:**

1. Click on Resources tab, Policy, Rules.
2. Click '+' on upper right side bar to add new rule.
3. Enter Name of rule.
4. Select the action from the "Action Selection" list for the rule.
5. Enter the rule-based on metric where CPU usage is less than 30 percent to decrease the CPU shares of the VM.

## Scalability Best Practices

This section provides best practices and recommendations for the deployment of CA Virtual Assurance. The purpose of the document is to assist with the planning of a roll out of CA Virtual Assurance within a production environment, with particular focus on:

- Monitoring and CA Virtual Assurance Management of VMware Environments
- Deployment of SystemEDGE and other Monitoring Software
- Initial and on-going configuration of SystemEDGE

**The following sections are included:**

1. [Remote Deployment and Policy Configuration Overview](#) (see page 274)
2. [Hardware Specifications](#) (see page 276)
3. [Database Considerations](#) (see page 276)
4. [Network Considerations](#) (see page 277)
5. [Scalability Recommendations and Limitations](#) (see page 277)
6. [Scenario Use Cases](#) (see page 283)


## Remote Deployment and Policy Configuration Overview

CA Virtual Assurance provides a comprehensive solution for remotely deploying the SystemEDGE agent to all managed systems. You can create deployment templates based on provided packages that contain customized installation parameters and simultaneously deploy these templates to numerous managed systems.

In addition CA Virtual Assurance provides a comprehensive solution for the ongoing configuration the SystemEDGE agents running on all managed systems. Policy Configuration allows a library of Policies to be created. These policies are applied to one or more systems running SystemEDGE and SRM AIM. When an agent managed by Policy Configuration is installed, it automatically requests a policy. As a result, the agent runs a controlled and consistent set of Policies. Each agent can then be updated individually, based on the Policy the agent is running, or the service the agent is a member of.

Remote Deployment and Policy Configuration share the Domain Server and Distribution Server technology. This technology provides a solution that is both scalable and able to be distributed across multiple data centers.

The diagram illustrates the basic architecture of the Remote Deployment and Policy Configuration components



## Hardware Specifications

This section lists the minimum hardware specifications for large-scale implementation of CA Virtual Assurance. For larger scale implementations, consider increasing the specification of the management servers.

- Domain Server: 2.6-GHz Dual-core Processor, 4-GB RAM, 100-GB disk.
  - Distribution Server: 1-GHz Single Core/Processor/Virtual Processor, 2-GB RAM, 100-GB disk, 100-Mb/sec Ethernet.
  - VC AIM Monitoring Server: Dual Core Processor: 2.6-GHz CPU, 4-GB RAM, 100-GB disk.
- Note:** 50 percent is the maximum usage allowed for CPU and memory.
- Target System: 1-GHz Single Core/Processor/Virtual Processor, 512-MB RAM, 2-GB disk, Single 100-Mb/sec Ethernet.

## Database Considerations

As the managed environment grows larger, more database activity can be expected. The product databases grow in size based on product usage, potentially consuming 30 GB or more, depending on the maintenance that is being done. We recommend the following general rule for data retention: for every 1000 machines in your monitored environment, increase the poll interval by 300 seconds.

**Note:** Using a dedicated standalone system for the database can improve its performance. Keep the database close to other Virtual Assurance for Infrastructure Managers components on the network, for example, on the same subnet, to improve response times.

## Network Considerations

When planning the roll out of CA Virtual Assurance, consider the quality of the network connections to decide where to locate management components. The following items influence the scalability and effectiveness of the solution:

- Network quality: Poor quality results in data loss, causing slow response, or failures.
- Network bandwidth: Lower bandwidth limits the rate at which data can be sent between the components.
- Network latency: Higher latency (delay) limits the rate of data transfer, in a similar way to low bandwidth.
- DNS: Badly configured DNS hinders the deployment and ongoing configuration of the monitoring agents.

We recommend using at least 100-Mb/sec Network Links between management components. Especially when using a remote DB. If the network speed is less than 100Mb/sec, consider introducing additional Distribution Servers collocated with the target systems.

## Scalability Recommendations

This section provides information about scalability recommendations and limitations.

Consider the following information:

- [Monitoring of VMware Environments](#) (see page 278)
- [CA Virtual Assurance Management of VMware Environments](#) (see page 279)
- [Remote Deployment and Policy Configuration Recommendations](#) (see page 281)
- [Domain Server Recommendations](#) (see page 283)
- [Distribution Server Recommendations](#) (see page 283)
- [Scalability Use Cases](#) (see page 283)

## vCenter AIM Monitoring Recommendations

The SystemEDGE agent provides a plug-in architecture through which it can load optional *application insight modules (AIMs)* when it initializes. AIMs are functional extensions to the SystemEDGE agent. For example, the vCenter AIM enables SystemEDGE to manage vSphere environments through VMware vCenter Servers.

The vCenter AIM (Application Insight Module) is a pluggable component implemented within the SystemEDGE framework. As such, the data that it publishes is available to multiple SNMP Managers. Products such as CA Virtual Assurance Manager, eHealth, and Spectrum IM can leverage this data.

Due to the possibility of this component being utilized outside of the CA Virtual Assurance Manager, scalability recommendations are specified separately. Two main scalability considerations are discussed:

### General Recommendation for vCenter AIM Monitoring

The general recommendation for scalability limitations for vCenter AIM Monitoring is as follows:

- Maximum Number of VMs is approximately  $240,000 / (x + 6)$
- Maximum Number of Objects is approximately  $2,000,000 / (x + 6)$

where x is the number of SNMP polls against AIM per hour.

## Scalability Limitations in Terms of Monitored Objects

In general, CPU usage is the primary concern when scalability limitations are considered. For vCenter AIM monitoring, consider the three main factors that influence the CPU usage:

- The dynamic nature of the vCenter Servers being monitored.

The level of activity of vCenter Servers influences the CPU consumption. The following scalability recommendations assume an average level of activity of the vCenter Server being monitored.

- The number of SNMP Managers, and polling intervals of those SNMP Managers.

Large numbers of SNMP Managers polling the vCenter AIM, or short poll intervals, result in increased CPU consumption.

- The ratio of object count in relation to VM count.

An *object* is any element of vSphere that is monitored by the vCenter AIM. For example, vCenter AIM monitors Datastores, Virtual Disks, Physical Network Interface Controllers, Virtual Switches, SCSI Controllers, ESX Host Hardware Sensors, and so on. The number of objects directly impacts CPU consumption. Due to the need to maintain the vCenter AIM cache and the additional overhead required for publishing this data. In real-world systems, there are typically from 6 through 11 times as many objects as there are Virtual Machines within a given vCenter.

Based on the preceding factors, a single SNMP Manager monitoring vCenter AIM with a 10-minute poll interval has a limitation of approximately 20000 VMs.

## Scalability Limitation in Terms of Monitored Servers

The vCenter AIM functions in a multiple vCenter Server environment. In fact, the framework of the vCenter AIM results in a slight reduction in CPU usage as the number of VMs per vCenter Server decreases. For example, CPU usage is lower against three vCenter Servers, each with 2000 VMs, than against a single vCenter Server with 6000 VMs.

With that noted, there are points at which the responsiveness of the vCenter AIM becomes the scalability limit. In general, vCenter AIM is able to monitor up to ten vCenter Servers.

## CA Virtual Assurance vCenter Management Recommendations

The VA-IM manager does far more than simply monitor and publish vCenter data. It has the responsibility for storing and managing historical data, performing active operations against the vCenter, running automation policy based upon vCenter data, reporting, and so on. As such, it often requires more resources than the vCenter AIM, which it uses as its main data collection mechanism.

## vCenter Management Limitations in Terms of Virtual Machines

Due to the fact that the vast majority of the CA Virtual Assurance manager resides within a single process space, Operating System limitations tend to be the primary culprit in scalability. Consider the following limiting factors:

- **Available Memory:** As the number of objects being managed increases, the amount of memory required to cache data and handle messaging increases rapidly. We recommend increasing the memory of manager to at least 8 GB.
- **Available CPU:** The CA Virtual Assurance manager requires significant CPU resources, especially in times of rapid environmental change, or during initial startup. We recommend supplying additional CPU (3.2 GHz or larger) for moderately large managed environments to improve the responsiveness of automated processes.
- **Operating System Limitations:** Much of the CA Virtual Assurance Manager resides within a single process space. As a result, the memory addressing space of a 32-bit operating system can become exhausted, even when system memory is not exhausted. To avoid this issue, we recommend using 64-bit processor and operating system for a moderately large managed environment.

### Examples:

The following examples provide requirements and scalability limit recommendations for the CA Virtual Assurance manager:

- **Minimum Requirements** (32-bit, 2.6-GHz CPU, 4-GB RAM, 100-GB disk)  
Scalability limit: 2500 Computer Systems (that is, VMs and ESX Hosts)
- **Recommended System** (64-bit processor and Operating System, 3.2-GHz CPU, 8-GB RAM, 100-GB disk)  
Scalability limit: 8000 Computer Systems (that is, VMs and ESX Hosts)



## Performance Considerations during Initial Discovery

Performing the initial discovery and database load of the vCenter environment you want to manage can be a time consuming process. During this process, the following take place:

1. The vCenter AIM parses the entire vCenter environment and publishes the results for the managers.
2. The CA Virtual Assurance manager retrieves the published data from the vCenter AIM and creates an internal cache for processing.
3. The internal cache is synchronized with the current CA Virtual Assurance manager database contents, with discoveries performed for Computer Systems not currently within the database.

During initial management of the vCenter server, the database has no Computer Systems in the database, so all of these objects are discovered and created. Consider the estimated time to complete initial discovery. Based upon baseline testing, the average throughput is: from eight to nine Computer Systems per minute.

### Examples:

The following examples provide the sizes of environments and corresponding estimated times to complete initial discovery:

- 2,500 Computer Systems - approximately five hours
- 8,000 Computer Systems - approximately 15 hours

During this initial population, CPU usage may be high for extended periods.

**Note:** The initial discovery process takes the vast majority of this time. However, the initial discovery is done once for the lifetime of the product. The vCenter AIM and CA Virtual Assurance internal caching processes, which are done whenever these processes are restarted, are considerably quicker. For example, 2500 Computer Systems are typically published through vCenter AIM and fully cached by the CA Virtual Assurance manager in approximately 5 minutes.

## Remote Deployment and Policy Configuration Recommendations

Consider the following aspects and recommendations to improve remote deployment and policy operations:

- Number of target machines  
For optimal performance, limit the deployment job size to 500 target machines in a batch.
- Number of Distribution Servers  
Deployment throughput is better when multiple distribution servers are used.

- Deployment package size

The smaller the deployment software package the better the throughput. The numbers recommended assume that all of the managed servers are required to have SystemEDGE and Advanced Encryption.

**Note:** A typical package size ranges from 10MB through 20MB.

- Quality and speed of the network

Low bandwidth, high packet loss and high latency between the Distribution server and the target affect the rate and reliability of Deployment and Configuration operations.

- The timescale for the rollout of Monitoring Software to the target systems

Stagger the deployment of monitoring software using collocated deployment servers. Staggering reduces the load on the network infrastructure. This load reduction can be achieved by creating a number of jobs, or by using the Staggering capability built into the solution.

If the monitoring software must be deployed over short time, we recommend deploying additional Distribution Servers in the environment.

- Frequency of agents reconfiguration (by Policy Configuration)

Reconfiguration of SystemEDGE Agents is expected to take around 30 seconds plus from 2 through 10 seconds per agent in a typical network environment. If agents must be reconfigured frequently, we recommend deploying additional Distribution Server in the environment.

- Geographical distribution of the target systems across multiple sites

Where the target systems are distributed across multiple sites, we recommend deploying a Distribution Server at each location. This recommendation is especially true if the remote data centers use a slow (< 100Mb/sec) or unreliable link. Deploying a local Distribution Server allows all Deployment and Configuration requests to be directed through the on-site Distribution Server. This limits the network traffic between the central and remote site.

- Communication ports

Remote Deployment and Policy Configuration rely on communications by CA-Messaging for Domain Server to Distribution Server and Distribution Server to Agent communications. CA-Messaging communicates over ports 4104(UDP) and 4105(TCP). For remote sites that are firewall protected, placing a Distribution Server on site allows all CA-Messaging communications to be set up as point-to-point.

**Note:** For agent discovery and ongoing monitoring, SNMP communications (typically port 161) are used. Open this port for direct communication from the managed systems to the manager.

- Management of Policy Configuration Policies and Templates

We recommend organizing the monitoring requirements into Templates based on the different workloads in your environment. Use a maximum of 1000 monitors per Policy or Template. Any number of templates can be applied to a system, but we recommend limiting the number of templates to 100 per system.

- Service Membership

To ease configuration operations, we recommend organizing servers into Services, with an upper limit of around 500 servers per Service. A server can be a member of multiple Services, and therefore we suggest creating services to represent different workloads. A template can be then applied directly to a Service.

## Domain Server Recommendations

*A deployment and configuration domain server (domain server) manages and controls all deployment and policy configuration operations.*

If the number of target systems exceeds 10000, we recommend running multiple instances of CA Virtual Assurance.

## Distribution Server Recommendations

*A deployment and configuration distribution (scalability) server helps ensure that deployment and policy configuration operations are carried out in an efficient and timely manner.*

Once the CA Virtual Assurance manager is installed, the next step in the rollout of Remote Deployment and Policy Configuration is to consider the number of Distribution Servers.

For Deployment operations, we recommend using one Distribution Server per 2000 target systems for a typical 100-Mb/sec network environment.

Where the SystemEDGE Agent is not being deployed using Remote Deployment (that is, only Policy Configuration is used), this number can reach 3000 systems per Distribution Server.

**Important!** We recommend performing a test deployment to at least one system for each distribution server before large deployment operations. We recommend deploying all possible packages using the distribution server to verify that there are no failures with the larger scale deployments.

## Scalability Use Cases

This section provides use cases that aim to represent typical production environments. We suggest comparing these cases with your own environment and following the recommendations of the closest one.

## Departmental Data Center

In this use case, monitoring is required for 1000 systems, all contained within one data center. All systems are in one location, within the firewall, and fast links exist between the computers.

Recommendations for this environment are:

- **Component Installation**

All CA Virtual Assurance manager components can be installed on the same system.

- **Initial Deployment**

Two jobs could be created to deploy the monitoring software to all target nodes. Depending on network load, the initial deployment of SystemEDGE (and Advanced Encryption if necessary), would be expected to complete in 8 to 12 hours.

Once deployment to remote systems has completed, the CA Virtual Assurance Manager discovers SystemEDGE. Policy Configuration delivers an initial policy to each agent. We expect the initial policy delivery to complete approximately eight hours after the completion of the job.

- **Service membership**

For ease of maintenance, it is suggested to align the monitored servers into Services, with approximately 200 servers per Service. Services could be aligned to business function, network topology, or other categorization as desired.

- **Applying Policies**

If necessary, applying policies to all monitored systems can be performed in a single operation.

## Multiple Data Centers

In this use case, 10000 systems, spread across multiple Data Centers, are being managed. The number of systems per data center varies from 500 through 2500. The Data Centers are geographically distributed across multiple locations. The links between the Data Centers including leased lines of less than 100Mb/sec. The systems run various workloads, and are managed by a number of different departments (application owners).

Recommendations for this environment are:

- **Component Installation**

Install the CA Virtual Assurance manager components on a dedicated server. This server must meet the minimum supported specification, but a quad-core server with at least 8-GB RAM is recommended.

We recommend installing the Database on a separate dedicated server with an increased specification of quad-core processor and 8-GB RAM.

To support the remote data centers, we recommend installing one Distribution Server in each data center. For a data center that contains more than 2000 servers, a second Distribution Server is recommended. Each Distribution Server can scale up to 3000 servers if only used for Policy Configuration.

- **Initial Deployment**

The following are relevant factors to consider while planning to deploy using multiple distribution servers:

- If possible, limit a single deployment job size to a maximum of 500 target systems.
- Verify that the nearest distribution server is chosen for the deployment operation.
- Although concurrent deployments are supported within a single distribution server, it is advised to only use concurrent deployments across multiple distribution servers. If you have 4 distribution servers, you could start a job to deploy to 500 machines for every distribution server.
- When you perform concurrent deployment using the same distribution server, verify that the second job does not deploy to the same set of target machines as the previous one.
- Where multiple packages are being delivered to many systems, consider splitting the package deliver to multiple jobs. For example, by deploying SystemEDGE and Advanced Encryption first.
- If you do see failures during a large-scale deployment, verify all prerequisites then use “resubmit job” to retry the deployment.
- To help with future deployments, we recommend saving the deployments as templates.

- Service membership

For ease of maintenance, we suggest aligning the monitored servers into Services, with a maximum of 500 servers per Service. For remote data centers, we recommend creating one or more Services (depending on data center size) to represent data center.

Where multiple departments use particular systems, the systems can be added to multiple services for ease of management by each department.

- Applying Policies

In this use case, the servers run various workloads. We therefore recommend that the base policy contains only the control settings. Hold your monitoring configuration in templates based on different monitoring requirements. These templates can then be applied to the required systems, either by manually selecting systems, or by applying templates to a service.

Splitting the monitoring requirements into templates allows the templates to be applied to the required systems, independently of each other. Select the systems manually, or apply templates to a service. We recommend applying templates in batches of 2000 to 2500 systems.

If the base policy is required to be changed, we recommend applying the policy to systems in batches of 2000 to 2500 systems.

**Note:** When templates are used, each delivery of a template or policy involves merging of all assigned templates with the base policy. The next step is delivering the resultant configuration to the agent. Therefore, where multiple templates are applied to a system, the time for delivery may be slightly increased.

## Large Environments

In this use case, approximately 21000 Agents are being managed. These agents are spread across three data centers, with each data center containing from 2000 through 10000 targets. The data centers are distributed, but have fast links between them. The managed systems are largely virtualized, run various workloads and are reprovisioned at times.

Recommendations for this environment would be:

- **Component Installation**

We strongly recommend running an instance of CA Virtual Assurance in each data center, such that each instance manages up to 10000 systems.

Within each data center, we recommend installing the CA Virtual Assurance manager components on a dedicated server. This server must meet the minimum supported specification, but we recommend an increased 8-GB RAM.

We suggest installing the Database on a separate, dedicated server with 8-GB RAM.

If a single instance of CA Virtual Assurance is required, we recommend using Manager and Database Servers with quad-core processors with 16-GB RAM.

To support Remote Deployment and Policy Configuration operations, we recommend installing a Distribution Server in each Data Center (one being installed on the Manager system).

- **Initial Deployment**

In this use case, we have one additional distribution server in the datacenter. Apart from that one difference in the setup, all the factors that were highlighted in the previous scenario apply equally to this scenario.

- **Service membership**

For ease of maintenance, we suggest splitting the monitored servers into multiple Services, with a maximum of 500 servers per Service.

- Applying Policies

We suggest limiting the base policy to control settings and 'base OS' monitors. Where different images are being used as the basis for virtual machines, a base policy can be created for each OS image. Verify SystemEDGE is configured to request this policy on registration.

For application-specific monitors, create templates based on individual monitoring requirements. To avoid index conflicts across templates, we suggest defining 'index ranges' up-front for each application. Alternatively, the base policy can be configured to "Automatically Resolve Indexes" under the 'Control settings' section.

Splitting the monitoring requirements into templates allows the templates to be applied to the required systems, independently of each other. You can either manually select systems, or apply templates to a service. We recommend applying templates in batches of 2000 - 2500 systems.

If the base policy is required to be changed, we recommend applying the policy to systems in batches of 2000 - 2500 systems.

**Note:** When templates are used, each delivery of a template or policy involves merging of all assigned templates with the base policy. The next step is delivering the resultant configuration to the agent. Therefore, where multiple templates are applied to a system, the time for delivery may be slightly increased.


**Important!** Contact CA Support if multiple instances of CA Virtual Assurance are deployed and if want to share the created policies between the different instances. CA Support can assist with the export and import of policies and templates between CA Virtual Assurance instances.



## How to Apply Policy and Layered Templates to Servers

From the CA Virtual Assurance user interface, you can control the SystemEDGE agent monitoring by creating a Base Policy and adding templates as layers to that policy. The diagram illustrates how to use Base Policy and Layered Templates:

### Apply Policy and Layered Templates to Servers



### Follow these steps:

[Layered Templates Concept](#) (see page 290)

[Create a Policy](#) (see page 291)

[Create Templates for Server Workload](#) (see page 301)

[Apply Policy and Templates to Servers and Verify Settings](#) (see page 315)

[\(Optional\) Manage the Base Policy and Templates for One or More Servers](#) (see page 317)


[\(Optional\) Update the Policy or Templates](#) (see page 318)

[\(Optional\) Apply Policy and Template Updates to Servers and Verify Updates](#) (see page 318)

## Layered Templates Concept

In an enterprise, the workload handled by a server or a server group varies. You can create multiple policies specific to the workload handled by a server or a server group. To assist in the creation of policies, templates are used to create application-specific monitors. The Base Policy and Layered Templates are combined to form a configuration file and applied to servers that you want to monitor. You can add or remove Layered Templates. Template updates can be applied directly to servers, without changing the base policy or reimporting the updated template into the Base Policy.

### Example: Apply Base Policy and Templates to Servers



You can use Layered Templates in the following scenarios:

**Disparate applications**

Create a library of templates for each server running a different set of applications. You can directly apply the template updates to each server.

**Dynamic environments**

The workload of the servers changes frequently in dynamic environments. You can use Layered Templates to segregate the monitors in logical groups. Based on the workload changes, you can directly apply the logical groups to systems or removed from systems.

**Shared servers**

In an enterprise setup, servers are shared across multiple departments. Each department manages and monitors applications on the shared server. You can use Layered Templates to independently manage and apply templates to systems of each department.

**Application maintenance**

You can split monitoring into multiple templates. In a server, you can remove a template for an application not in use, without affecting the monitoring of the remaining system.

## Create a Policy

Create a Base Policy to define a set of Monitors, MIB Extensions, Traps & Communities, and Control Settings to control the agent monitoring.

Common settings in Traps & Communities and Control Settings are available for policies only. If you use Layered Templates, common settings are specified in the Base Policy.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.

The Available Policies page appears.

2. Click + (New) on the Available Policies toolbar.

The New SystemEDGE Policy dialog appears.

3. Enter a name and an optional description for the policy, the system type and whether to base it on an existing policy and click Ok.

The policy is created, and a configuration screen appears in the right pane.

4. Click Save Policy.

The policy is created and saved.

**Note:** You can also use the existing default policy as a Base Policy, if necessary.

### Define SystemEDGE Policy Control Settings

You can control the following agent behavior using the SystemEDGE policy control settings:

- Security settings
- SNMP settings
- MIB table population
- UNIX settings
- Performance monitoring settings

You can segregate these common control settings from specific server workload configurations by adding them to the Base Policy.

You can apply the control settings defined in the policy to all systems you want to monitor with this configuration.

#### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Control Settings.

The Controls page appears.

4. (Optional) Click Use Defaults.

The default selections pane appears. You can change the default settings.

5. Configure the following control settings:

**SNMP**

Lets you define the following basic SNMP properties:

**Bind Address**

Specifies an interface to which the agent binds and listens for incoming SNMP requests. Valid addresses are IPv4 or IPv6 address.

**Note:** The corresponding default \_port is specified during installation.

**Bind Port**

Specifies the trap port the agent binds to for sending SNMP traps. If no bind\_address is specified, the agent binds to all available UDP addresses.

**Default:** Port selected by the system

**IP Family**

Specifies the agent communication method: IPv4 only, IPv6 only, or both. By default, the agent tries using IPv4 and then IPv6.

**FIPS Mode**

Specifies the agent to use FIPS-compliant encryption. Select Non-FIPS Mode to enable the CA eTrust Public Key Infrastructure libraries, and if this method fails, fall back to the internal minimum security solution. Select FIPS Co-existence Mode to enable FIPS-compliant encryption, and if this method fails, fall back to the CA eTrust Public Key Infrastructure Libraries. If they fail, select FIPS Only Mode to enable the RSA BSAFE Crypto-C Micro Edition FIPS-compliant libraries and perform no encryption.

**Default:** Non-FIPS Mode

**Trap Source**

Specifies the source address used to send traps. Valid addresses are IPv4, IPv6 address, or a host name.

**Default:** Host name of the agent

**Security Settings**

Lets you define the following security preferences:

**Authentication Traps**

Sends an authentication failure trap when the agent receives an SNMP message with a community name that the agent cannot recognize.

**Default:** Disabled

**Process Sets**

Permits access to processes and other software running on agent systems in the Process table and Running Software table. Allowing SNMP Sets on these tables can cause security issues.

### **Remote Shell Group**

Permits management systems to instruct the agent remotely to run shell scripts and programs on the agent system through the Remote Shell group. The disclosure of this type of information can post a potential security risk.

### **Execution Action**

Enables the execution of action commands with the monitoring tables when a threshold breach occurs. The capability to run action commands and scripts can be a security issue.

### **MIB Table Population**

Populates the following tables in the Systems Management MIB:

- Process Table
- User Group Table
- Who Table
- Trap Community Table
- Monitor Mirror Table
- Aggregate Mirror Table
- Top Processes Table

Each table either contains sensitive information that you can expose in a MIB or nonessential information that you can disable to save disk space. The default settings enable population of all tables except for the process table.

### **Miscellaneous**

Lets you define the following miscellaneous settings:

#### **Allow agent to be Updated using SNMP**

Permits agent updates using SNMP Sets (for example, removes write communities). If you permit SNMP Sets on the agent, any updates through this method cause a notification of an SNMP Set change. These updates also cause an exception when viewing policy details for the system.

### **Notify Manager of Configuration Updates**

Enables the agent to send a notification to the manager for any SNMP Set request that the agent processes.

### **Warm Start Discovery**

Enables an agent rediscovery of all devices after every warm start configuration update. If you manage a system with many devices, a discovery after every warm start can consume too much time and too many resources.

### **Use Perl Compatible Regular Expressions**

Perl Compatible Regular Expressions (PCRE) enables you to specify i18n compatible regular expressions while defining monitors that support regular expressions. The examples of regular expressions are log file, process, process group, Windows services and Windows events. You can also use this option to create more complex regular expressions. This option is provided in SystemEDGE agent 5.1.0 and above versions.

### **Automatically Resolve Index Conflicts**

Enables you to resolve Index conflicts. When you apply the layered templates to all systems, indexes are assigned to the monitors added in the template. If the assigned indexes conflict with existing indexes either within the base policy or another template, this option reassigns unique index values.

**Note:** Indexes contained within the base policy are always maintained in the delivered configuration. If this option is disabled, you cannot resolve conflicting indexes. However, when you apply layered templates to the systems, the conflicting indexes are displayed as errors on the layered templates that caused the conflicting indexes.

### **Historical Performance Monitoring**

Lets you define the following settings for the Performance Cube AIM, which collects history information into Systems Performance cubes for historical performance management:

#### **Collection Interval**

Specifies how often to collect information from the History table into performance cubes.

### **Index Range Start**

Specifies the beginning of the reserved range of indexes, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

### **Index Range End**

Specifies the end of the reserved range of indexes, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

### **UNIX Control Settings**

Lets you define the following settings for agents running on UNIX systems:

#### **Sub-program Group**

Specifies a group name other than root under which to run subprograms.

#### **Sub-program User**

Specifies a user name other than root under which to run subprograms.

#### **Linux Freemem Include**

Specifies whether to include system buffers, disk cached memory, or both in free memory calculation.

### **Query System Devices**

Lets you enable querying of the following system device metrics:

- Serial device status
- Floppy disk status
- Disk size, capacity, description, and other properties (Probe Disks)
- NFS file system status
- HP-UX graphics status

Querying these metrics can cause issues with potential agent blocking. The default settings enable querying of only serial device status and NFS file system status.

6. Click Plugins.

The Plugins pane appears. This pane controls which AIMs to load with the agent.

7. Do one of the following:

- Select 'Load all available plugins' to load all AIMs available on the agent system.
- Select 'Load plugins selected in the table'.
- Click + (New) on the External Plugins toolbar to add an AIM to the External Plugins table.



**Note:** For more information about available AIMs, see the *SystemEDGE User Guide*.

AIM loading is configured.

8. Click Aggregate Monitors.

Configure aggregate monitors as described in [Configure Object Aggregation](#) (see page 297).

The control settings are defined.

9. Click Save Policy.

The policy is saved.

**More Information:**

[Configure Object Aggregation](#) (see page 108)

## Configure Object Aggregation

By default, SystemEDGE aggregates monitors into a managed object that contain the same values for the object class, instance, and attribute properties. For example, all monitors with a class of SysHealth, an instance of CPU, and an attribute of SysTime are combined into an aggregate managed object.

You can configure the agent to aggregate objects on higher levels when defining SystemEDGE policy. You can also configure other aspects of agent behavior related to object aggregation and the state management model.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Control Settings.

The Controls page appears.

4. Click Aggregate Monitors.

The Aggregate Monitors page appears.

5. Select one or more of the check boxes to specify aggregation levels.

These represent higher aggregation levels than the default, up to aggregating all monitors into one top-level agent object. Specifying aggregation levels lets you create a tiered object architecture that propagates status up to the level you specify.

6. Configure the following additional settings, and click Save Policy:

### **Send legacy traps for all aggregated monitors**

Specifies whether to send legacy traps for all monitors that make up a managed object. By default, the agent only sends a state change trap for the monitor with the highest severity, even if other monitors in the object experience threshold breaches.

### **Execute commands of all aggregated monitors**

Specifies whether to execute action commands for all monitors that make up a managed object. By default, the agent only runs an action command for the monitor with the highest severity, even if other monitors in the object experience threshold breaches.

Aggregation settings are configured. Apply or reapply the policy for the changes to take effect.

### **More Information:**

[Define SystemEDGE Policy Control Settings](#) (see page 158)

## Define Traps and Communities

SNMP settings define the communities that the agent uses and the destinations to which it sends traps.

### **Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.  
The Available Policies page appears.
2. Select the policy in the Available Policies table.  
The Summary page for the policy appears.
3. Click the Traps and Communities tab.  
The Communities page appears.
4. Select one of the following and click Action, select Apply:
  - Include only Server specific SNMP settings
  - Include Server specific SNMP settings and all Default settings
  - Select 'Include Server specific SNMP settings and selected Default settings

The SNMP settings are updated and the community page in the Communities table displays the following:

**Name**

Specifies the name of the community string.

**Port**

Specifies the port of SNMP.

**SNMP Version**

Specifies the SNMP version that the community uses.

**Access Rights**

Specifies whether the community should have read write or read-only permissions.

**Note:** Add at least one read-only and one read write community.

**Community/User**

Specifies the community name.

**Authentication Protocol**

Specifies the protocol to authenticate SNMPv3 data.

**Privacy Protocol**

Specifies the protocol to authenticate SNMPv3 data.

**Access Control List**

Specifies a space separated list of IP addresses to restrict community usage to those addresses only. If you leave the list blank, the agent grants access to any system that uses the associated community name. Access lists are only for communities that use SNMPv1.

**Note:** For information about defining SNMPv2c and SNMPv3 access lists, see the *SystemEDGE User Guide*.

5. (Optional) Add, update, or delete other communities as necessary.

6. Click Save Policy.

The policy is saved.

7. Click Trap Destinations.

The Trap Destinations page appears.

8. Define a trap destination using the following controls and click Add:

**Trap Type**

Specifies the type of trap to send, depending on the SNMP version.

**Destination**

Specifies the IPv4 or IPv6 address to which to send traps.

**Port**

Specifies the UDP port to which to send traps.

**Community**

Specifies the community name sent with the traps.

**Encoding**

(Optional) Specifies how to include the source address you defined in the Trap Source field of the Control Settings pane in traps. This parameter is important if the trap source translates to an IPv6 address. Enter the encoding parameter in a three digit format XYZ, assuming leading zeros.

**Default:** 000

**X**

Controls extending the four byte IPv4 source address field (SNMPv1 traps only). Enter 0 to not extend the source address field to include the 16 byte IPv6 address, and enter 1 to extend the source address field.

**Y,Z**

Controls the inclusion of source information into the trap's varbind (Y) or UDP packet (Z; SNMPv1 traps only). Enter one of the following for these digits:

**0:** Do not modify the trap's varbind or the outer UDP packet.

**1:** Include the trap\_source parameter as is in the varbind or packet (IPv4/IPv6 address or host name).

**2:** Include the trap\_source parameter preferably as an IPv4 address (then IPv6 address, then host name).

**3:** Include the trap\_source parameter preferably as an IPv6 address (then IPv4 address, then host name).

**4:** Include the trap\_source parameter preferably as a host name (then IPv4, then IPv6).

**5:** Follow the preference for 2 and include the host name.

**6:** Follow the preference for 3 and include the host name.

**7:** Follow the preference for 1 and include the host name (if trap\_source is an IPv6 address).

### Trap Source

(Optional) Specifies the IPv4 or IPv6 address or the host name to use as trap source.

**Default:** Global Trap

The trap destination appears in the Defined Trap Destinations table.

9. (Optional) Add, update, or delete other trap destinations as necessary.
10. Click Save Policy.

The policy is saved.

**Note:** For more information, see the *SystemEDGE User Guide*.

## Create Templates for Server Workload

Create templates that are specific to the workload of a server. You can specify Monitors and MIB Extensions.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates, and click SystemEDGE.

The Template List page appears.

2. Click + (New) on the Template List toolbar.

The New SystemEDGE Monitoring Template dialog appears.

3. Enter a name and an optional description for the template, the system type, and whether to base it on an existing template, and click Ok.

The template is created, and the Summary page appears.

4. A template is a collection of monitors and MIB extensions. To add monitors to the template, see the section [Add Monitors to the Template or the Policy](#) (see page 302). To add MIB extensions to the template, see section [Define MIB Extensions](#) (see page 313).

5. Click Save Template.

The template is created and saved.

## Add Monitors to a Template or the Policy

Add monitors to the template that are specific to the workload handled by a server or a server group. The following procedure is similar for adding monitors to a policy.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click Monitors and select the monitor you want to add.

To create monitors, you define the settings, which specify the threshold and severity values for the following monitors:

- [Create a Threshold Monitor](#) (see page 302)
- [Create a Process Monitor](#) (see page 305)
- [Create a Log File Monitor](#) (see page 307)
- [Create a Windows Event Monitor](#) (see page 308)
- [Create a History Monitor](#) (see page 310)
- [Create a Process Group Monitor](#) (see page 311)

4. (Optional) Repeat the process for any additional monitors.

5. Click Save.

The monitor is loaded to the policy or the template.

## Create a Threshold Monitor

Create a threshold monitor that lets the agent monitor the servers or the server groups against specified thresholds. The agent sends a trap when thresholds are breached.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

The Summary page appears with a list of monitors managed by the policy.

4. Click Threshold.  
The Threshold Monitors page appears.
5. Click + (New) on the Threshold Monitors toolbar.  
The Threshold Monitor Details: New dialog appears.
6. Configure the following threshold settings:

**Index**

Defines the table index that you want to use.

**Platform**

Specifies the platform.

**Description**

Defines an optional description.

**Object Class**

Specifies the object class to monitor. The values refer to the available MIB tables.

**Object Class Name**

Defines the object class name to use for the object state model. Value is an arbitrary string, for example, FileSystems.

**Object Attribute**

Specifies the object attribute to monitor. The values refer to the available attributes of the table selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this threshold monitor.

**Object Attribute Name**

Defines the object attribute name to use for the object state model. This is an arbitrary string, for example, PercentUsed.

**Object Instance**

Specifies the object instance to monitor. This value, for example, .3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this threshold monitor. For some object classes, the name of the instance itself can be given (for example, C: instead of .3, or /var for a Unix machine).

**Object Instance Name**

Defines the object instance name to use for the object state model. Value is an arbitrary string, for example, SysVol\_C.

**Interval**

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

**Severity**

Specifies the severity to use for the object state model.

**Operator**

Specifies the operator to use.

**Value**

Defines the value to use.

**Sample Type**

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

**State**

Specifies if the monitor maintenance entry is active or inactive.

**Start Time**

Defines the start time when the monitor is switched off and the maintenance window begins.

**Stop Time**

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Threshold Monitor settings are saved.

8. Click Save Template.

The Threshold Monitor is loaded to the template.



## Create a Process Monitor

Create a process monitor that lets the agent monitor a process, service, or process table objects against specified thresholds. The agent sends a trap when thresholds are breached or the state of a process (running or stopped) changes.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

The Summary page appears with a list of monitors managed by the policy.

4. Click Process.

The Process Monitors page appears.

5. Click + (New) on the Process Monitors toolbar.

The Process Monitor Details: New dialog appears.

6. Configure the following process settings:

#### Index

Defines the table index that you want to use.

#### Platform

Specifies the platform.

#### Description

Defines an optional description.

#### Object Class Name

Specifies the object class name to use for the object state model. Value is an arbitrary string, for example, Process.

#### Object Attribute

Specifies the object attribute to monitor. The values define the available attributes for process monitoring.

#### Object Attribute Name

Defines the object attribute name to use for the object state model. Value is an arbitrary string, for example, MemUsedPercent.

### **Object Instance**

Specifies the object instance to monitor. This is the regular expression (dependent from optional settings) to use for matching processes by name, or Windows services by name. Pattern should uniquely match a single process (service). Arguments can be included (see optional settings).

### **Object Instance Name**

Specifies the object instance name to use for the object state model. Value is an arbitrary string, for example, ApacheServer.

### **Interval**

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

### **Severity**

Specifies the severity to use for the object state model.

### **Operator**

Specifies the operator to use.

### **Value**

Defines the value to use.

### **Sample Type**

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

### **State**

Specifies if the monitor maintenance entry is active or inactive.

### **Start Time**

Defines the start time when the monitor is switched off and the maintenance window begins.

### **Stop Time**

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save  
The Process Monitor settings are saved.
8. Click Save Template.  
The Process Monitor is loaded to the Policy.

## Create a Log File Monitor

Create a log file monitor that lets the agent monitor any UTF-8 encoded system or application log file by searching for strings specified as regular expressions. The agent sends a trap when a match occurs.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

The Summary page appears with a list of monitors managed by the policy.

4. Click Log File.

The Log File Monitors page appears.

5. Click + (New) on the Log File Monitors toolbar.

The Log File Monitor Details: New dialog appears.

6. Configure the following process settings:

#### Index

Defines the table index that you want to use.

#### Monitor Type

Specifies the monitor type that you want to use.

#### Platform

Specifies the platform.

#### Description

Defines an optional description.

#### Log File/Directory Name

Defines the path to the file or the directory to monitor.

#### Search Filter

Specifies the search filter.

#### Interval

Defines the evaluation interval for the monitor in minutes

#### Severity

Specifies the significance of the monitor on a match.

The Maintenance Window page lets you define the following settings:

**State**

Specifies if the monitor maintenance entry is active or inactive.

**Start Time**

Defines the start time when the monitor is switched off and the maintenance window begins.

**Stop Time**

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings page lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Log File Monitor settings are saved.

8. Click Save Template.

The Log File Monitor is loaded to the Policy.

## Create a Windows Event Monitor

Create a windows event monitor that lets the agent monitor the Windows event log entries using different filters (event source). The agent sends a trap when a match occurs.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

The Summary page appears with a list of monitors managed by the policy.

4. Click Windows Event.

The Windows Event Monitors page appears.

5. Click + (New) on the Windows Event Monitors toolbar.

The Windows Event Details: New dialog appears.

6. Configure the following process settings:

**Index**

Defines the table index that you want to use.

**Platform**

Specifies the platform.

**Description**

Defines an optional description.

**Event Log**

Specifies the event log to read.

**Event Type**

Specifies the event type to match.

**Source Filter**

Defines the source filter to use.

**Description Filter**

Defines the description filter to use.

**Severity**

Specifies the significance of the monitor on a match.

The Maintenance Window subtab lets you define the following settings:

**State**

Specifies if the monitor maintenance entry is active or inactive.

**Start Time**

Defines the start time when the monitor is switched off and the maintenance window begins.

**Stop Time**

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Windows Event Monitor settings are saved.

8. Click Save Template.

The Windows Event Monitor is loaded to the Policy.

## Create a History Monitor

Create a history monitor that lets the agent provide the historical data collection for manager-side baseline and trend analysis. The agent uses the metrics to provide a picture of average system performance during a specific time interval.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

The Summary page appears with a list of monitors managed by the policy.

4. Click History.

The History Monitors page appears.

5. Click + (New) on the History Monitors toolbar.

The Historical Details: New dialog appears.

6. Configure the following process settings:

#### Index

Defines the table index that you want to use.

#### Platform

Specifies the platform.

#### Description

Defines an optional description.

#### Object Class

Specifies the object to monitor. The values refer to the available MIB table values.

#### Object Attribute

Specifies the object attribute to monitor. The values refer to the available attributes of the table selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this History entry.

#### Object Instance

Defines the object instance to monitor. This value (for example, 0.3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this History entry.

**Interval**

Defines the collection interval in a multiple of 30 seconds.

**Buckets**

Defines the number of samples to collect.

**Add to Performance Cube check box**

Specifies whether to collect performance cube data for this entry.

The Maintenance Window page lets you define the following settings:

**State**

Specifies if the monitor maintenance entry is active or inactive

**Start Time**

Defines the start time when the monitor is switched off and the maintenance window begins.

**Stop Time**

Defines the stop time when the monitor is switched on again and the maintenance window ends.

7. Click Save  
The History Monitor settings are saved.
8. Click Save Template.  
The History Monitor is loaded to the Policy.

## Create a Process Group Monitor

Create a process group monitor that lets the agent define a group of processes and monitors that group for changes. If the process group changes, the agent sends a trap.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.  
The Template List page appears.
2. Select the template in the Template List.  
The Summary page for the template appears.
3. Click the Monitors tab.  
The Summary page appears with a list of monitors managed by the policy.
4. Click Process Group.  
The History Monitors page appears.

5. Click + (New) on the Process Group Monitors toolbar.

The Process Group Details: New dialog appears.

6. Configure the following process settings:

**Index**

Defines the table index that you want to use.

**Platform**

Specifies the platform.

**Description**

Defines an optional description.

**Process Name**

Defines the process name. This is the regular expression (dependent from optional settings) to use for matching processes by name.

**Interval**

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

**User Name**

Defines the user name to match in addition to any process name regular expression.

**Group Name**

Defines the group name to match in addition to any process name regular expression.

**Severity**

Specifies the significance of the monitor on a group change

The Maintenance Window page lets you define the following settings:

**State**

Specifies if the monitor maintenance entry is active or inactive

**Start Time**

Defines the start time when the monitor is switched off and the maintenance window begins.

**Stop Time**

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings page lets you define the flags that can be used for the different monitor entries or history control entries.



7. Click Save  
The Process Group Monitor settings are saved.
8. Click Save Template.  
The Process Group Monitor is loaded to the Policy.

## Define MIB Extensions

Defining MIB extensions provide functional benefits that are not available in local file manipulation. The policy configuration feature provides field names and the list of key properties such as object type.

When you configure a policy or a monitoring template, click the MIB Extensions tab to add the following objects:

- MIB Extensions
- Windows Performance
- Windows Registry

**Note:** To add MIB Extensions to a template or a policy, see [Add MIB Extensions to a Template or a Policy](#) (see page 313). MIB Extensions within templates are supported for the purposes of applying the MIB Extensions directly to monitored systems. MIB Extensions for use within policies should be created directly in the Policy itself.

## Add MIB Extensions to a Template or a Policy

Define MIB extensions for a template or policy using the policy configuration feature.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, and expand Monitoring Templates or Policies.
2. From the Templates List or Available Policies page, click the template or policy name.  
The Summary page appears.
3. Click the MIB Extensions tab.  
The MIB Extensions page appears.
4. Define the MIB extension attribute using the following controls and click Add:

#### Index

Defines the attribute leaf number.

#### Type

Specifies the attribute type.

**Extension Command**

Defines the full path or the name (including parameters) of the script or binary to execute.

**Access Rights**

Specifies the attributes access rights.

5. Click the Windows Performance tab.

The Windows Performance pane appears.

6. Define the Windows Performance attributes using the following controls and click Add:

**Index**

Defines the attribute leaf number.

**Type**

Specifies the attribute type.

**Object**

Specifies the performance registry object.

**Counter**

Specifies the performance registry counter.

**Instance**

Defines the performance registry instance.

7. Click the Windows Registry tab.

The Windows Registry pane appears.

8. Define Windows Registry attribute using the following controls and click Add:

**Index**

Defines the attribute leaf number.

**Type**

Specifies the attribute type.

**Key**

Defines the registry key in HKEY\_LOCAL\_MACHINE.

**Value**

Defines the attribute value.

**Note:** For more information, see the *SystemEDGE User Guide*.

9. Click Save Template or Policy.

The configuration is saved.

## Delete Monitors from Templates or a Policy

You can delete a monitor from a policy or template.

**Follow these steps:**

1. Click the Resources tab, open the Configure pane, and expand Monitoring Templates or Policies.
2. From the Templates List or Available Policies page, click the template or policy name.

The Summary page appears.

3. Click the Monitors tab.

The Summary page appears with a list of monitors managed by the policy.

4. Click the appropriate monitor tab and select the monitor you want to delete.
5. Click Action and select Delete.

A warning message appears.

6. Click Ok to confirm the deletion.
7. (Optional) Repeat the process for any additional monitors.
8. Click Save Policy.

The monitor is deleted from the Policy.

**Note:** You cannot delete a template, or the policy with the template which is used by a server or a server group.

## Apply Policy and Templates to Servers and Verify Settings

After you create the template, you can apply the policy, with the template directly to the servers or the server groups across the enterprise.

**Follow these steps:**

1. Select the policy in the Available Policies table or select a template from the Template List.

The Summary page for the policy or template appears.

2. Select the Managed Machines tab.

The list of managed machines appears.

3. Click Action and select Apply.

Tabs appear for selecting systems on which to apply the policy. The 'Update machines running this policy/template' tab lets you apply the policy to systems that are already running the policy or template. The 'Apply to Machines not running this policy/templates' tab lets you apply policy or a template to systems.

4. (Optional) Do one of the following from the 'Update machines running this policy/template' tab:

- Select 'Update all machines using this policy/template' to deploy the policy or template on all systems currently running it. This option is useful if you have made configuration changes that you want to apply globally.
- Select 'Update selected groups of machines' to update only systems that meet any of the following criteria:

- Running an out of date version of the policy or template
- Exceptions have occurred
- Running a current version of the policy

Select any of these options. Exceptions occur when a user applies a point configuration change to an agent that is not represented in the applied policy or template.

- Select Advanced and manually select systems in the Select Machines pane to which to reapply the policy or template.

5. (Optional) Select systems from the 'Apply to Machines not running this policy/template' tab to which to apply the policy or template.

6. Click Apply Policy or Apply Template.

The application is initiated.

7. Verify if the servers behave as expected. If necessary, you can update and apply the updated policies and templates.

## (Optional) Manage the Base Policy and Templates for One or More Servers

Manage the templates and the base policy for a single or multiple servers. You can replace the current base policy, add templates, or remove templates.

### Follow these steps:

1. Click the Resources tab, open the Explore pane, and select the server for which you want to change the policy configuration.

The resources page for the server appears.

2. Select Monitoring Software, Policies.

The table displays the list of policies and templates applied to the server.

3. Click  (Modify Policy) to replace the current base policy for this server by another available base policy.

The Modify Policy dialog appears listing all available base policies.

4. Select the appropriate policy and click Apply.

The new base policy for the selected server has been applied. The status of the policy changes from Delivery Requested, Delivered, to Configured.

5. Click  (Modify Template) to add or remove templates from the configuration of the selected server.

The Modify Templates dialog appears listing the available templates in the left pane and the applied templates in the right pane.

6. Select the templates that you want to add or remove, use the arrows to make your assignments, and click Apply.

The new set of templates has been applied to the configuration. The status of the templates change from Delivery Requested, Delivered, to Configured.

The new configuration has been applied.

You can also manage multiple servers as a group.

### Follow these steps:

1. Create a service at the datacenter level that specifies the group of servers.

The new service appears in the Explore pane.

2. Select the service.

The service page appears.

3. Select Monitoring Software, Policies.

The table displays the list of policies and templates applied to the servers.

The following steps are identical to the procedure for single servers.

4. Complete the configuration.

## (Optional) Update the Policy or Templates

If necessary, you can update the existing policy or templates by adding or deleting monitors from the policy or the template. The update procedures are similar to the creation process.

### Follow these steps:

1. Add or delete monitors that are specific to the server workload. To add monitors to the template or the Policy, see [Add Monitors to the Template or the Policy](#) (see page 302). To delete monitors from the Policy, see [Delete Monitors from Templates or a Policy](#) (see page 315).
2. [Define MIB Extensions](#) (see page 313).
3. [Define SystemEDGE Policy Control Settings](#) (see page 292).

The policy or templates are updated.

## (Optional) Apply Policy and Template Updates to Servers and Verify Updates

After you update the template, apply the template updates directly to servers or server groups across the enterprise.

### Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates, and then select SystemEDGE.

The Summary page displays a list of the SystemEDGE Monitoring Templates.

2. Select the Template Name.

The Summary page appears with the Template information.

3. Click Action and select Apply.

Tabs appear for selecting machines on which to apply the monitoring template. The 'Update machines running this template' tab lets you apply the monitoring template to machines that are already using the template. The 'Apply to Machines not running this template' tab lets you apply the monitoring template to machines without using any template.

4. (Optional) Under Existing Machines, select machines from the 'Update machines running this template' tab options.
5. (Optional) Under Selected Machines, select the machines to which the template is re-applied.
6. (Optional) Select machines from the 'Apply to the Machines not running this Template' tab to apply the template.
7. Click Apply.

The template application is initiated and the view Status link appears.

8. Click View Status link to verify whether the SystemEDGE monitoring template updates are applied to servers.

The page appears with the list of servers to which the SystemEDGE monitoring template updates are applied.

The Layered Template Updates have successfully been applied to the servers or the server groups.

9. Verify if the servers behave as expected. If necessary, you can update and apply the updated policies and templates again.





# Appendix A: FIPS 140-2 Encryption

---

This section contains the following topics:

[FIPS Overview](#) (see page 321)

## FIPS Overview

The Federal Information Processing Standards (FIPS) 140-2 publication is a security standard for the cryptographic libraries and algorithms a product should use for encryption. FIPS 140-2 encryption affects the communication of all sensitive data between components of CA products and between CA products and third-party products. FIPS 140-2 specifies the requirements for using cryptographic algorithms within a security system protecting sensitive, unclassified data.

CA Virtual Assurance uses the Advanced Encryption Standard (AES) adapted by the US government. CA Virtual Assurance incorporates the RSA Crypto-J v3.5 and Crypto-C ME v2.0 cryptographic libraries, which have been validated as meeting the FIPS 140-2 Security Requirements for Cryptographic Modules.



# Appendix B: Troubleshooting

---

This section contains the following topics:

[Adjusting Poll Interval Settings for Solaris Zones Environments](#) (see page 323)  
[Attributes Show a Value of Zero](#) (see page 324)  
[Browsers Do Not Display Consecutive Spaces in Events](#) (see page 324)  
[CA NSM Fails to Receive Traps on Windows](#) (see page 324)  
[Cisco UCS Folder Does Not Display in UI](#) (see page 325)  
[DB Transaction Log Sizes Increase Unexpectedly](#) (see page 325)  
[Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero](#) (see page 326)  
[Empty Task ID for the dpmvc virtualswitch Command](#) (see page 327)  
[Local and Remote Monitors Do Not Show the Same Values](#) (see page 327)  
[Remote Deployment to Solaris Lists SPARC and x86 Systems](#) (see page 328)  
[PMM Stops Polling an AIM](#) (see page 328)  
[Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear](#) (see page 329)  
[Resetting the vCenter Server Password Causes Data Collection to Fail](#) (see page 329)  
[Solaris Zones AIM Reset if a Monitored System is Down](#) (see page 329)  
[Unable to Connect to Microsoft SQL Server](#) (see page 330)  
[Upgrading SystemEDGE](#) (see page 330)  
[User Interface is Unresponsive on Provisioning and Policy Screens](#) (see page 330)  
[vCenter Server AIM Attributes Show Zero](#) (see page 331)  
[vCenter Server Connection Failed](#) (see page 331)  
[vCenter AIM Instance Status Icon Shows Disabled](#) (see page 333)  
[vCenter AIM Instance Status Icon Shows Discovery in Progress](#) (see page 333)  
[vCenter AIM Instance Status Icon Shows Error](#) (see page 334)  
[vCenter AIM Instance Status Icon Shows Multiple Instances](#) (see page 335)  
[vCenter AIM Instance Status Icon Shows Warning](#) (see page 336)  
[VM Usage Values Do Not Update Immediately After Power Down](#) (see page 336)  
[Blank Query Results Tab after Upgrade](#) (see page 337)

## Adjusting Poll Interval Settings for Solaris Zones Environments

### Symptom:

I do not know how to adjust poll interval settings for Solaris Zones environments.

### Solution:

Increase the poll interval of the Solaris Zones AIM if the number of systems and zones increases. For example, if the host and zone count is greater than 100, set the default poll interval to 240.

## Attributes Show a Value of Zero

**Symptom:**

Attributes show a value of zero.

**Solution:**

SystemEDGE rounds values down to zero, if they are smaller than one.

**Note:** The zoneAimStatHostDiskSvc MIB attribute always shows a value of zero.

## Browsers Do Not Display Consecutive Spaces in Events

**Symptom:**

Browsers do not display more than one consecutive space character in event descriptions.

**Solution:**

Browsers do not display more than one consecutive space, because additional spaces are truncated according to the HTML specification. Use caution when cutting and pasting events from the browser into rules as the event descriptions can differ.

## CA NSM Fails to Receive Traps on Windows

**Symptom:**

After a few days of receiving traps from CA Virtual Assurance, the NSM Event Console stops displaying traps in the Event Console Log. The issue is in the Windows SNMP Trap Service. WinSNMP applications stop responding when you run third-party security scanning software program on a Windows Server 2003-based computer.

**Solution:**

To resolve this issue, see the KB article on the Microsoft Support website <http://support.microsoft.com/kb/931565>.

## Cisco UCS Folder Does Not Display in UI

**Symptom:**

After the product installation with Cisco UCS services configured, the Cisco UCS folder does not appear in the user interface.

**Solution:**

Open Services on the server where the UCS AIM is configured, and verify that SystemEDGE is running; if the SystemEDGE service is stopped, restart it. Start nodecftutil.exe to verify access information for the UCS Manager node. Use a MIB Browser to verify data polling from UCS Manager. If UCS access information is not populated, review the sysedge log for additional information.

## DB Transaction Log Sizes Increase Unexpectedly

**Symptom:**

In data centers with numerous managed objects, configuration changes, and metrics data collection activities, the Management DB and Performance DB transaction logs can increase unexpectedly. This issue can cause disk space to become low in environments with limited resources.

**Solution:**

To resolve this issue, see the KB article on the Microsoft Support website <http://support.microsoft.com/kb/873235>

The transaction log files, aom2.ldf and dpm.ldf, are located in the directory C:\Program Files\Microsoft SQL Server\...\MSSQL\Data in default Microsoft SQL Server installations.

**Note:** If the database log file is reduced in size, restart the Apache service to improve performance.

## Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero

**Symptom:**

Some Solaris Zones AIM MIB values always show N/A or zero.

**Solution:**

These MIB attributes of the Solaris Zones AIM are deprecated and remain for backward compatibility. The deprecated MIB attributes are:

- zoneAimStatHostDiskMode
- zoneAimStatProcessorSetContainerList
- zoneAimStatProcessorSetResourcepoolId
- zoneAimStatProcessorSetResourcePoolIdList
- zoneAimStatProcessorSetResourcepoolName
- zoneAimStatProjectFSSEnabled
- zoneAimStatResourcePoolContainerList

## Empty Task ID for the dpmvc virtualswitch Command

**Symptom:**

When I run the dpmvc virtualswitch command, the result shows an empty task ID.

**Solution:**

This operation does not run asynchronously, and the result gets back immediately. However, the PMM treats the operation as a tasked operation. Therefore the response contains a task ID, but it is always an empty string ("").

For example, when you run the following command from the CLI, you get an empty task ID:

```
dpmvc virtualswitch -vs_add -vc_server MYVC5 -switch_name XYZ  
-esx_host_name MYESX -ws_user admin -ws_password ca_admin
```

CLI output:

```
...  
SC URL: https://VASManager/aip/sc  
VC URL: https://VASManager:443/aip/vc  
Task ID:  
Command execution successful
```

Other commands like dpmvc faulttolerance or dpmvc distributedswitch run asynchronously and you get a task ID.

## Local and Remote Monitors Do Not Show the Same Values

**Symptom:**

Local and remote monitors do not show the same values for the same attributes at the same time.

**Solution:**

For seamless local and remote monitoring, the monitored object names have been chosen identical. However, the APIs used are different and can return different values.

SystemEDGE on a remote machine runs independently from the RM AIM on the server, and the start point of their poll schedulers cannot be synchronized.

The monitored metrics are highly volatile. For example, two samples can differ with a high probability.

## Remote Deployment to Solaris Lists SPARC and x86 Systems

### Symptom:

The computers listed in the Deployment UI are typically filtered to the chosen operating environment for which you are deploying. However, you can see computers other than the chosen operating environment listed under the following situations:

- When you deploy to either a Solaris x86 or a Solaris SPARC server, the servers listed are for all Solaris architectures regardless of whether you selected Solaris x86 or Solaris SPARC as the target operating environment.
- When you deploy to any computer that is unclassified.

### Solution:

Verify that the target computer matches the chosen agent architecture for a successful deployment. If you proceed by selecting all computers listed, deployment succeeds for the matching architectures and fails on mismatched architectures.

## PMM Stops Polling an AIM

### Symptom:

The PMM stops polling an AIM (all instances of AIM) and sends CAAM6504 message. Additionally, the UI shows the affected instance as Critical in the AIM panel, Administration tab.

The polling stops for Cisco UCS, Microsoft clusters, IBM LPAR, and Solaris Zones AIMS that monitor multiple instances.

The reason is that the PMM cannot keep polling the AIM. When an AIM in such state is restarted, the MIB is not populated with the data from the affected instance. The PMM would assume that data are no longer available and removes from the management database.

### Solution:

- Verify why the instance is not in the Ready state. You can check for the following:
  - Invalid or expired credentials
  - Network connection
  - Hardware issue on the server
- Disable the instance in the Administration tab of the corresponding platform (not available on Cisco UCS).
- Remove the instance from the AIM.



## Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear

### Symptom:

When I remove a vCenter Server from management, objects of another managed vCenter Server disappear unexpectedly.

### Solution:

To avoid product management issues, do not install the vCenter AIM on a VM that is managing another vCenter Server. If you remove the monitoring and management of the vCenter associated with that VM from CA Virtual Assurance, it removes the objects associated with the vCenter including the VM system that is running the AIM.

## Resetting the vCenter Server Password Causes Data Collection to Fail

### Symptom:

After resetting the VMware vCenter Server password for the user that CA Virtual Assurance is using to communicate with VMware vCenter Server, data collection does not work.

### Solution:

Update the vCenter AIM configuration with the new password. You can update the password from the Administration tab in the user interface or through NodeCfgUtil on the server on which the vCenter AIM runs.

## Solaris Zones AIM Reset if a Monitored System is Down

### Symptom:

Solaris Zones AIM reset if a monitored system is down.

### Solution:

If you reset the AIM while one of its monitored systems is down, the AIM polls that system at each polling interval. The AIM does not update the properties until the system is up again.

## Unable to Connect to Microsoft SQL Server

**Symptom:**

Attempts to authenticate credentials to the Microsoft SQL Server 2005 and 2008 (Evaluation Edition) fail during product installation. The error message, Failed to establish connection to MSSQL displays.

**Solution:**

This issue occurs because TCP/IP is disabled by default on the Evaluation Edition. Enable TCP/IP.

## Upgrading SystemEDGE

**Symptom:**

When I upgrade SystemEDGE to 5.6.0, the AIMs of the previous CA Virtual Assurance release do not run.

**Solution:**

Upgrade Advanced Encryption and all AIMs to CA Virtual Assurance 12.6 level. SystemEDGE 5.6.0 does not load AIMs of previous CA Virtual Assurance releases.

## User Interface is Unresponsive on Provisioning and Policy Screens

**Symptom:**

If the database server is restarted while you are on the Provisioning page or Policy page, the user interface goes blank or is unresponsive.

**Solution:**

Log out of the CA Virtual Assurance user interface and log back in.

## vCenter Server AIM Attributes Show Zero

**Symptom:**

vCenter Server Attributes show zero.

**Solution:**

The following object values are only retrievable when the vCenter Server AIM is installed on the local vCenter Server instance. When the AIM is remote, these parameters show zero (0).

- vmvcAimStatServerCPUUsage [1.3.6.1.4.1.546.16.52.2.2.12.0]
- vmvcAimStatServerMemUsage [1.3.6.1.4.1.546.16.52.2.2.17.0]
- vmvcAimStatServerTotalPhysMem [1.3.6.1.4.1.546.16.52.2.2.18.0]
- vmvcAimStatServerUsedPhysMem [1.3.6.1.4.1.546.16.52.2.2.19.0]

## vCenter Server Connection Failed

**Symptom:**

After I have added a new vCenter Server connection under Administration, Configuration, the validation of the connection to the vCenter Server failed.

**Solution:**

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used vCenter Server connection data (server name, user, password, protocol, port) is still valid. If necessary, update the connection data.
- Verify, if the vCenter Server system is running and accessible.
- Verify, if the VMware Management Service on the vCenter Server system is running properly.

**To update the vCenter Server connection data**

1. Click  (Add) or  (Edit) that is associated with the failed connection.

The New or Edit vCenter Server dialog appears.

2. Add the valid server name, user, password, protocol, and port. Enable Managed Status and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the vCenter Server cannot be established, continue with the next procedure.

### To verify, if the vCenter Server system is running and accessible

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

```
nslookup <vCenter Server Name>
ping <IP Address of vCenter Server>
```

2. Verify the output of the commands to find out whether the vCenter Server has a valid DNS entry and IP address.

If the vCenter Server is not in the DNS, add the vCenter Server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

If the vCenter Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress <vCenter Server Name>
```

Enter the correct IP address and vCenter Server name. For example:

```
192.168.50.50 myvCenter
```

4. Click  (Validate) in the upper-right corner.

If the vCenter Server credentials and connection data are correct and you can ping the vCenter Server, the connection can still fail. In this case, it is possible that the vCenter Server causes the problem. If the connection to the vCenter Server cannot be established, continue with the next procedure.

### To verify, if the VMware Management Service on the vCenter Server system is running properly

1. Contact the vSphere Administrator to access the vCenter Server system.
2. Log in to the vCenter Server system and open Administrative Tools, Services from the Start menu.

The Services window opens.

3. Select the service *VMware VirtualCenter Server*. Start or restart the service.
4. Change to the CA Virtual Assurance user interface, vCenter Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Virtual Assurance validates the vCenter Server connection.

If the connection to the vCenter Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

Work with the vSphere administrator or VMware support to fix the vCenter Server connection problem.

## vCenter AIM Instance Status Icon Shows Disabled

**Symptom:**

After CA Virtual Assurance has discovered vCenter AIM instances in the network, the status icons of several instances show  (Disabled). This vCenter AIM instance is not managed.

This status appears, if CA Virtual Assurance has discovered a vCenter AIM with the following relationships:

- The vCenter AIM is configured for a vCenter Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a vCenter Server that has not been configured in the vCenter Servers pane.

**Solution:**

To change the status of the AIM instance to ready, do one of the following:

- Add the missing vCenter Server connection to the CA Virtual Assurance manager.
- Edit the existing vCenter Server connection and change its managed status to enabled.

## vCenter AIM Instance Status Icon Shows Discovery in Progress

**Symptom:**

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Discovery in Progress).

**Solution:**

Wait until the discovery process of the vSphere environment has completed. The discovery duration depends on the amount of managed objects related to virtual and physical resources in vSphere. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job has completed, CA Virtual Assurance adds a vCenter Server folder to the resources tree. Then you can start managing vSphere and its entire virtual infrastructure.

## vCenter AIM Instance Status Icon Shows Error

### Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

### Solution:

The following procedures address the most common issues which can cause a connection failure to the vCenter AIM:

- Verify, if the vCenter AIM Server is accessible.
- Verify, if SystemEDGE is running. Start or restart SystemEDGE if necessary.

### To verify, if the vCenter AIM server system is accessible

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

```
ping servename
```

2. Verify the output of the commands to find out whether the vCenter AIM server has a valid DNS entry and IP address.

If the vCenter AIM server is not in the DNS, add the vCenter AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

If the vCenter Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress servename
```

Enter the correct IP address and vCenter AIM server name. For example:

```
192.168.50.51 myvCenterAIM
```

4. Click  (Validate) in the upper-right corner of the vCenter AIM Server pane.

If the error status remains unchanged, continue with the next procedure.

**To verify, if SystemEDGE is running**

1. Log in to the vCenter AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.

The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.

2. Start or restart SystemEDGE.

Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.

3. Change to the CA Virtual Assurance user interface, vCenter AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Virtual Assurance validates the vCenter AIM Server connection.

If the error status remains unchanged, verify whether the data you gathered according to the requirements for this scenario is still valid.

## vCenter AIM Instance Status Icon Shows Multiple Instances

**Symptom:**

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Multiple AIMs manage this instance).

**Solution:**

Verify that your CA Virtual Assurance manager manages each vCenter Server with one vCenter AIM instance only. If a CA Virtual Assurance manager manages a vCenter Server through multiple AIM instances, management problems would occur. CA Virtual Assurance stops monitoring the associated vCenter Server.

Decide which AIM instance you want to use to manage the vCenter Server and remove the other instances from the vCenter AIM Servers pane.

**Follow these steps:**

1. Select the AIM instance you want to delete and click  (Delete).

The Delete Item dialog appears.

2. Click Yes.

Repeat these steps with other multiple instances until you have unique relationships between manager and AIM instance established.

## vCenter AIM Instance Status Icon Shows Warning

**Symptom:**

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Warning). Other instances on the AIM are still in error state.

**Solution:**

No specific actions are required for the associated instance. The warning informs you that other instances on the same AIM are in error state. The warning disappears when the problems of the other instances are resolved. If one or more instances of an AIM are in error state, then all other instances in the AIM show warning state.

## VM Usage Values Do Not Update Immediately After Power Down

**Symptom:**

VM usage values do not update immediately after power down.

**Solution:**

After VMs are powered off, usage values do not drop to 0 until the next successful poll. Polling can take up to 5 minutes, which is the default data collection and recording interval.



## Blank Query Results Tab after Upgrade

**Symptom:**

Remote Monitoring Query results show blank values after the upgrade.

**Solution:**

RM PMM require remote system names to comply with in Fully Qualified Domain Name (FQDN) notation while adding systems. However, the RM AIM leaves existing system non-FQDN names. This name mismatch shows blank query results. You can fix the name mismatch as following:

**Before Upgrade Conversion**

- Log in to refresh UI and remove all non-FQDN systems from Remote Monitoring.  
Deletes any associated systems, queries, instances, and monitors from both the manager (database); and the SystemEDGE agent with RM AIM plugin.
- Re-add these systems using FQDN notation and specify the same configuration sets to re-create associated queries, instances, and monitors.
- Upgrade.

**After Upgrade Conversion**

- Log in to the SystemEDGE agent machine and run the Refresh RM AIM plugin.
- Locate the `rmonwbem.cf` file containing the current Remote Monitoring configuration in the data directory path and make a copy of this file. For example, save as `rmonwbem-upgrade.cf`
- Log in to the Refresh UI and remove all non-FQDN systems from Remote Monitoring.  
Deletes any associated systems, queries, instances, and monitors from both the manager (database) and the RM AIM agent machine.
- Now upgrade. On the agent machine, run the `rmonwatch add` command with `rmonwbem-upgrade.cf` as input file. This process re-adds all systems and associated queries, instances, and monitors with FQDN notation.

**Note:** The After Upgrade Conversion approach has the advantage to re-add the systems automatically, and configure systems from the UI.

Query results show the values after the upgrade conversion.



# Glossary

---

Hyper-V is the Microsoft hypervisor-based server virtualization technology for Windows Server 2008 R2. Separate virtual machines (VMs) run on a single physical server and can run multiple different operating systems, such as Windows or Linux.

## ADES AIM

Active Directory and Exchange Server AIM (ADES AIM) lets you monitor Exchange Server and Active Directory environments. This AIM enables configuration of Active Directory and Exchange Server environments, and monitoring of the key performance indicators.

## application insight module, AIM

The SystemEDGE agent provides a plug-in architecture through which it can load optional *application insight modules (AIMs)* when it initializes. AIMs are functional extensions to the SystemEDGE agent. For example, the vCenter AIM enables SystemEDGE to manage vSphere environments through VMware vCenter Servers.

## autoshell

The *AutoShell* provides a command line and scripting environment that you can use to automate complex recurring and management tasks. AutoShell is not a programming language, but is a combination of a scripting language and a command line shell. AutoShell is based on the standardized scripting language ECMA-Script (JavaScript). While JavaScript is mostly known as a scripting language that is used on web pages, it does not need to run in a browser. It is a standalone scripting language implementing support for object orientation, XML and regular expression processing. AutoShell uses an out-of-the-box version of the Mozilla Spidermonkey JavaScript interpreter which also provides JavaScript functionality to the Mozilla Firefox web browser.

## autoshell loadable module, ALM

An *autoshell loadable module (ALM)* is an extension to the AutoShell core. Depending on the selected components of a CA Virtual Assurance installation, the required ALMs are installed automatically. For example, ALMs allow you to manage platforms like LPAR, Solaris Zones, or vCenter Server through AutoShell.

## blade (UCS)

Server that is attached to a Cisco UCS chassis.

## chassis (UCS)

Hardware frame that holds Cisco UCS switches and blades.

---

**Cisco Nexus 1000V Switch**

*Cisco Nexus 1000V Switch* is a Distributed Virtual Switch that can run in a VMware vSphere environment. The Cisco Nexus 1000V Switch consists of the Virtual Ethernet Module (VEM) and the Virtual Supervisor Module (VSM). On each ESX or ESXi host associated with a Cisco Nexus 1000V Switch, VEM replaces the VMware vSwitch and runs as a module in the hypervisor kernel. VSM controls multiple VEMs as one logical switch and runs in a VM on an ESX or ESXi host. For further details, see the Cisco Nexus 1000V Switch documentation at <http://www.cisco.com/go/1000vdocs>. CA Virtual Assurance VM provisioning supports VMware vNetwork Distributed Switches and Cisco Nexus 1000V Switches.

**Cisco Unified Computing System (UCS)**

*Cisco Unified Computing System (UCS)* provides data center hardware and virtualization services.

**Cluster**

A *cluster* consists of two or more independent computer systems that are linked together and work as a single entity. Clustering is used for parallel processing, load balancing, and fault tolerance.

**cmdlet**

A *cmdlet* is a command that must start with the first non-white character in a line. Because of this restriction they can only be used standalone and not as part of a broader JavaScript expression. In particular, they cannot be used as an rvalue (right hand side operand of an assignment operator).  
? is an example for an AutoShell cmdlet.

**container (Solaris)**

A Solaris *Container* provides complete runtime environments for applications. Resource management and Solaris Zones are parts of a container.

---

### **CPU shares (VMware)**

Shares are specified as natural numbers and express a proportional weight to each virtual machine.

Specifying shares makes sense only with regard to sibling virtual machines, vApps, or resource pools which have the same parent in the hierarchy. When you assign shares to a virtual machine, you always specify the priority for that virtual machine relative to other powered-on virtual machines.

For example, when competition occurs, a virtual machine with 2000 shares receives more CPU time than a virtual machine with 1000 shares. Shares are configured relative to the other shares; thus, only the proportion of shares matters, not the values of the shares. Three virtual machines with share values of 1000, 2000, 3000 act the same as three virtual machines with share values of 1, 2, 3. You can use any number scheme you prefer.

### **datacenter (VMware)**

A *datacenter* serves as a container for your hosts, virtual machines, resource pools, or clusters. If their virtual configurations meet the requirements of specific departments, datacenters can represent organizational structures, such as geographical regions or separate business functions. You can also use datacenters to create isolated virtual environments for testing or to organize your environment.

### **datastore (VMware)**

A *datastore* specifies a virtual representation of combinations of underlying physical storage resources in a datacenter. These physical storage resources can be provided by local disks on a server, by SAN disk arrays, and so on.

### **Dell EqualLogic**

*Dell EqualLogic* comprises virtualized iSCSI SAN solutions to use virtual storage for virtual servers.

### **dvPort group (VMware)**

Each VMware vNetwork Distributed Switch has one or more *dvPort Groups* assigned to it. dvPort Groups group multiple ports under a common configuration and provide a stable point for VMs connecting to labeled networks. A unique network label identifies each dvPort Group. The network labels are unique to the current datacenter.

A dvPort Group specifies port configuration options for each member port on a vNetwork Distributed Switch. dvPort Groups define how a connection is made to a network.

### **dvUplink port (VMware)**

*Distributed Virtual Uplinks (dvUplinks)* provide a level of abstraction for the physical NICs (vmnics) on the ESX Hosts. Each physical NIC is mapped to a dvUplink. For each host associated with a VMware vNetwork Distributed Switch, each physical NIC (uplink) is assigned to the vNetwork Distributed Switch through one uplink port.

---

**dynamic reconfiguration connector index, DRC-index (LPAR)**

Each slot in a physical system unit has a *DRC-index* assigned to it. The deploy process requires this number to perform the actual creation of the LPARs. The management console (HMC or IVM) and the system uses this index to identify uniquely each slot on the system. The DRC-index is not assigned to a slot until the unit is powered up.

**ESX/ESXi host (VMware)**

An *ESX or ESXi host* is a physical computer that uses ESX or ESXi Server virtualization software to run virtual machines. Hosts provide the CPUs and memory resources that virtual machines use and give virtual machines access to storage and network connectivity.

**fair share scheduler, FSS (Solaris)**

The *fair share scheduler (FSS)* specifies a scheduler class that allocates CPU time based on shares. Shares define the portion of the system's CPU resources allocated to a project.

**fault tolerance, FT (VMware)**

VMware vSphere lets you enable *Fault Tolerance (FT)* on a VM defined to a cluster which is configured for High Availability (HA). Fault Tolerance creates a secondary VM on another ESX Server in the cluster. The secondary VM operates in lock-step mode with the primary VM that is executing the workload. If there is a failure, the secondary VM immediately takes over the workload execution from the point of failure. CA Virtual Assurance discovers and manages primary and secondary VMs in a cluster.

**Fibre Channel, FC**

*Fibre Channel* is a standardized gigabit-speed technology for transmitting data between computer devices. Fibre Channel is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

**funclet**

*Funclets* maintain the verbose command like syntax with optional clauses, stringification and so on. Funclets are often used like cmdlets, that is, standalone in a single line. They can return a value that can be processed as part of a broader expression.

**global zone (Solaris)**

A *global zone* is a zone that is contained on every Solaris system. If non-global zones exist on the system, the global zone is the default zone for the system and for systemwide administration.

**Hardware Management Console, HMC (LPAR)**

The *Hardware Management Console (HMC)* is an external appliance used to perform management tasks on IBM POWER Servers. HMC can be used to create or change logical partitions, including dynamically assigning resources to a partition. The HMC communicates with the server firmware layers of POWER servers, providing a single point of control in large PowerVM environments.

---

**Host Bus Adapter, HBA**

A *Host Bus Adapter (HBA)* is the interface card which connects a host to a *Storage Area Network (SAN)*.

**I18n (Internationalization)**

*I18N (internationalization)* is the modification of a software product so that it can potentially handle multiple languages, time and date formats, writing conventions like the formatting of numbers (decimal separator, digit grouping), and so on. CA Virtual Assurance uses UTF-8 encoding to display language-specific characters like the German ü (umlaut), the French è (grave accent), or Japanese characters in input and output data.

**Integrated Virtualization Manager (IVM, LPAR)**

The *Integrated Virtualization Manager (IVM)* is an enhancement of the Virtual I/O Server (VIOS) and allows you to manage a single POWER Server. IVM lets you create and manage LPARs. IVM enables management of VIOS functions and provides a web-based user interface.

**Internet Small Computer Systems Interface, iSCSI**

*iSCSI* is used to facilitate data transfers over intranets and to manage storage over large distances. iSCSI encapsulates SCSI commands in IP packets, which are routed just like any other IP packet on the network. When the IP packet reaches its destination, the iSCSI device removes the encapsulation and interprets the SCSI command.

**L10n (Localization)**

*L10N (localization)* is the implementation of a specific language for an already internationalized software.

**lightweight process, LWP (Solaris)**

*Lightweight processes (LWP)* belong to the Solaris 10 kernel thread model. LWPs form the execution context for a user thread by associating a user thread with a kernel thread. In the Solaris 10 kernel, kernel services and tasks run as kernel threads. When a user thread is created, the associated LWP and kernel threads are also created and linked to the user thread. Resource control allows to set bounds for LWPs.

**logical memory block, LMB (LPAR)**

A *logical memory block (LMB)* specifies the granularity of physical and logical memory assigned to an LPAR (for example: 256 MB).

**Logical Partition, LPAR**

A *Logical Partition (LPAR)* is a subset of hardware resources, virtualized as a separate system. A physical system can be partitioned into multiple LPARs, each providing a separate operating system and applications. The number of logical partitions depends on the hardware configuration of the system. LPARs are typically used for different environments, such as databases, web servers, and so on. LPARs communicate as separate systems in the network.

---

**Management Information Base (MIB)**

A *Management Information Base (MIB)* is a data store that describes properties of a resource. MIBs are written in ASN.1, which is a language specified by a management standard and complies with OSI's structure of management information (SMI) standards for defining SNMP MIBs.

**MIB objects, MIB attributes**

A *MIB object* is an entity defined in a MIB that represents one or more resource objects or data items. MIB objects include groups, tables, and individual attributes, and they must be defined in accordance with the structure for management information (SMI).

**NetApp filer**

A *NetApp filer* is a type of disk storage device which owns and controls a filesystem, and presents files and directories to hosts over the network.

**network installation manager, NIM (LPAR)**

A *Network Installation Manager (NIM)* provides a central point of management for installing and maintaining AIX images for LPARs and individual servers. It also facilitates the installation of all of those instances from the same master image, from different images, from installation media or from a previous mksysb of that instance. An instance refers to an OS image, regardless of whether it is an LPAR or on a physical server.

**non-global zone (Solaris)**

A *non-global zone* provides a virtualized operating system environment in a single instance of the Solaris operating system. The Solaris Zones software partitioning technology virtualizes operating system services.

**onTap**

The *onTap* framework is a free service-oriented web application framework.

**platform management module, PMM**

A *Platform Management Module (PMM)* is a web service which is responsible for providing connection and operational support for the corresponding environment. Supported environments are for example: VMware vSphere, Microsoft Hyper-V, IBM PowerVM, Solaris Zones, Cisco UCS, or Microsoft Cluster Service. A PMM manages connections with the servers of these environments, performs environment-related operations, retrieves data from the corresponding AIM, and populates the CA Virtual Assurance Management Database.

**policy based configuration**

*Policy based configuration* provides the ability to create agent configuration policy that you can deploy to sets of managed machines in one operation.

**poll interval**

The *poll interval* is the length of time between consecutive polls of a resource group.

**POWER processors (LPAR)**

*POWER processors* are RISC-based and used as the CPU in many of IBM servers, mini-computers, workstations, and supercomputers.



---

**processor set, pset (Solaris)**

*Processor sets* define disjoint groups of CPUs. Each processor set can contain zero or more processors. It is a resource element in the resource pools configuration.

**project (Solaris)**

A *project* defines a container associated with a host. It is an abstraction layer that helps to organize and manage the collection of physical system resources.

Projects are collections of tasks, which are collections of processes. A new task is started in a project when a new session is opened by a login, cron, newtask, setproject or su command. Each process belongs to only one task, and each task belongs to only one project.

Projects and tasks are the basic entities which are used to identify workloads in the Solaris 10 operating system. A project is associated with a set of users and a set of groups. Users and groups can run their processes in the context of a project they are a member of, but they can be members of more than one project. The project is the basic entity against which the usage of resources can be restricted. The task is the entity to which a process is associated and the project is associated with a set of tasks.

**regular expressions**

*Regular expressions* are text patterns used for matching. Regular expressions are strings that include a mix of plain text and special characters to indicate the kind of matching required.

**resource control (Solaris)**

*Resource control* can be set up for Solaris Zones directly by defining bounds on the consumption of specific resources for a workload. A workload is an aggregation of all processes of an application or group of applications.

Resource controls are stored in the /etc/project file or in a zone's configuration through the zonecfg command described in zonecfg(1M).

**resource pool (Solaris)**

A *resource pool* defines a configuration mechanism for partitioning system resources. A resource pool is an association between resource groups which can be partitioned.

**resource pool (VMware)**

A *resource pool* defines partitions of physical computing and memory resources of a single host or a cluster. You can partition any resource pool into smaller resource pools to divide and assign resources to specific groups or for specific purposes. You can also hierarchically organize and nest resource pools.

**service profile**

Set of configuration information about Cisco UCS hardware, including interfaces, fabric connectivity, and network and server identity.

**shared memory (Solaris)**

*Shared memory* defines the total amount of memory that can be used by the processes that run in a project.

---

## Simple Network Management Protocol (SNMP)

The *Simple Network Management Protocol (SNMP)* is the standard management protocol for the Internet. SNMP management applications and agents use the get request, set request, get-next request, get response, and trap PDUs to communicate with each other. MIBs, which keep track of network and system resources and applications, define the data they exchange.

## SNMPv3

*SNMPv3* is a protocol that has the following three levels of communication:

**noAuthNoPriv:** Mirrors SNMPv1 and SNMPv2 in that messages are accompanied by a username, which must be consistent between sender and receiver.

**AuthNoPriv:** Uses a consistent username and a password.

**AuthPriv:** Uses a username, password, and an encryption key that encrypts the body of the message.

## Storage Area Network, SAN

A *storage area network (SAN)* is an architecture to connect remote computer storage devices to servers in such a way that the devices appear as locally connected to the operating system.

## stringification

*Stringification* takes a sequence of characters and turns it into a proper JavaScript literal string.

## task (Solaris)

A *task* represents a set of work over time. Each task is associated with one project.

## time-sharing scheduler, TS (Solaris)

A *time-sharing scheduler (TS)* specifies a scheduler class that tries to provide every process with equal access to available CPUs. It allocates CPU time on a priority basis.

## trap

A *trap* is an unsolicited message that an SNMP agent can send to one or more managers to notify management applications of agent and resource events. SNMP traps are generic (common to all types of SNMP agents) or enterprise-specific (unique to the agent that sends it).

## UCS

See *Cisco Unified Computing System (UCS)*.

## UCS Manager

Software module that manages UCS hardware (switches, chassis, and blades).

---

**vApp (VMware)**

A *vApp* is a specific resource pool which treats a collection of VMs as a single unit. vApp uses the Open Virtualization Format (OVF). OVF is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels associated with it. CA Virtual Assurance can perform operations on a vApp. An operation on a vApp is propagated to all VMs in the vApp.

**vCenter Server (VMware)**

VMware *vCenter Server* provides the central point of control for configuring, provisioning, and managing a virtual vSphere environment. vCenter Server runs as a service on Microsoft Windows Servers and Linux Servers.

**vCenter Server Agent (VMware)**

The VMware *vCenter Server Agent* connects ESX Servers with a vCenter Server.

**vCenter Server Database (VMware)**

The VMware *vCenter Server Database* stores persistent information about the physical servers, resource pools, datacenters, and virtual machines managed by the VirtualCenter.

**virtual disk (VMware)**

A *virtual disk* defines the disk drive in a virtual guest operating system. A virtual disk is a specific file or a set of files that reside on the local host or on a remote file system. It behaves like a physical disk drive in an operating system.

**Virtual I/O Server, VIOS (LPAR)**

A *Virtual I/O Server (VIOS)* is a special logical partition configured to own all physical I/O resources and provides its virtualization capabilities to other LPARs. LPARs access disk, network, and optical devices through the Virtual I/O Servers as virtual devices. Each PowerVM system with virtualized resources has a Virtual I/O Server.

**virtual machine hardware version 7 (VMware)**

*Virtual Machine Hardware Version 7*, specifies a virtual hardware generation from VMware and is the default for VMs created with vSphere. It supports hot plug, for example, for CPU and memory. If hot plug is enabled in the VM, CA Virtual Assurance supports hot plug for CPU and memory as well.

**Note:** For information about VMware Virtual Machine Hardware Version 7, see the VMware documentation.

**virtual machine, VM (VMware)**

A *virtual machine (VM)* is a software-based computer that runs an operating system and applications like a physical computer. A virtual machine consumes resources dynamically on its physical host, depending on its workload. Because virtual machines are flexible computing units, their deployment comprises a wide range of environments like datacenters, clusters, cloud computing, test environments, desktops, or laptops. Their primary strength lies in datacenters, where they are used for server consolidation, workload optimization, and energy efficiency.

---

**virtual NIC (VMware)**

A *virtual NIC* is a virtual Ethernet adapter on a virtual machine. The guest operating system communicates with the virtual Ethernet adapter through a device driver as if the virtual Ethernet adapter was a physical Ethernet adapter. The virtual Ethernet adapter has its own MAC address, one or more IP addresses, and responds to the standard Ethernet protocol like a physical NIC.

**virtual switch (VMware)**

A *virtual switch* works like a physical switch. Each ESX Server has its own virtual switches that connect to virtual machines through port groups. These virtual switches also have uplink connections to the physical Ethernet adapters on the ESX server. Virtual machines communicate with the outside world through physical Ethernet adapters connected to virtual switch uplinks.

**vNetwork Distributed Switch, vDS (VMware)**

A *VMware vNetwork Distributed Switch* abstracts the configuration of virtual switches from the host to the datacenter level. A vNetwork Distributed Switch operates as a single virtual switch that spans across all hosts in a datacenter which are associated with that switch. vNetwork Distributed Switches consist of distributed port groups which are similarly configured to port groups on standard switches, but extend across multiple hosts. These properties allow virtual machines to maintain a consistent network configuration as they migrate among multiple hosts.

Like a vNetwork Standard Switch, each vNetwork Distributed Switch is a network hub that VMs can use. A vNetwork Distributed Switch can forward traffic internally between VMs or link to an external network by connecting to physical NICs (uplink adapters). For further details, see the vNetwork Distributed Switches documentation at <http://pubs.vmware.com>.

CA Virtual Assurance VM provisioning supports VMware vNetwork Distributed Switches and Cisco Nexus 1000V Switches. You can manage Virtual Distributed Switches through the vNetwork panel, AutoShell, or CLI commands.

**vNetwork Standard Switch, vSwitch (VMware)**

CA Virtual Assurance manages policies and properties of standard vSwitches which are abstracted network devices. A *VMware vNetwork Standard Switch (vSwitch)* operates on a single host and virtual machines on that host can be attached to the standard switch. A vSwitch can route traffic internally between VMs and link to external networks. vSwitches combine the bandwidth of multiple network adapters and balance communications traffic among them. A vSwitch can handle physical NIC failover.

**zone (Solaris)**

Solaris *Zones* define a virtualized operating system environment that you can set up for Solaris 10 systems. Zones virtualize operating system services and provide an isolated, secure environment for applications. Each Solaris system contains a global zone that is the default zone for the system. For example, you can create, delete, modify, halt, or reboot non-global zones.

# Index

---

## (

- (Optional) Apply Policy and Template Updates to Servers and Verify Updates • 318
- (Optional) Manage the Base Policy and Templates for One or More Servers • 317
- (Optional) Update the Policy or Templates • 318

## A

- Access Control • 235
- Access the CA EEM User Interface • 27
- Access the User Interface • 24
- Action Types • 133
- Active Directory • 25
- Active Directory and Exchange Server AIM • 245
- Add a Cisco UCS Manager Server • 96
- Add a Logical Partition for an IBM AIX Computer • 187
- Add a New vCenter Server Connection to the Manager • 260
- Add a Solaris Zone • 186
- Add a Virtual Machine (Hyper-V Server) • 185
- Add a Virtual Machine (vCenter Server) • 182
- Add MIB Extensions to a Template or a Policy • 313
- Add Monitors to a Template or the Policy • 302
- Add the AIM Instance for the vCenter Server • 264
- Additional Provisioning Information • 189
- ADES AIM • 339
- Adjusting Poll Interval Settings for Solaris Zones Environments • 323
- Advantages • 248
- Advantages of Remote Monitoring • 233
- Agent Configuration • 118
- Agent Configuration Without Write Community • 204
- Agent Policy Dashboard Views • 173
- Agent Visualization • 119
- Agent-less Monitored Systems • 234
- application insight module, AIM • 339
- Application Insight Modules (AIMs) • 57
- Apply Policy and Templates to Servers and Verify Settings • 315
- Apply Policy to Machines • 169
- Architecture • 19, 236, 246
- Assign User Groups Access Rights to Services • 35

- Assign Users to Groups • 34
- Associate Service Profiles with Blades • 155
- Attributes Show a Value of Zero • 324
- Audience • 15
- Audit Trail • 194
- Automation • 236
- autoshell • 339
- autoshell loadable module, ALM • 339

## B

- blade (UCS) • 339
- Blank Query Results Tab after Upgrade • 337
- Browsers Do Not Display Consecutive Spaces in Events • 324

## C

- CA NSM Fails to Receive Traps on Windows • 324
- CA Product References • 3
- CA Virtual Assurance vCenter Management Recommendations • 279
- Cancel Network Discovery • 62
- Change the CA Embedded Entitlements Manager Administrator Password (EiamAdmin) • 29
- Change the Database Administrator (sa) Password • 30
- Change the Domain Server a Distribution Server Connects To • 193
- Change the System User Password for Active Directory Security • 32
- Change the System User Password for Native Security • 31
- chassis (UCS) • 339
- Cisco Nexus 1000V Switch • 340
- Cisco UCS • 93
- Cisco UCS Folder Does Not Display in UI • 325
- Cisco UCS Overview • 95
- Cisco Unified Computing System (UCS) • 340
- Cluster • 340
- cmdlet • 340
- Collection Engine and Report Support for Remote Monitoring Metrics • 243
- Compatibility Libraries for Linux • 221
- Configuration • 235
- Configuration Prerequisites • 239

---

- Configuration Requirements for Managing Environments • 41
- Configure ADES AIM with NodeCfgUtil in Command Mode • 251
- Configure ADES AIM with NodeCfgUtil in Dialog Mode • 252
- Configure AIMs with NodeCfgUtil in Command Mode • 47
- Configure AIMs with NodeCfgUtil in Dialog Mode • 43
- Configure and View Applied Policies • 171
- Configure CA SDM • 126
- Configure CA Virtual Assurance to Forward Events • 55
- Configure Data Collection for a Data Center • 144
- Configure Data Collection for a Server • 145
- Configure Data Collection for a Virtual Resource • 147
- Configure Object Aggregation • 108, 297
- Configure Performance Thresholds • 149
- Configure SNMPv1 Traps by Editing the sysedge.cf File • 52
- Configure the CA SDM Ticket Status Setting • 127
- Configure the Environment to Enable ADES AIM Monitoring • 250
- Configure the Metric Filter • 150
- Configure the Service Poller • 98
- Configure the SNMP Data Poller • 97
- Configure Windows for SNMP • 54
- Configuring Data Collection • 141
- Configuring Remote Monitor Systems • 240
- Configuring Resources • 41
- Contact CA • 4
- container (Solaris) • 340
- Conventions • 16
- CPU shares (VMware) • 341
- Create a Blade Power Action • 152
- Create a Custom Action • 135
- Create a History Monitor • 310
- Create a Log File Monitor • 307
- Create a New Package Wrapper • 205
- Create a Policy • 291
- Create a Process Group Monitor • 311
- Create a Process Monitor • 305
- Create a Rule • 129
- Create a Rule for CPU Metric to Decrease Allocation • 273
- Create a Rule for CPU Metric to Increase Allocation • 273

- Create a Service • 62
- Create a Threshold Monitor • 302
- Create a Windows Event Monitor • 308
- Create an Action for CPU Metric • 272
- Create Automation Policy • 128
- Create CA EEM Users • 27
- Create Default User Groups • 28
- Create Templates for Server Workload • 301
- Create User Groups • 33

## D

- Dashboard Status Views • 122
- Database Considerations • 276
- Databases • 22
- datacenter (VMware) • 341
- datastore (VMware) • 341
- DB Transaction Log Sizes Increase Unexpectedly • 325
- Default Package Wrappers • 203
- Define a Schedule • 138
- Define an Action Sequence • 136
- Define MIB Extensions • 166, 313
- Define SRM Control Settings • 168
- Define SystemEDGE Policy Control Settings • 158, 292
- Define Traps and Communities • 163, 298
- Delete a Network • 61
- Delete Monitors from Templates or a Policy • 315
- Delete User Groups • 36
- Dell EqualLogic • 341
- Departmental Data Center • 284
- Deploying/Installing SystemEDGE Agents Using Custom Ports • 197
- Deployment Components • 192
- Deployment Credential Restrictions • 195
- Deployment Dashboard Views • 214
- Deployment Jobs • 212
- Deployment Management Certificate on Linux or UNIX • 221
- Deployment Management Certificate on Windows • 220
- Deployment Package Configuration File • 201
- Deployment Package Library • 198
- Deployment Package Library Maintenance • 200
- Deployment Packages • 202
- Deployment Primer Installation on Linux or UNIX • 220
- Deployment Primer Installation on Windows • 219

---

- Deployment Restrictions • 195
- Deployment Sizing Key Factors • 196
- Deployment to Windows Vista, Windows 2008 and Windows XP Computers Running Firewall Software • 223
- Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero • 326
- Device Management for VMs • 71
- Discover a Network • 60
- Discover a System • 59
- Discover Environments • 42
- Discover Resources • 59
- Discovery • 41
- Distributed Virtual Switches (DVS) • 72
- Distribution Server Recommendations • 283
- Documentation Changes • 5
- Domain Server Recommendations • 283
- dvPort group (VMware) • 341
- dvUplink port (VMware) • 341
- dynamic reconfiguration connector index, DRC-index (LPAR) • 342

## E

- Edit a Service • 64
- Empty Task ID for the dpmvc virtualswitch Command • 327
- Enable Maintenance Mode • 122
- ESX Host Fault Tolerance Attributes • 77
- ESX/ESXi host (VMware) • 342
- Event Forwarding • 54
- Example
  - How to Monitor a Specific Windows Performance Registry Metric • 175
  - How to Monitor User-specific Metrics (MIB Extensions) • 173

## F

- fair share scheduler, FSS (Solaris) • 342
- Fault Tolerance • 75
- Fault Tolerance Requirements • 76
- fault tolerance, FT (VMware) • 342
- Features • 247
- Features and Benefits • 234
- Fibre Channel, FC • 342
- FIPS 140-2 Encryption • 321
- FIPS Overview • 321
- funclet • 342

## G

- General Recommendation for vCenter AIM Monitoring • 278
- global zone (Solaris) • 342

## H

- Hardware Management Console, HMC (LPAR) • 342
- Hardware Specifications • 276
- Host Bus Adapter, HBA • 343
- Hot-plug Support for VMs • 70
- How CA Embedded Entitlements Manager Works with CA Virtual Assurance • 26
- How to Apply Policy and Layered Templates to Servers • 289
- How to Configure the vCenter Server Management Components • 256
- How to Create or Update a Service Profile • 100, 154
- How to Create SRM Policy • 166
- How to Create SystemEDGE Policy • 156
- How to Use Policy Actions to Identify Performance Issues • 271

## I

- l18n (Internationalization) • 343
- IBM PowerVM (LPAR) • 86
- Imaging Services • 177
- Infrastructure Deployment Process • 215
- Install ADES AIM • 248
- Integrated Virtualization Manager (IVM, LPAR) • 343
- Integration • 236
- Interaction Between AIX LPAR Management Components • 87
- Interaction Between Cisco UCS Management Components • 94
- Interaction Between Solaris Zones Management Components • 83
- Interactions Between Hyper-V Server Management Components • 85
- Interactions Between vCenter Server Management Components • 69
- internet Small Computer Systems Interface, iSCSI • 343
- Introduction • 15

## K

- Key Performance Indicator Metrics • 235
- Key Points About Metrics Collection • 141

---

## L

- L10n (Localization) • 343
- Large Environments • 287
- Layered Templates Concept • 290
- lightweight process, LWP (Solaris) • 343
- Local and Remote Monitors Do Not Show the Same Values • 327
- logical memory block, LMB (LPAR) • 343
- Logical Partition, LPAR • 343
- Logical Volume Management in VMs • 77
- LPAR Monitoring • 88

## M

- Manage a Renamed vCenter Server • 51
- Manage a Resource • 62
- Manage Multiple vCenter AIMs • 49
- Managed Mode and Legacy Mode • 110
- Managed Virtual Environments • 66
- Management DB • 22
- Management Information Base (MIB) • 344
- Manager Components • 57
- Managing Policies • 125
- Managing Resources • 57
- Managing Systems Using Remote Monitoring • 243
- Managing Traps • 225
- Managing Users • 25
- Manual Installation of the Infrastructure Deployment Primer Software • 219
- MIB objects, MIB attributes • 344
- Microsoft Cluster Overview • 91
- Microsoft Cluster Service • 90
- Microsoft Hyper-V Server • 84
- Modify a Package Wrapper • 205
- Modify Cluster Properties • 92
- Monitor Configuration • 59
- Monitor Definition • 157
- Monitor Distributed Virtual Switches Through Events • 73
- Monitor vApps Through Events • 74
- Multiple Data Centers • 285

## N

- Native Security • 26
- NetApp filer • 344
- Network Considerations • 277
- network installation manager, NIM (LPAR) • 344
- non-global zone (Solaris) • 344

- Notes on Infrastructure Deployment Using IPv6 Addresses • 218

## O

- onTap • 344
- Overview • 19, 233, 245

## P

- Package Filter • 200
- Password Management • 29
- Perform a Point Agent Configuration • 118
- Performance • 248
- Performance Considerations during Initial Discovery • 281
- Performance DB • 22
- platform management module, PMM • 344
- PMM Stops Polling an AIM • 328
- policy based configuration • 344
- Policy Configuration • 155
- poll interval • 344
- POWER processors (LPAR) • 344
- Prerequisites for Automatically Deploying CA Virtual Assurance Infrastructure • 216
- processor set, pset (Solaris) • 345
- project (Solaris) • 345
- Protocols for Transferring Packages Employed by IDManager • 219
- Provide the Deployment Management Certificate to a Primer Installation • 220
- Provisioning Resources • 177

## R

- Register a Cluster • 91
- Register a UCS AIM Server • 96
- regular expressions • 345
- Related Publications • 15
- Remote and Multi-instance vCenter Server Support • 70
- Remote Configuration and Deployment • 111
- Remote Deployment • 190
- Remote Deployment and Policy Configuration Overview • 274
- Remote Deployment and Policy Configuration Recommendations • 281
- Remote Deployment Architecture • 191
- Remote Deployment to Solaris Lists SPARC and x86 Systems • 328



---

- Remote Deployment to UNIX/Linux Using Non Privileged User Account • 195
- Remote Monitoring • 117, 233
- Remove a Cluster • 92
- Remove a UCS Server • 97
- Remove Users or User Groups from a User Group • 34
- Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear • 329
- Report Availability • 231
- Reporting • 231
- Requirements for LPAR Management • 52
- Requirements for Solaris Zones Management • 52
- Resetting the vCenter Server Password Causes Data Collection to Fail • 329
- Resilience • 235
- Resource Allocation • 79
- Resource Allocation Best Practices • 81
- Resource Allocation Limit • 81
- Resource Allocation Reservation • 80
- Resource Allocation Shares • 80
- resource control (Solaris) • 345
- resource pool (Solaris) • 345
- resource pool (VMware) • 345
- Review Interactions Between vCenter Server Management Components • 258
- Review Policy Application Progress • 170
- Review Requirements • 257
- Rule Planning • 128
- Rules and Actions • 125
- Run Reports • 231

## S

- Scalability • 193, 236
- Scalability Best Practices • 274
- Scalability Limitation in Terms of Monitored Servers • 279
- Scalability Limitations in Terms of Monitored Objects • 279
- Scalability Recommendations • 277
- Scalability Use Cases • 283
- Scenarios and Best Practices • 255
- Search for Users or User Groups • 36
- Security and Maintenance • 122
- Security Considerations for Active Directory • 26
- service profile • 345
- Service Profiles • 99, 153
- Service Response Monitoring • 113

- Set User Group Permissions • 36
- shared memory (Solaris) • 345
- Simple Network Management Protocol (SNMP) • 346
- SNMP Trap Receiver • 225
- SNMP Trap Receiver Configuration File • 226
- SNMP V3 Engine ID • 56
- SNMPv3 • 346
- Solaris Zones • 82
- Solaris Zones AIM Reset if a Monitored System is Down • 329
- SRM Tests • 114
- State Management Model • 107
- Stateless Monitoring • 109
- Storage Area Network, SAN • 346
- stringification • 346
- SystemEDGE AIMS • 111
- SystemEDGE Features • 103
- SystemEDGE Trap Forwarding • 229
- Systems Management • 101
- Systems Management MIB • 105

## T

- task (Solaris) • 346
- Tasks for External Directory Users • 33
- Test Definition • 167
- Threshold Definition • 167
- time-sharing scheduler, TS (Solaris) • 346
- Track Deployment Job Status • 212
- trap • 346
- Troubleshoot the vCenter AIM Instance Connection • 265
- Troubleshoot the vCenter Server Connection • 261
- Troubleshooting • 323
  - vCenter AIM Instance Status Icon Shows Disabled • 270
  - vCenter AIM Instance Status Icon Shows Discovery in Progress • 267
  - vCenter AIM Instance Status Icon Shows Error • 268
  - vCenter AIM Instance Status Icon Shows Multiple Instances • 267
  - vCenter AIM Instance Status Icon Shows Warning • 269
  - vCenter Server Connection Failed • 262

## U

- UCS • 346

---

- UCS Action Types for Policies • 152
- UCS Manager • 346
- UCS Trap Management • 98
- Unable to Connect to Microsoft SQL Server • 330
- Uninstall the ADES AIM • 254
- Upgrading SystemEDGE • 330
- Use a Predefined Action Type • 131
- Use Case
  - Adding a New Rule to a Service • 140
  - Adding a Server to a Service • 139
  - Defining an Action • 140
- Use Case Scenario • 238
- Use Cases for Policies • 139
- User Access Control • 25
- User Interface • 23
- User Interface is Unresponsive on Provisioning and Policy Screens • 330
- User-scoped Authentication for vCenter Server • 37
- Using Multiple Distribution Servers • 192
- Using Multiple Multi-instance AIMS • 101

## V

- vApp (VMware) • 347
- vApp Support • 74
- vCenter AIM Instance Status Icon Shows Disabled • 333
- vCenter AIM Instance Status Icon Shows Discovery in Progress • 333
- vCenter AIM Instance Status Icon Shows Error • 334
- vCenter AIM Instance Status Icon Shows Multiple Instances • 335
- vCenter AIM Instance Status Icon Shows Warning • 336
- vCenter AIM Monitoring Recommendations • 278
- vCenter Management Limitations in Terms of Virtual Machines • 280
- vCenter Server (VMware) • 347
- vCenter Server Agent (VMware) • 347
- vCenter Server AIM Attributes Show Zero • 331
- vCenter Server Connection Failed • 331
- vCenter Server Database (VMware) • 347
- vCenter Server in a Cluster • 75
- Verify the vCenter Server Folder Appearance in the Resources Tree • 270
- View Deployment History • 213
- View Managed Object States • 120
- View SRM Tests • 121
- View SystemEDGE Monitors • 120

- virtual disk (VMware) • 347
- Virtual I/O Server, VIOS (LPAR) • 347
- virtual machine hardware version 7 (VMware) • 347
- virtual machine, VM (VMware) • 347
- virtual NIC (VMware) • 348
- virtual switch (VMware) • 348
- Visualization • 235
- VM Fault Tolerance Properties • 76
- VM Usage Values Do Not Update Immediately After Power Down • 336
- VMware vCenter and vSphere User Permissions • 178
- VMware vCenter Provisioning • 178
- VMware vCenter Server • 67
- VMware vCenter Server Architecture • 77
- vNetwork Distributed Switch, vDS (VMware) • 348
- vNetwork Standard Switch, vSwitch (VMware) • 348
- vNetwork Standard Switches (vSwitch) • 71

## Z

- zone (Solaris) • 348