CA Virtual Assurance for Infrastructure Managers

Administration Guide

Release 12.9

CA Technologies Product References

This document references the following CA Technologies products:

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

Contact CA Technologies

Contact CA Support

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

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

Providing Feedback About Product Documentation

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

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

Contents

Chapter 1: Introduction

This section contains the following topics:

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, and MIB attributes.

Performance Metrics Reference

Describes the performance metrics that are available for monitoring the systems performance of the supported platforms.

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.

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.

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

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 user interface.

CA Virtual Assurance integrates with the following external technologies:

- Cisco Unified Computing System
- Citrix XenDesktop
- Citrix XenServer
- Huawei GalaX
- IBM PowerHA
- IBM PowerVM for AIX virtualization management
- Microsoft Cluster Service
- Microsoft Hyper-V Server for virtual machine (VM) management and optional integration with Microsoft System Center Virtual Machine Manager (SCVMM) for provisioning
- Red Hat Enterprise Virtualization (RHEV)
- Solaris Zones virtualization management
- VMware vCenter Server for virtual machine (VM) management
- VMware vCloud

CA Virtual Assurance leverages the following existing technologies:

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

The following diagram displays the product architecture:

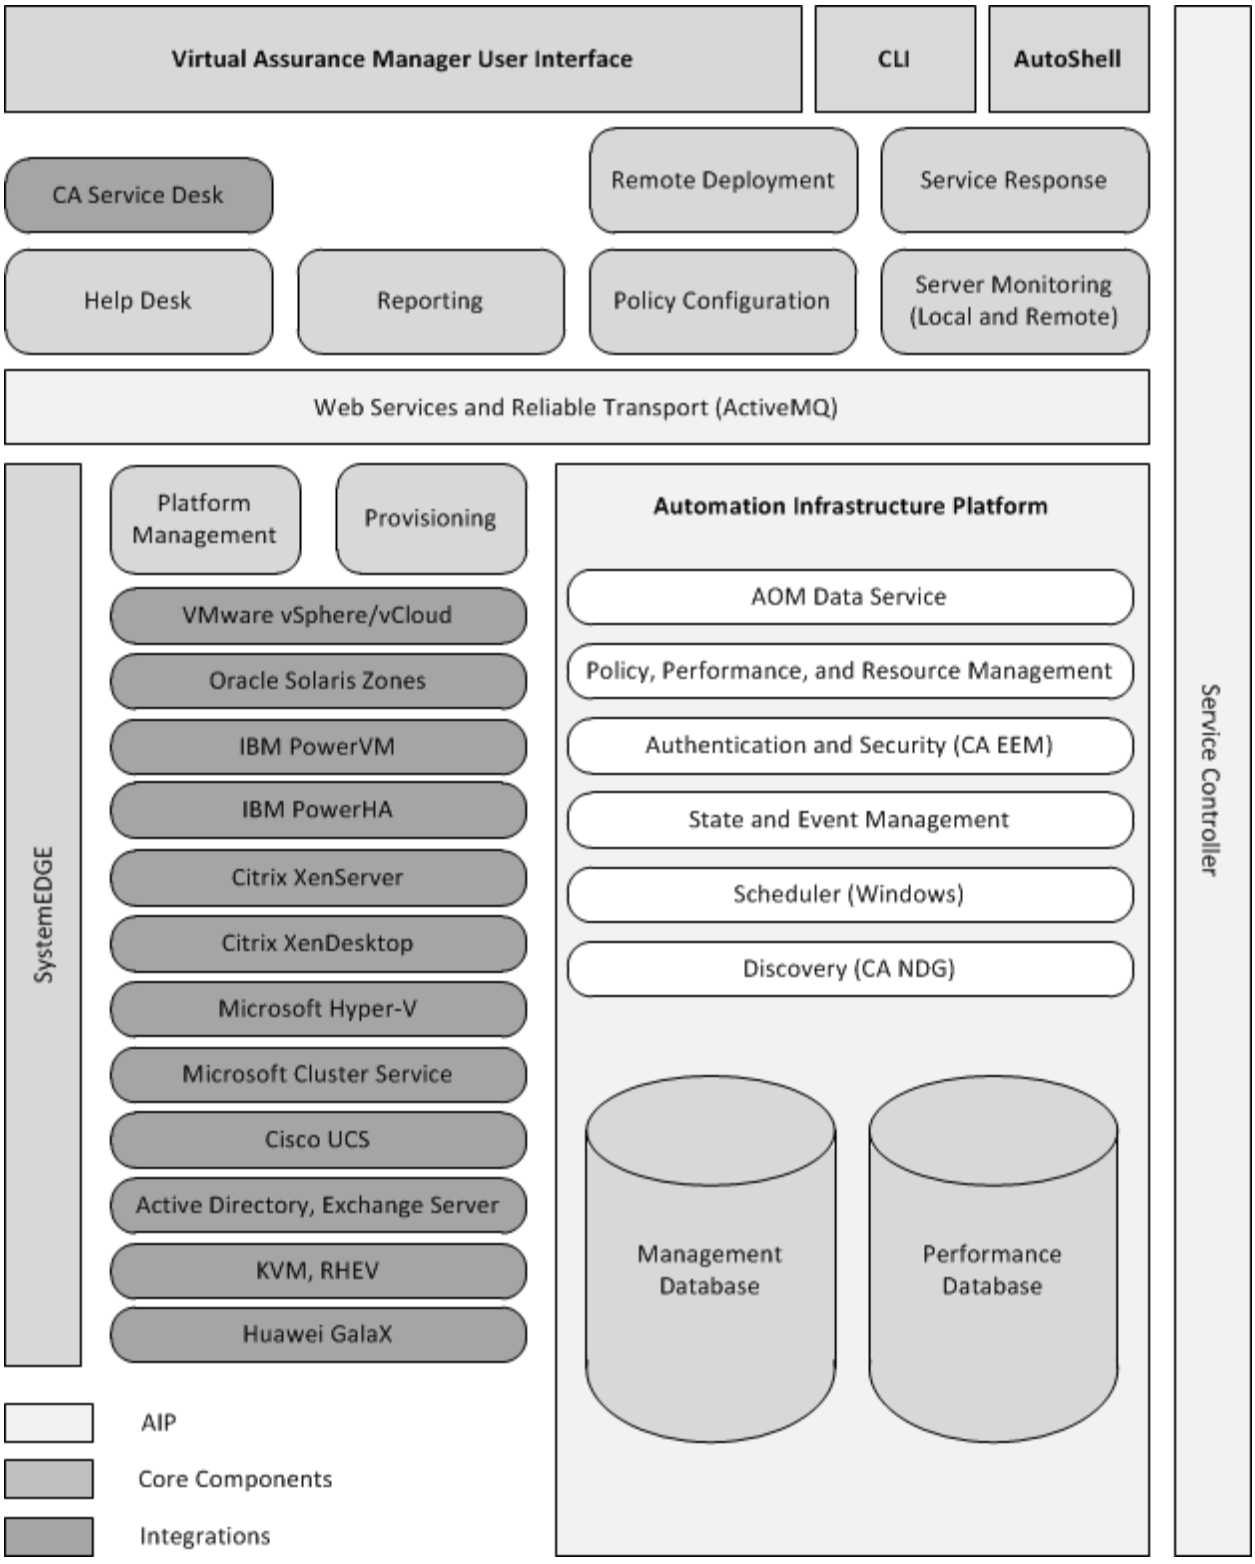

Each component registers with the service controller on startup. The service controller is a central component for identifying the location and status of all components. After a component is registered, it provides its Web Services Description Language (WSDL) location and publishes its events for other 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 EEM 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 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.

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:

- VMware vCenter Server integration for VM provisioning.
- Hyper-V integration for VM provisioning that is based on local templates on Hyper-V servers out of the box.
  - Integration with Microsoft System Center Virtual Machine Manger (SCVMM) for reuse of existing SCVMM image libraries.
- Solaris Zones integration for zones provisioning.
- Citrix XenServer integration for VM provisioning.
- VMware vCloud Director integration for VM provisioning.
- Red Hat Enterprise Virtualization integration for KVM provisioning.

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.

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. and locate the dpmutil set|get mgmtdb Command—Configure the Management Database.

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 -perfdb 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. To automate repetitive work, you can use the AutoShell interface.

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 EEM 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 access component user interfaces directly to perform more advanced functions. For example, you may want to open CA EEM to use native security. You can log in to the server where the component is installed to access its user interface directly. Alternatively, 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.

eHealth Integration Overview

As a system administrator, you use eHealth for performance management and reporting in your physical IT infrastructure. You want to extend these capabilities into managed physical and virtual server environments.

CA Virtual Assurance uses Application Insight Modules (AIMs) to manage and monitor physical and virtual system environments.

Note: The supported environments vary according to the versions of CA Virtual Assurance and eHealth. For more information, see the eHealth documentation.

The following diagram provides an overview of how system administrators set up CA Virtual Assurance to enable eHealth to monitor CA Virtual Assurance managed environments:

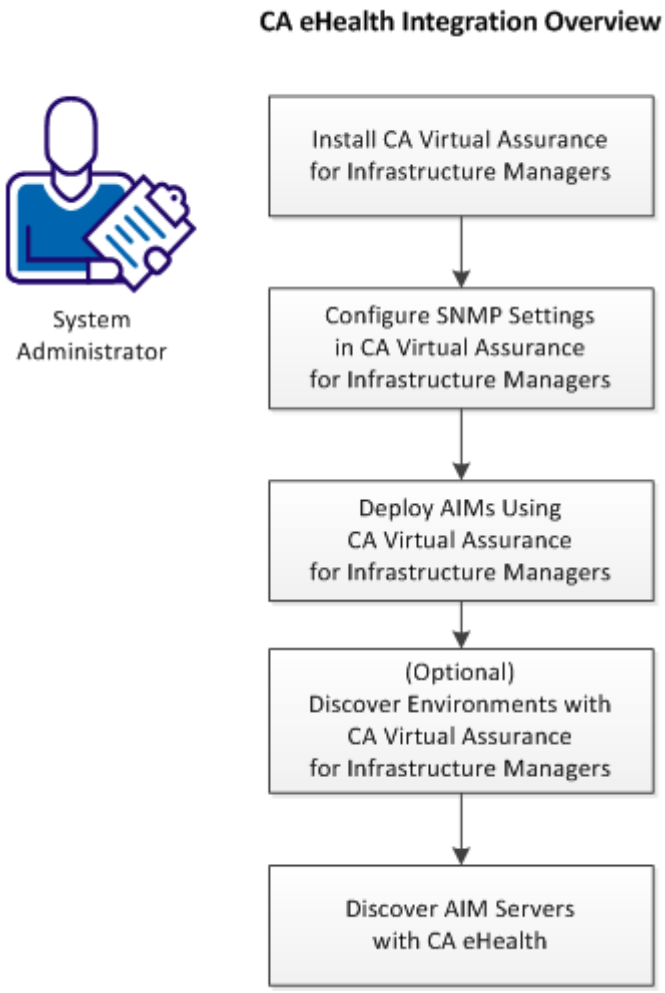

1. Install CA Virtual Assurance.

   Note: For more information, see the *CA Virtual Assurance Installation Guide*.

2. Configure SNMP in CA Virtual Assurance.

   Verify that the CA Virtual Assurance SNMP configuration is consistent with the eHealth port settings. If CA Virtual Assurance and eHealth are installed on different networks, verify that appropriate firewall ports are open.

   Note: For more information, see the *CA Virtual Assurance Administration Guide*.

3. Deploy the AIMs for your environments from CA Virtual Assurance.

   Use the appropriate AIMs and agent deployment method for your infrastructure and environment.

   Note: For more information, see the *CA Virtual Assurance Administration Guide*.

4. (Optional) Discover the environments with CA Virtual Assurance.

   Discover the components of the environment, to enable server management in CA Virtual Assurance.

   Note: For more information, see the *CA Virtual Assurance Administration Guide* and *Online Help*.

5. Discover the AIM servers with eHealth.

   The AIMs expose the infrastructures that they support to eHealth.

   Note: For more information, see the eHealth documentation.

eHealth displays the discovered infrastructure in its navigation panels, and provides performance dashboards and near real-time performance and availability reporting for the components of your environment.

Manage components of your environments and implement policy-based management using CA Virtual Assurance.

The following diagram represents an overview of the integrated infrastructure:

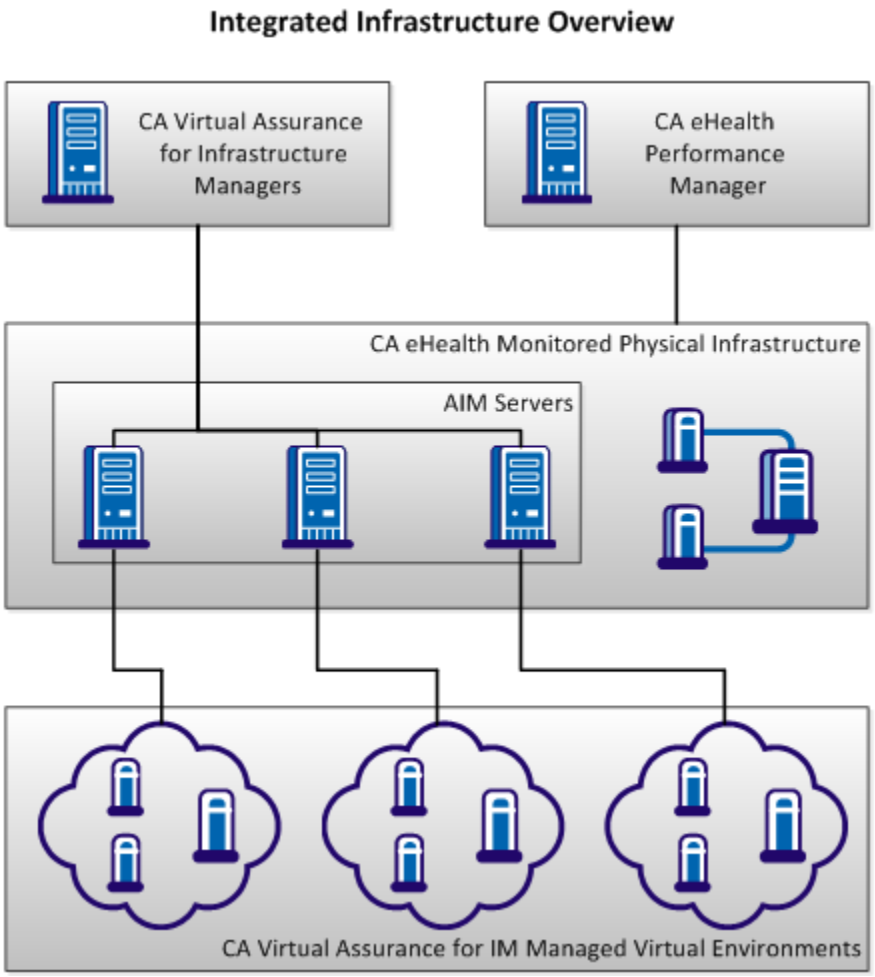

Spectrum Infrastructure Manager Integration Overview

As a system administrator, you use Spectrum Infrastructure Manager to manage and monitor your physical IT infrastructure, isolating faults to root cause. You want to extend these capabilities into virtual environments.

CA Virtual Assurance uses Application Insight Modules (AIMs) to manage and monitor physical and virtual system environments.

Note: The supported environments vary according to the versions of CA Virtual Assurance and Spectrum Infrastructure Manager. For more information, see the Spectrum Infrastructure Manager documentation.

The following diagram provides an overview of how system administrators set up CA Virtual Assurance to enable Spectrum Infrastructure Manager to model and monitor CA Virtual Assurance managed environments:

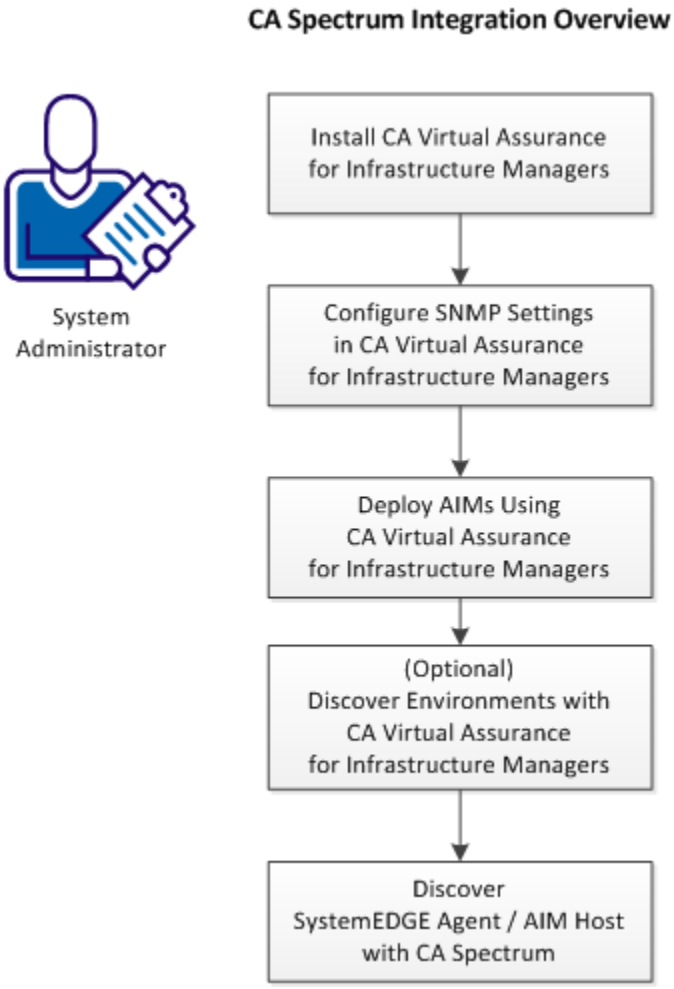

1. Install CA Virtual Assurance.

   Note: For more information, see the *CA Virtual Assurance Installation Guide*.

2. Configure SNMP in CA Virtual Assurance.

   Verify that the CA Virtual Assurance SNMP configuration is consistent with the Spectrum Infrastructure Manager port settings. If CA Virtual Assurance and Spectrum Infrastructure Manager are installed on different networks, verify that appropriate firewall ports are open.

   Note: For more information, see the *CA Virtual Assurance Administration Guide*.

3. Deploy the AIMs for your environments from CA Virtual Assurance.

   Use the appropriate AIMs and agent deployment method for your infrastructure and environment.

   Note: For more information, see the *CA Virtual Assurance Administration Guide*.

4. (Optional) Discover the environments with CA Virtual Assurance.

   Discover the components of the environment, to enable server management in CA Virtual Assurance.

   Note: For more information, see the *CA Virtual Assurance Administration Guide* and *Online Help*.

5. Discover the SystemEDGE agent / AIM host with Spectrum Infrastructure Manager.

   The AIMs expose the infrastructures that they support to Spectrum Infrastructure Manager.

   Note: For more information, see the Spectrum Infrastructure Manager documentation.

Spectrum Infrastructure Manager displays the discovered infrastructure in its OneClick navigation panels, provides infrastructure modeling, and availability reporting for the components of your environment.

Manage components of your environments and implement policy-based management using CA Virtual Assurance.

The following diagram represents an overview of the integrated infrastructure:

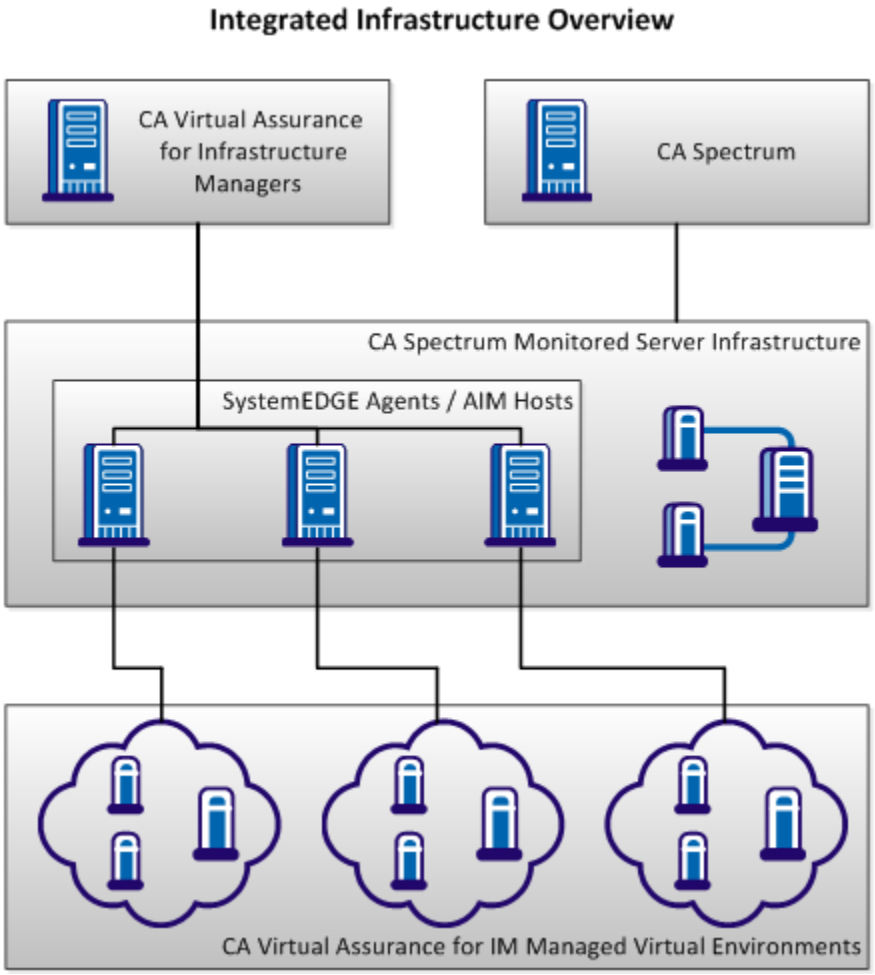

Chapter 3: Managing Users and User Groups

This section contains the following topics:

User Access Control

CA EEM 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*. Locate the documentation from Start, Programs, CA, Embedded Entitlements Manager, Documentation where CA EEM 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 EEM.

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 EEM Works with CA Virtual Assurance

CA EEM includes the following key objects:

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

CA EEM 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 EEM with CA Virtual Assurance.

Access the CA EEM User Interface

Log in to the CA EEM home page to use native security. The CA EEM 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 EEM user interface

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

   The CA EEM Log In window appears.

   Note: If you receive a security certificate request, bypass it and continue. To eliminate these messages, 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 EEM 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*.

To create CA EEM users

1. Click Manage Identities on the CA EEM 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 CA EEM 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 EEM administrator
- Database sa user (SQL authentication)

Change the CA EEM Administrator Password (EiamAdmin)

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

To change the administrator password (EiamAdmin) in CA EEM

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 EEM.

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 EEM 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 –mgmtdb`

   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 EEM and CA Virtual Assurance. If you change the password for sys_service in CA EEM later, also change it in CA Virtual Assurance to ensure that all CA Virtual Assurance services continue running.

To change the sys_service password in CA EEM

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 EEM.

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 EEM. 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 EEM 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 EEM.

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 EEM UI. Retry the dpmutil commands using valid credentials.

User Group Management

The User Group page provides access to user and user group authorization and user access control to product functions.

More information:

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.

Create a User Group

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.

More information:

Assign Users to Groups (see page 41)

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 39)

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.

Assign External Directory User Groups to User Groups

You can add user groups from an external directory to an existing CA Virtual Assurance user group when you want to grant existing access rights. The administrator user group is a predefined group and appears in the list by default.

To assign external directory user groups to user groups

1. Click Administration.

   The Administration page appears.

2. Click User Groups.

   The User Groups node appears on the left pane.

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

   The user group page appears on the right pane.

4. Select the Membership submenu.

   A series of membership panes appears below the tab.

5. Select User Group from the Identity list.

   The criteria for searching for users appear in the Attribute list.

6. Enter the user group name to add from the external directory in the Value text box, and click Search.

   If the user is located or a message notifies you that no match has been found, the user group appears in the Available User/User Group pane. User group name is identified by [Global Groups] in the Available user/user group list.

7. Select the user group to add from the Available User/User Group pane and click the right arrow  .

   The user group moves to the Selected User/User Group pane.

8. Click Save to finish adding user groups.

   Users are immediately granted the access privileges assigned to the user group with which they are associated.

Set User Group Privileges

You can use the Administration page to control user group access to services. Users that are granted administrator rights have access to all services.

Note: If you restrict a user group from a specific page, restrict also the user group from all actions on that page.

To set user group permissions

1. Click Administration.

   The Administration page appears.

2. Click User Group.

   The User Groups page appears.

3. Select the user group for which to set permissions, and click the Privileges tab.

   The Privileges page appears.

4. Click the check boxes for the tabs and actions to which you want to grant or restrict access, and click Save.

   The user group privileges are updated.

   Note: If you restrict a user group from a specific page, restrict also the user group from all actions on that page.

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.

Set User Group Permissions for Services

You can control user group access to services. Users with administrator rights have access to all services by default.

To set user group permissions

1. Click Administration.

   The Administration page appears.

2. Click User Group.

   The User Groups page appears.

3. Click a user group in the left pane.

4. Click Service Access.

   The right pane displays the resources for which services are enabled or disabled.

5. Enable or disable resources for service access as needed.

6. Click Save.

   The Service Access list is updated.

Set Run Command Script Privileges

You can grant or invoke access to a user group to an individual command script action by using the Administrator. The command script actions must already be created in the Actions & Rules page (Policy).

To set run command script action privileges

1. Click Administration.

   The Administration page appears.

2. Click User Group.

   The User Groups page appears.

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

   Tabs appear in the right pane.

4. Select the Privileges tab.

   A list of privileges appears, with check boxes to select or unselect privileges.

5. Expand the Policy folder, select Run Command Script, and click Save.

   The command script privileges are updated.

Import External Directories

You can import an external directory service that provides authentication of user names and passwords as a user group.

To import an external directory

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.

4. Select Membership.

   The Users/User Group page appears.

5. Select User Groups from the Identity drop-down list, type a name or partial name of an external directory in the Value text box, and click Search.

   A confirmation message notifies you if the search is unsuccessful or populates the Available User/User group section with the User Group that is found. The external directory has been imported to CA Virtual Assurance.

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.

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 *AIPAdmins*, 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.

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, 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.

Chapter 4: Managing Systems Performance

This section contains the following topics:

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 Systems Performance LiteAgent

Supported Virtual Environments and Operating Systems

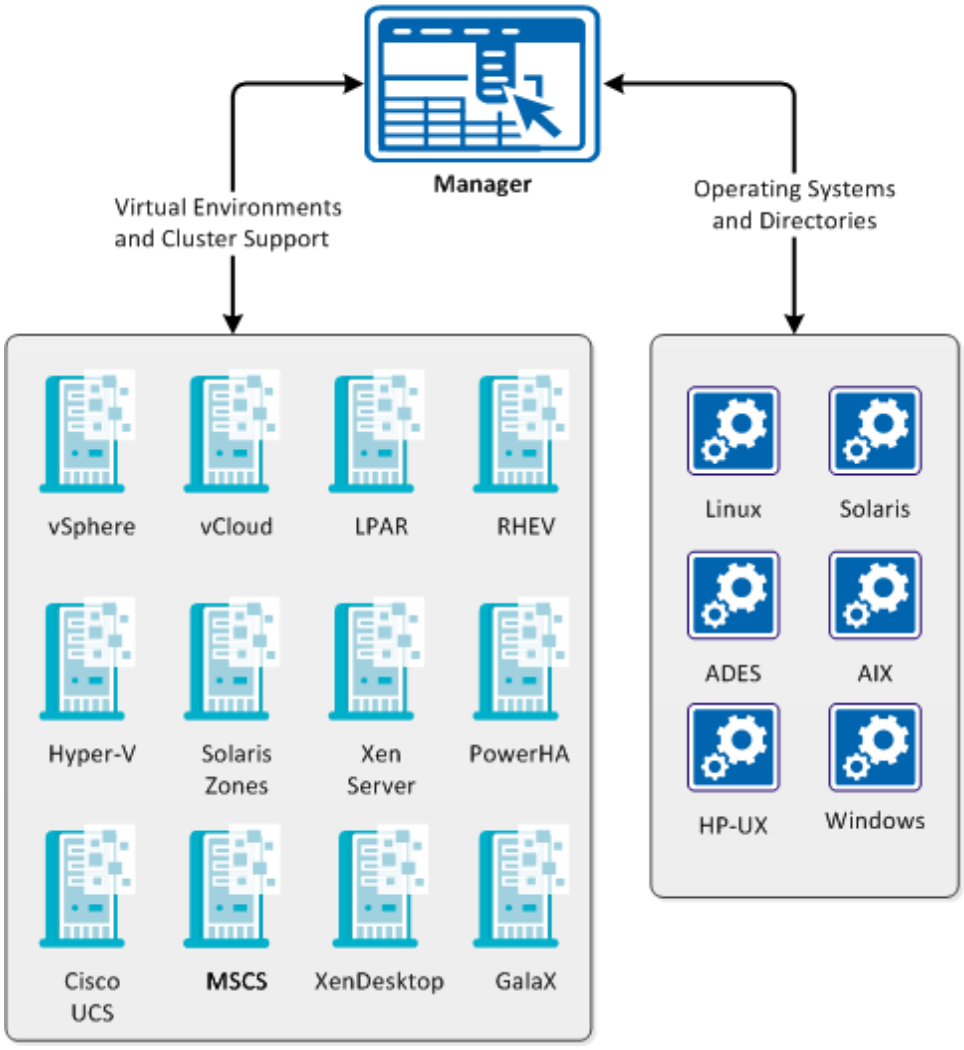

SystemEDGE is the base for systems management in CA Virtual Assurance and provides the following benefits:

- 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

More Information

Agent Configuration (see page 70)
Security and Maintenance (see page 73)
Agent Visualization (see page 77)
SystemEDGE Features (see page 62)
Service Response Monitoring (see page 74)
Interaction Between Remote Monitoring Components (see page 562)

Discovery

You can discover and add servers or entire subnets that you want to manage, including previously unmanaged or newly added servers.

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 discover, manage, or assign this system to a service.

To discover a system

1. Select Resources, Manage, Discover Server.
2. Complete the System Name field to specify 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.

Delete a System

Deleting a discovered system removes the system from CA Virtual Assurance.

To delete a system

1. Select Resources, Manage, Manage Systems.
2. From the drop-down menu on top of the right pane select one of the following:
   - Bare Metal Servers
   - Managed Servers
   - Unmanaged Servers

   A list of discovered systems appears.
3. Select the system or multiple systems to delete and select Delete from the Actions drop-down menu.

   Note: You can also delete all systems on the System page by selecting Delete All.

   A message prompts for confirmation.
4. Click Yes.

   The system is deleted from the System page and no longer appears as discovered in CA Virtual Assurance.

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.

Enhanced Discovery and SNMP Information

You can specify the credentials and SNMP information to discover a system or network.

To discover using enhanced information

1. Select Resources, Manage, Discover Network or Discover System.

   The Specify Discovery Type and Target section appears.

   After you enter the required details in the Discovery Type and Discovery Method section, click Next. The Enhanced Discovery and SNMP Information section appears.

2. Complete the following fields for the Enhanced Discovery section:

   Enhanced Discovery

   Select this option to specify enhanced credentials for discovery.

   Discovery Credentials

   Select *one* of the options to specify credentials.

   Specify Credentials

   Select this option to specify the credentials such as User Name and Password.

   Select Saved Credentials

   Select the existing saved credentials from the Available list.

3. Complete the following fields for the SNMP Information section, then click Next:

   Override SNMP Defaults

   Select this option to override the SNMP defaults for discovery.

   SNMP Settings

   Select *one* of the options to specify credentials.

   Specify Credentials

   Select this option to specify the credentials such as SNMP Version and Community String.

   Select Saved Credentials

   Select the existing saved credentials from the Available list.

4. Click Finish

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

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.

Rediscover a Network

If systems are added to a discovered network, or if the network has changed in other ways since the last discovery, you can rediscover that network .

To rediscover a network

1. Select Resources, Manage, Manage Discovered Networks.

   A list of discovered networks appears in in the right pane.

2. Select a network to rediscover and click (Rediscover) on the Network List toolbar.

   A discovery begins on the network. Any changes to the network (systems added or removed) are reflected when the discovery is complete.

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 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.

Services

You can group existing managed servers to a service and monitor this group.

More information:

Create a Service (see page 57)
Edit a Service (see page 58)
Remove Server from Services (see page 59)
Delete Services (see page 60)

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.

Remove Server from Services

You may not want a server to belong to a particular service anymore. You can remove servers from services.

To remove a server from a service

1. Click Resources.

   The Resources page appears.

2. Expand the Data Center folder and the CA Virtual Assurance Services folder in the Explore pane.

   The discovered and managed resources in the Data Center appear.

3. Select a server to be removed from service.

4. Click Remove from Service in the Quick Start tab.

   A message prompts you to confirm that you want to remove the server.

5. Click Yes.

   The server is removed from the service.

Delete Services

When you delete a service, its server collection is deleted, but the servers within the service remain managed within CA Virtual Assurance.

To delete a service

1. Select Resources, Manage, Manage Services.

   A list of services appears in in the right pane.

2. Select a service and click - (Delete) on the Services toolbar.

   A message prompts for confirmation.

3. Click Yes.

   The service is deleted.

Managed and Unmanaged Resources

You can specify whether to monitor a resource by changing its monitored state. If you change an object configuration to Unmanaged, the PMM processes the request and it 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 you change its configuration to Managed, the PMM processes the request and it 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.

Important! The managed or unmanaged status of a resource is different to the managed or unmanaged mode of SystemEDGE. Setting a Computer System to unmanaged enables SystemEDGE maintenance mode, if it is installed on that system.

Unmanage Managed Resources

You can stop managing currently managed servers.

Follow these steps:

1. Click Resources.

   The Resources page appears.

2. Expand the Data Center folder and the CA Virtual Assurance Services folder in the Explore pane.

   The discovered and managed resources in the data center appear.

3. Right-click the server and select Management, Unmanage.

   A message prompts you to confirm that you want to unmanage the server.

4. Click OK.

   The unmanaged server does not appear in the managed resource list. To view the unmanaged server, open the Unmanaged folder in Explore pane.

Manage Unmanaged Resources

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.

Delete Managed Resources

You can delete resources that you no longer want to manage.

To delete a managed resource

1. Click Resources.

   The Resources page appears.

2. Expand the Data Center folder and the CA Virtual Assurance Services folder in the Explore pane.

   The discovered and managed resources in the data center appear.

3. Right-click the server and select Management, Delete from System.

   A message prompts you to confirm deletion.

4. Click OK.

   The deleted server does not appear in the managed or unmanaged server list.

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

State Management Model (see page 66)
Managed Mode and Unmanaged Mode (see page 68)
Configure Object Aggregation (see page 223)
Stateless Monitoring (see page 67)
Systems Management MIB (see page 64)

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 (msgqueTable, shmemTable, 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.

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 Unmanaged 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.

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 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 Citrix XenDesktop

Provides capabilities to monitor your Citrix XenDesktop environment. This AIM can run on any Windows system where SystemEDGE is installed.

AIM for Citrix XenServer

Provides capabilities to monitor your Citrix XenServer environment. This AIM can run on any Windows system where SystemEDGE is installed. The AIM gets an entire view of all set-up XenServers and resource pools by communicating directly with the XenServers through XML RPC.

AIM for Active Directory and Exchange Server

Provides capabilities to monitor the Active Directory and Exchange Server on both off-premise and on-premise infrastructure. The AIM enables Domain and Exchange server Management, maintenance and upgrade.

AIM for Huawei GalaX

Provides capabilities to monitor your Huawei GalaX environment. This AIM can run on any Windows system where SystemEDGE is installed.

AIM for IBM PowerHA

Provides capabilities to monitor an IBM PowerHA, formerly known as High Availability Cluster Multiprocessing system.

AIM for IBM PowerVM (LPARs)

Provides capabilities to 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 server on which the AIM runs.

AIM for KVM

Provides capabilities to monitor your RHEV environment. This AIM can run on any Windows system where SystemEDGE is installed. The AIM communicates with the RHEV manager to get an entire view of all KVM servers that are registered with the manager.

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 Microsoft Hyper-V

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

AIM for Remote Monitoring

Provides capabilities to monitor remote Windows systems. Remote monitoring is also referred to as agent-less monitoring.

AIM for Service Response Monitor

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

AIM for Solaris Zones

Provides capabilities to monitor Solaris 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 server on which the AIM runs.

AIM for VMware vCenter Server

Provides capabilities to 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 VMware vCloud Director

Provides capabilities to monitor virtual systems that are under VMware vCloud Director control. This AIM can run on any Windows system where SystemEDGE is installed.

More Information

NodeCfgUtil Overview (see page 713)

Agent Configuration

The following 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

Configuration Overview (see page 165)
Perform a Point Agent Configuration (see page 71)

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 and existing self or process monitor.

Monitoring Software Settings

The Monitoring Software page lets you set non-policy related information for an individual server, server group, or service.

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Select a system or a service.

3. Click Monitoring Software.

   The Machine Details pane appears.

4. Modify the settings according to your needs, and click Apply:

   System Description

   Defines a description of the system.

   System Contact

   Defines a contact for the system.

   System Location

   Defines the system location.

   System Log Location

   Enables you to add a log location or the log file.

   Note: If you only specify the Log File name, the file is added to the port folder of SystemEdge.

   SystemEDGE Log Level

   Specifies the SystemEDGE log level.

   The settings are updated.

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 73)

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 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 disable the Maintenance Mode, and click Apply.

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

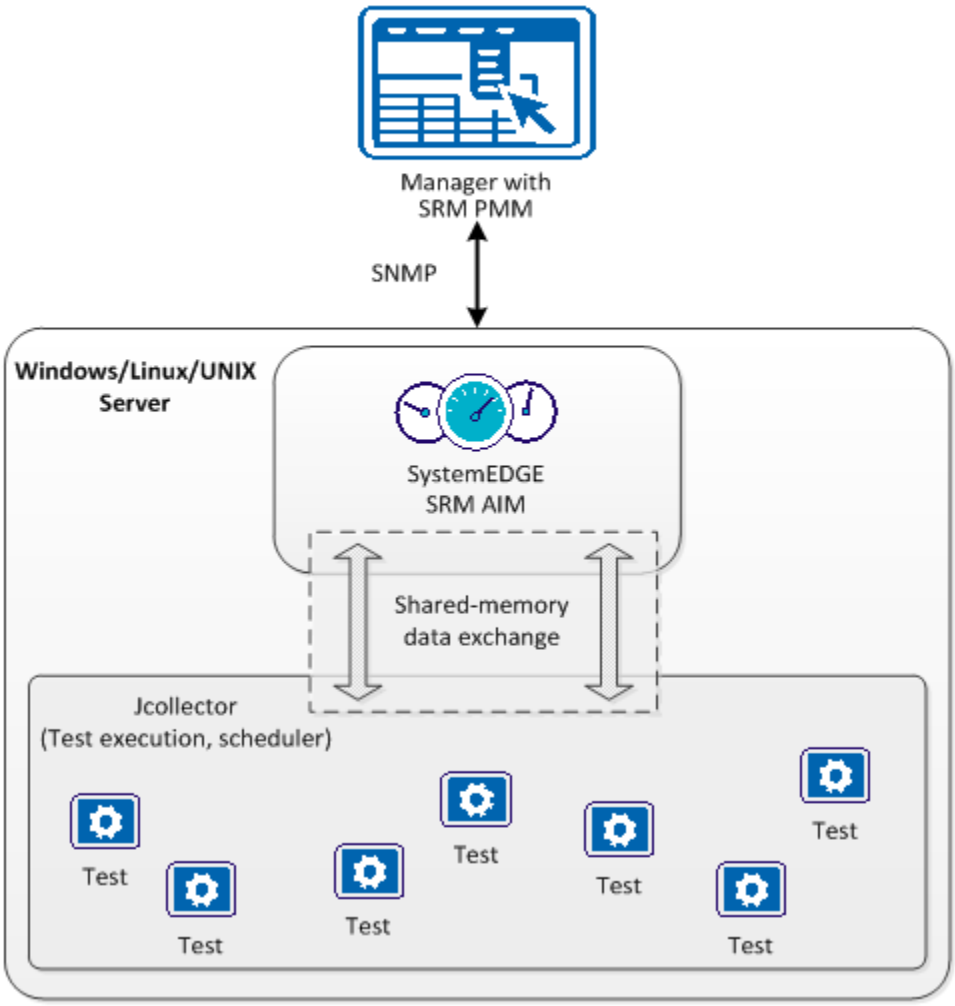

The svcrsp.cf configuration file contains the test specifications. The SRM AIM reads this configuration file and makes the test specifications available in the shared memory segment. The SRM Jcollector component reads each test configuration from the shared memory. Jcollector executes the tests, collects 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 75)

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.

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 78)
View Service Response Tests (see page 79)
View SystemEDGE Monitors (see page 77)

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 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

View Managed Object States

The CA Virtual Assurance user interface displays all SystemEDGE managed objects for managed systems.

To view managed object states, Click Resources, expand an appropriate folder in the Explore tree, and select a managed system on which SystemEDGE runs. The system Summary page appears in the right pane.

The System 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, In Maintenance, Destroy)
- Object class, instance, and attribute
- Current monitored value, operator, threshold value, and monitoring machine name.

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 Service Response Tests

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

To view Service Response tests

1. Click Resources, expand Managed, and select a system.

   The system Summary page appears in the right pane.

2. Click Details, and click Service Response.

   The Service Response Tests pane appears.

The Service Response 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

Chapter 5: Managing SystemEDGE and Application Insight Modules (AIMs)

This chapter explains how you can set up and can configure your monitoring software in your environment. This chapter also provides details about appropriate user permissions and how to change SystemEDGE to managed mode or to unmanaged mode.

This section contains the following topics:

User Permissions and Access Requirements Reference

The following sections summarize the access requirements to Install CA Virtual Assurance components and monitor your environments using CA Virtual Assurance. Each section includes information about the required communication ports. If a distributed installation ranges across firewalls, you can use this list to verify that the required communication ports are open.

This documentation is intended for:

- Administrators who install, configure, and use CA Virtual Assurance to manage virtual environments.
- Operators who use CA Virtual Assurance to monitor virtual environments.

Active Directory and Exchange Server (ADES)

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

(Exchange 2007) Requires Domain Administrator or Exchange Administrator role.

(Exchange 2010) Requires Exchange Organization Management role.

Communication Ports

PowerShell Ports: 80, 443, 5985, and 5986

ADSI Ports: 3268, 389

Cisco UCS

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Requires UCS Manager user account with sufficient privileges to run the following UCS operations: blade power operations, service profile operations, pool operations, policy operations, Import/Export operations.

Note: We recommend giving the UCS user admin privilege.

If admin privilege cannot be given, assign the following roles to the UCS user:

- Ext-lan-config
- Ext-san-config
- Service-profile-config
- Service-profile-config-policy
- Service-profile-ext-access
- Service-profile-network
- Service-profile-network-policy
- Service-profile-qos

- Service-profile-qos-policy
- Service-profile-security
- Service-profile-security-policy
- Service-profile-server
- Service-profile-server-oper
- Service-profile-server-policy
- Service-profile-storage
- Service-profile-storage-policy
- Operations
- Server-equipment

Communication Ports

HTTP Port: 80

HTTPS Port: 443

Citrix XenDesktop

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Citrix XenDesktop version 5.6 requires Active Directory account with at least a read-only administrator role in XenDesktop.

Communication Ports

WinRM Port: 5985, 5986

SNMP Port: 161

WMI Port: 135

Citrix XenServer

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

(XenServer 6.0 and higher) Requires root or Active Directory subject with a read-only role.

Communication Ports

HTTPS Port: 443

SNMP Port: 161

Huawei GalaX

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Huawei GalaX monitoring requires administrator user credentials and a corresponding p12 file from the GalaX environment.

Communication Port

HTTP Port: 8773

Hyper-V

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Requires local Administrator Account.

SCVMM Monitoring

Requires System Center Virtual Machine Monitoring (SCVMM) Administrator role.

Communication Port

Windows RPC Endpoint Mapper Port: 135

DCOM/WMI Port: dynamically assigned during RPC Endpoint negotiation

For more information, see The default dynamic port range for TCP/IP http://support.microsoft.com/kb/929851.

IBM PowerHA

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring PowerHA

Requires an account with the rights to execute the following CLI commands:

- clstat
- clRGinfo -s
- cldump
- cllsnw
- cltopinfo
- cllsif
- clshowsrv -v
- vmstat

Communication Port

Secure Shell TCP Port: 22

IBM PowerVM

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

The following requirements depend on your existing environment that you want to monitor with CA Virtual Assurance:

Monitoring Hardware Management Console (HMC)

Requires the hmcsuperadmin task role account. We recommend defining users whose resource role only includes the P-Servers you want them to manage.

Note: HMC monitoring requires *both* HMC *and* VIOS configuration.

Monitoring Virtual IO Server (VIOS)

Requires the padmin user account on VIOS that you want to monitor.

Monitoring Integrated Virtualization Manager (IVM)

Requires the padmin user account on the IVM that you want to monitor.

Note: IVM Monitoring requires IVM configuration.

Communication Port

Secure Shell TCP Port: 22

Microsoft Cluster Server

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Requires a domain administrator account or a cluster node local account. If a domain user is used, it must be in the domain administrators group. If a cluster node local account is used, the user must be a member of the administrators group.

Important! Set up the same cluster node local credentials on all nodes. If the cluster service is moved to a different node and the node has different credentials, the AIM is unable to connect.

Communication Port

Windows RPC Endpoint Mapper Port: 135

DCOM/WMI Port: dynamically assigned during RPC Endpoint negotiation

For more information, see The default dynamic port range for TCP/IP http://support.microsoft.com/kb/929851.

Oracle Solaris Zones

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Requires root user access.

Communication Port

Secure Shell TCP Port: 22

Zone CPU Cap

CA Virtual Assurance monitors the restriction of the CPU utilization to the defined percentage in the cpu cap resource control.

Solaris 11

CA Virtual Assurance monitors Solaris 11 zones.

Red Hat Enterprise Virtualization

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Requires a corresponding Red Hat Enterprise Administrator Role with Super User privileges.

Note: You can use the Microsoft Active Directory (AD) user or the Red Hat Enterprise IPA user.

Communication Port

REST API Port: 8443

Remote Deployment Agent

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

Installation on Windows

Requires Windows System Administrator privileges.

Installation on Linux

Requires root access or use of sudo or pfexec.

Cross-platform Remote Deployment

Uses Infrastructure Deployment (ID).

ID manager component

(Windows targets) Requires mapping of a windows share Admin$ on target machines.

(Unix or Linux targets) Requires SSH connection between the manager and target to be successful.

Communication Ports for Remote Deployment (Windows)

CIFS UDP Port: 137 (Inbound/Outbound)

CIFS UDP Port: 138 (Inbound/Outbound)

TCP Port: 135 (Inbound)

CIFS TCP Port: 139 (Inbound/Outbound)

CIFS TCP Port: 445 (Inbound/Outbound)

CAM UDP Port: 4104 (Inbound/Outbound)

CAM TCP Port: 4105 (Configurable)

Communication Ports for Remote Deployment (UNIX, Linux)

CAM UDP Port: 4104 (Inbound/Outbound)

Secure Shell TCP Port: 22 (Inbound)

TCP Port: 135 (Inbound)

CAM TCP Port: 4105 (Configurable)

Remote Monitoring

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Remote Monitoring

Requires credentials with access to Windows Management Instrumentation (WMI).

Communication Port

Windows RPC Endpoint Mapper Port: 135

DCOM/WMI Port: dynamically assigned during RPC Endpoint negotiation

For more information, see The default dynamic port range for TCP/IP http://support.microsoft.com/kb/929851.

As a best practice the Remote Monitoring systems must be a member of an AD Domain. 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 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.

SystemEDGE and Advanced Encryption

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Installation on Linux or Unix

Requires root access or use of sudo configuration for a nonprivileged user account.

Monitoring

Permission to edit and load the cf file, or permission to use Remote Configuration.

Communication Ports

UDP Port: 161 (SNMP Get/Set Requests); alternative port: 1691

UDP Trap Port: 162 (Outbound)

SystemEDGE in Managed Mode uses CAM:

CAM UDP Port: 4104

CAM TCP Port: 4105

VMware vCenter

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

(Valid for AIM) Requires read-only user access in VC for the AIM component.

(Valid for PMM) Requires the set of privileges that are specified for vSphere vCenter Server.

Important! The user role must match with type of operation that is being performed otherwise operation does not work.

Communication Port

HTTPS Port: 443

VMware vCloud

This section lists required access permissions to install and monitor your environments and the required communication ports. Verify that the listed communication ports are open.

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

(Valid for AIM) Requires the System Administrator role.

(Valid for VMware WS) Operations are limited to the user role.

System Administrator@System

Grants full access.

Organizational Access@org_name

Limits operation at the organization level and the role assignment.

Communication Port

REST API Port: 8443

How to Configure SNMP and Access Control Lists

This section explains the differences between global and server-level SNMP settings, how to apply Access Control Lists, and how to configure SNMP to discover systems successfully.

More information:

SNMP Consistency (see page 93)
Global and Server-level SNMP Settings (see page 94)
How to Configure SNMPv1/v2 Settings and Access Control Lists (see page 96)
How to Manage Server-level SNMP Settings (see page 107)
How to Configure SNMPv3 (see page 111)
Configure CA Virtual Assurance to Forward Events (see page 117)

SNMP Consistency

Consistent SNMP settings are necessary for discovering systems and networks correctly. If none of the SNMP settings of SystemEDGE on a remote system exists on the CA Virtual Assurance manager, CA Virtual Assurance cannot discover the required resources on that system. CA Virtual Assurance requires at least valid read-only SNMP credentials to discover a system.

If you remotely deploy the SystemEDGE agents and you configure the agents through Policy Configuration, the SNMP consistency condition between manager and managed systems is automatically fulfilled.

If you configure the SNMP settings locally on the remote server, verify the consistency of the SNMP settings. If none of the SNMP settings on the remote server is specified on the manager, specify the missing credentials as a global SNMP object in CA Virtual Assurance and discover the remote system.

The SNMP scenarios and procedures in this chapter assume that the SystemEDGE agents run in managed mode. In managed mode, SystemEDGE is configured through Policy Configuration in CA Virtual Assurance.

More information:

Global and Server-level SNMP Settings (see page 94)

Global and Server-level SNMP Settings

Categories like server-level or global SNMP settings only exist on the CA Virtual Assurance manager. Policy Configuration delivers a collection of these settings through a policy to managed servers. These SNMP settings finally appear in the sysedge.cf configuration file on each of the managed target servers. SystemEDGE does not distinguish between server-level or global SNMP settings. This information is stored on the manager only. The manager knows which version of the policy has been applied to a particular managed server.

If necessary, you can add your own global or server-level SNMP settings to the CA Virtual Assurance manager.

In most cases, the global SNMP settings mechanism provides you the appropriate flexibility to manage the SNMP settings on your servers. In specific cases, it can be necessary to use server-level SNMP settings. Policy Configuration provides you the full flexibility when you create the collection of SNMP settings for your policy. You can select global or server-level SNMP settings as appropriate.

Global SNMP settings populate the drop-down lists for the following fields in the SystemEDGE package wrapper for Remote Deployment:

- Port
- Read Community
- Read-Write Community

Alternatively, you can edit the fields inline.

The available SNMPv1 community strings depend on the port setting. When you select the port number first, then you get automatically the valid community strings in the drop-down lists for that port. If no credentials are specified in the package wrapper, the installer defaults to the read-only string of public. The credentials in the package wrapper are valid from the point SystemEDGE is installed, until SystemEDGE on the managed server registers with Policy Configuration to enter the managed mode. SystemEDGE loads the settings from the policy.

Note: SystemEDGE requires at least one SNMPv1 community for its installation. After CA Virtual Assurance has discovered SystemEDGE on the server, CA Virtual Assurance can treat these SNMPv1 settings as server-level SNMP settings.

The following option under Administration, Configuration, Deployment & Configuration controls whether the SNMP settings in the package wrapper become server-level SNMP settings or not.

- Create server-specific SNMP settings when a SystemEDGE Agent registers.

If this option is enabled, CA Virtual Assurance uses the SNMPv1 settings for the installation as server-level SNMP credentials.

For each Remote Deployment job, you can specify a policy that is applied to the target systems during the deployment process. If you do not specify a particular policy, CA Virtual Assurance uses the SystemEDGE default policy. If you have defined multiple SystemEDGE policies, you can determine the default policy in the SystemEDGE Policies pane from the list of existing policies.

More information:

How to Configure SNMPv1/v2 Settings and Access Control Lists (see page 96)
How to Configure SNMPv3 (see page 111)
How to Manage Server-level SNMP Settings (see page 107)

How to Configure SNMPv1/v2 Settings and Access Control Lists

The following diagram provides an overview of the required actions when you specify SNMP settings for your environment. The diagram includes strategies for common and exceptional cases. Exceptional cases are indicated as optional in the diagram.

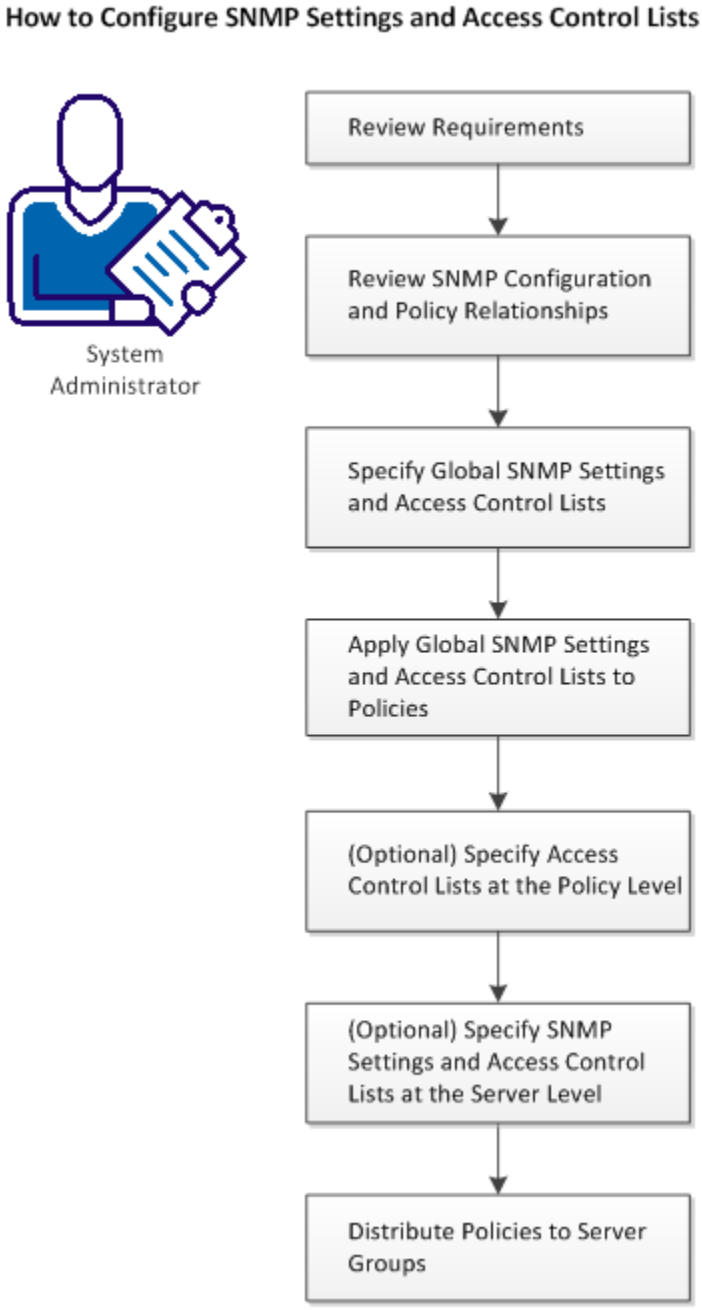

Review Requirements (SNMPv1/2)

Review the following requirements before you start configuring the SNMP settings on CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, and Windows Server operating systems.
- You have a basic understanding of CA SystemEDGE.
- You can access a CA Virtual Assurance manager installation that includes the Monitoring Agent (CA SystemEDGE).
- You can access the monitoring agents (CA SystemEDGE) on managed nodes.
- You can access the CA Virtual Assurance user interface.
- CA Virtual Assurance has discovered all relevant systems.
- SystemEDGE runs in managed mode on all systems that you want to configure.

More information:

Review SNMP Configuration and Policy Relationships (see page 97)

Review SNMP Configuration and Policy Relationships

An SNMP Settings object for SNMPv1/v2 comprises of a name, the community string, the type of operation (read-only or read-write), SNMP version, port, timeout, retry limit, and Access Control List (ACL).

An ACL specifies a list of manager systems for a group of managed systems on which SystemEDGE runs. The CA Virtual Assurance manager distributes SNMP settings and ACLs through Policy Configuration to the managed systems. These managed systems accept SNMP requests only from the manager systems listed in the ACL. If no ACL is specified, the managed systems accept SNMP requests from any system.

If ACLs are defined, the CA Virtual Assurance manager is also automatically added to the list of ACLs. The CA Virtual Assurance manager always has connectivity.

In most cases, the same SNMP credentials are used across many or all systems. To manage and apply those credentials appropriately, you can specify the SNMP credentials and ACLs at the global level. Consistent SNMP credentials and ACL settings on manager and agent systems are necessary to discover and manage systems properly. You specify global SNMP Settings objects under Administration, Configuration, SNMP.

In exceptional cases, you can add ACLs at the policy level or you can specify SNMP credentials and ACLs entirely at the system level. If you want to modify SNMP settings at the system level, change the settings for each affected system.

Only those SNMP settings are applied to a target system, which use the same port as the target system.

The following diagram illustrates the policy architecture:

Policy Architecture

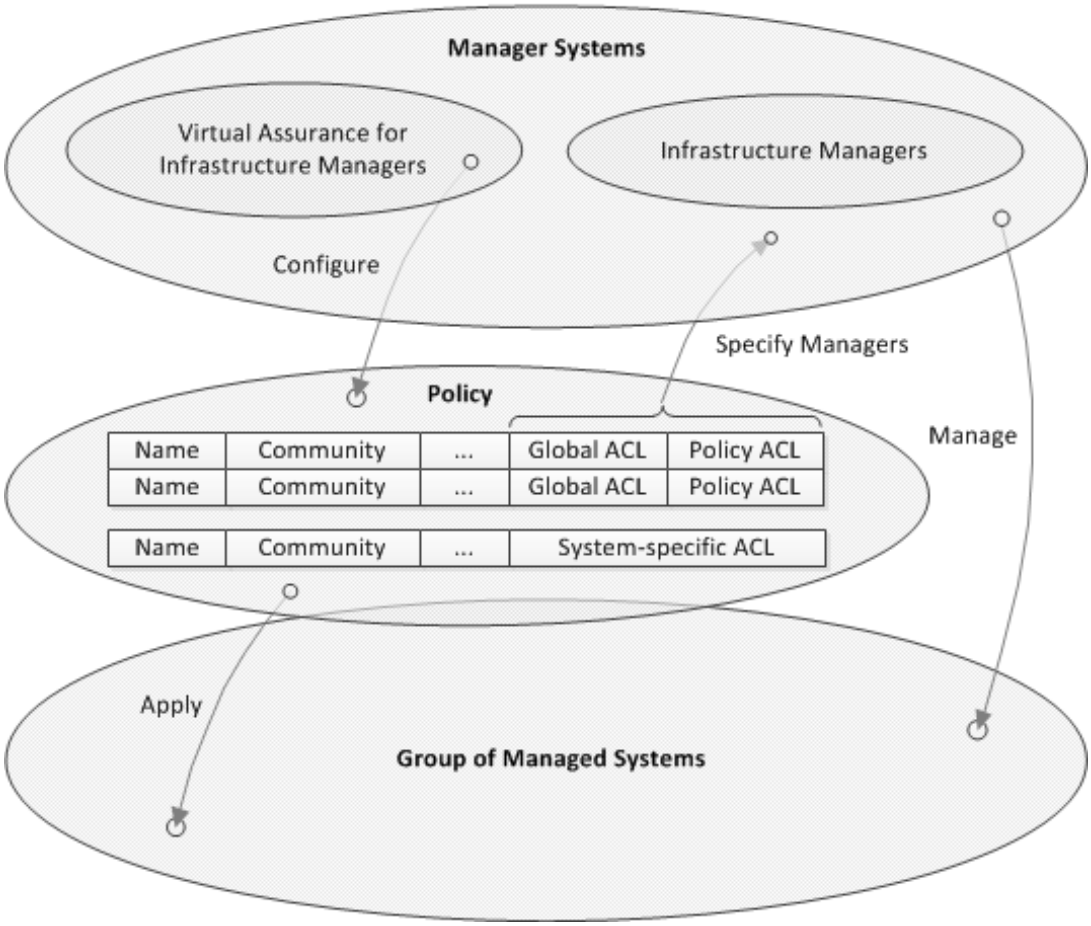

You can configure SNMP settings at the global, policy, or system level and you can assign these settings to a policy (upper left arrow). The policy can be applied through CA Virtual Assurance to a group of managed systems. The Access Control Lists (ACLs) specify the names of the manager systems which manage the group of managed systems. If you add all required manager systems into an ACL, the managed systems respond only to SNMP requests from these managers.

More information:

Specify Global SNMP Settings and Access Control Lists (see page 99)
Apply Global SNMP Settings and Access Control Lists to Policies (see page 100)
(Optional) Specify Access Control Lists at the Policy Level (see page 101)
Distribute Policies to Server Groups (see page 103)
(Optional) Specify SNMP Settings and Access Control Lists at the Server Level (see page 102)
Example for Three Server Groups (see page 104)

Specify Global SNMP Settings and Access Control Lists

You can specify Access Control Lists (ACLs) for SystemEDGE SNMP credentials at the global, policy, and system level. You can minimize the dependency on system-specific SNMP objects by associating the ACLs to global SNMP objects.

Edit ACLs at the global level using Administration, Configuration, SNMP from the user interface.

Follow these steps:

1. Navigate to Administration, SNMP in the user interface.

   The SNMP Settings page appears.

2. (Optional) Click Actions, New to create an SNMP Settings object.

   The New SNMP Settings dialog appears. An SNMP Settings object consists of a name, the community string, the type of operation (read-only or read-write), SNMP version, port, timeout, retry limit, and Access Control List (ACL). Use the port number that is specified on the managed nodes to which you want to apply the SNMP settings.

3. Specify the required data to create an SNMP Settings object, and click OK.

4. Select the SNMP Settings object to which you want to add an ACL and click the Edit icon.

   The Edit SNMP Settings dialog with ACL panel appears.

5. Specify the names or IP addresses of the manager systems under the Policy Configuration SystemEDGE Access Control List panel and click OK.

   The ACL for a particular global SNMP Settings object is specified.

You can apply the global SNMP Settings object with its Access Control List to policies.

More information:

Apply Global SNMP Settings and Access Control Lists to Policies (see page 100)
(Optional) Specify Access Control Lists at the Policy Level (see page 101)
Distribute Policies to Server Groups (see page 103)
(Optional) Specify SNMP Settings and Access Control Lists at the Server Level (see page 102)

Apply Global SNMP Settings and Access Control Lists to Policies

After you completed your global SNMP settings with appropriate ACLs, apply the SNMP settings to policies.

Follow these steps:

1. Navigate to Resources, Configuration in the user interface.

   The Policy page appears.

2. Expand Policy, Policies, SystemEDGE in the navigation pane.

   The SystemEDGE page appears listing the available policies.

3. (Optional) Click  to create a policy.

   The New SystemEDGE Policy dialog appears.

4. Specify the required data to create a policy and click OK.

5. Open the policy that you want to apply to one or more managed systems and click Traps & Communities.

   The Communities page appears with the table of SNMP settings and the following options:

   - Include only Server Communities
   - Include Server Communities and all Default Communities (Global Communities)
   - Custom Selection

   Note: The only default (global) SNMP settings from the table that are included in the configuration are those settings with a port that matches the agent port.

6. Select one of the three options and verify that you have at least one community with the appropriate port specified for each of your target systems.

   If you select the first option *Include only Server Communities,* verify that appropriate server-level SNMP settings exist for the target system. The available server communities which you can select are generic:

   - Server Read
   - Server Write

   They represent the existing read and write credentials at the server level.

If you select the second option, all global SNMP settings from the table and server-level settings are applied to the target systems.

If you select the third option, only the selected SNMP settings from the table are applied to the target systems. This option allows you to select global settings only.

7. Click Save Policy.

You can distribute the policy to the appropriate server group or specify additional ACLs at the policy or server level if necessary.

More information:

(Optional) Specify Access Control Lists at the Policy Level (see page 101)
Distribute Policies to Server Groups (see page 103)
(Optional) Specify SNMP Settings and Access Control Lists at the Server Level (see page 102)

(Optional) Specify Access Control Lists at the Policy Level

After you specified your global SNMP settings with optional global ACLs for a policy, you can define ACLs at the policy level.

Follow these steps:

1. From the policy page, select the second or third option, to apply global SNMP settings from the table:
   - Include Server Communities and all Default Communities (Global Communities)
   - Custom Selection
2. Select a global SNMP Settings object from the table and click the *View* or *None Defined* link.

   The ACL dialog opens.
3. Add the names or IP addresses of the manager systems into the *Policy-specific SNMP Access Control Lists* field and click OK.

   The servers of the server group, to which you want to apply this policy, accept SNMP requests from these manager systems.
4. Click Save Policy.

You can distribute the policy to the appropriate server group or specify additional ACLs at the system level if necessary.

More information:

Distribute Policies to Server Groups (see page 103)
(Optional) Specify SNMP Settings and Access Control Lists at the Server Level (see page 102)

(Optional) Specify SNMP Settings and Access Control Lists at the Server Level

In exceptional cases, you can specify SNMP settings and an Access Control List for particular managed systems.

Follow these steps:

1. Navigate to Resources, Explore in the user interface.

   The Explore pane appears.

2. Expand the Explore tree and right-click the system for which you want to specify SNMP credentials and an Access Control List.

3. Select Policy, Configure SNMP Settings from the pop-up menu.

   The SNMP Settings dialog lists the valid SNMP settings for that system.

4. Click Add.

   The New SNMP Settings dialog appears.

5. Specify Name, Port, Community String, the type of operation (read-only or read-write), SNMP version, port, timeout, and retry limit. Use the port number of the installed SystemEDGE on the server. Click OK.

6. Close the dialogs, change to the selected system page, and click the Monitoring Software, SNMP Access Control tab.

   The specified system-specific SNMP Community Settings are listed.

7. Click the Edit link from the Access Control List column.

   The system-specific Access Control List dialog appears.

8. Enter manager system names into the SNMP Access Control List field and click OK.

   The managed system accepts SNMP requests from manager systems that you list in the ACL.

9. Click Save.

Distribute the policy to the appropriate server group.

More information:

Distribute Policies to Server Groups (see page 103)

Distribute Policies to Server Groups

After you completed your SNMP settings with appropriate ACLs, apply the policy to systems in the network.

Follow these steps:

1. Navigate to Resources, Configuration in the user interface.

   The Policy page appears.

2. Expand Policy, Policies, SystemEDGE in the navigation pane.

   The SystemEDGE page appears listing the available policies.

3. Select the policy that you have previously saved with appropriate SNMP settings.

   The policy page opens.

   Note: If you do not want to apply existing server-level SNMP settings and ACLs through Policy Configuration to managed systems, clear the Server Read and Server Write entries in the Server Communities pane.

4. Click Action, Apply.

   The Select Machines page appears.

5. Select all systems which you want to configure with that policy, and click Apply.

   You can view the delivery status or return to the Policy page.

   The new settings are applied to the target systems.

More information:

Example for Three Server Groups (see page 104)

Example for Three Server Groups

The following example illustrates a use case that consists of three server groups, global SNMP settings, and ACLs specified at the global and policy level.

The datacenter consists of the following server groups:

- Infrastructure Manager Servers: CA Virtual Assurance system, SQL Server systems, CA EEM system, one or more distribution servers, three infrastructure manager systems (im1.ca.com, im2.ca.com, im3.ca.com). These systems are managed through va-im.ca.com, im1.ca.com.
- Sales Servers: All servers that belong to the Sales department, managed through va-im.ca.com, im1.ca.com, im2.ca.com.
- Development Servers: All servers that belong to the development department, managed through va-im.ca.com, im1.ca.com, im3.ca.com.

Server Group	Global Community Settings	Global Access Control Lists	Policy Level Access Control Lists
Infrastructure Manager Servers	_public_	va-im.ca.com, im1.ca.com	-
	admin	va-im.ca.com, im1.ca.com	-
Sales Servers	_public_	va-im.ca.com, im1.ca.com	im2.ca.com
	admin	va-im.ca.com, im1.ca.com	im2.ca.com
Development Servers	_public_	va-im.ca.com, im1.ca.com	im3.ca.com
	admin	va-im.ca.com, im1.ca.com	im3.ca.com

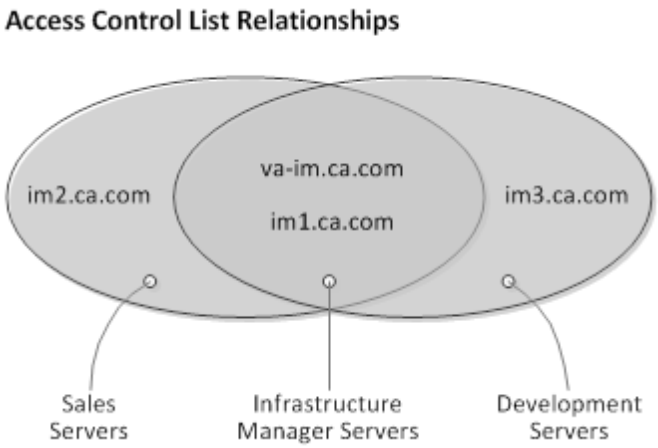

Follow these steps:

1. Specify the following global SNMP objects under Administration, SNMP:

   infrastructure-read: port 161, read-only access, community _public_, ACL: va-im.ca.com, im1.ca.com

   infrastructure-write: port 161, read-write access, community _admin_, ACL: va-im.ca.com, im1.ca.com

   sales-read: port 161, read-only access, community _public_, ACL: va-im.ca.com, im1.ca.com

   sales-write: port 161, read-write access, community _admin_, ACL: va-im.ca.com, im1.ca.com

   development-read: port 161, read-only access, community _public_, ACL: va-im.ca.com, im1.ca.com

   development-write: port 161, read-write access, community _admin_, ACL: va-im.ca.com, im1.ca.com

2. Create three policies (one for each server group) that are based on the default policy: infrastructure, sales, and development
3. Change to the infrastructure policy page, select the third option to apply global SNMP settings from the table:
   - Custom Selection
4. Add infrastructure-read and infrastructure-write global SNMP objects to the infrastructure policy.
5. Save the policy.

6. Change to the sales policy page, select the third option to apply global SNMP settings from the table:
   - Custom Selection
7. Add sales-read and sales-write global SNMP objects to the sales policy.
8. For the sales-read and sales-write objects, click the View links.

   The corresponding ACL dialog opens.
9. Add im2.ca.com to the sales-read and sales-write objects (Policy-specific SNMP Access Control List) and click OK.
10. Save the policy.
11. Change to the development policy page, select the third option to apply global SNMP settings from the table:
    - Custom Selection
12. Add development-read and development-write global SNMP objects to the development policy.
13. For the development-read and development-write objects, click the corresponding View link.

    The corresponding ACL dialog opens.
14. Add im3.ca.com to the development-read and development-write objects and click OK.
15. Save the policy.
16. Apply each policy (infrastructure, sales, development) to its associated server group.

More information:

Review SNMP Configuration and Policy Relationships (see page 97)

How to Manage Server-level SNMP Settings

The following diagram provides an overview of the required actions when you want to manage server-level SNMP settings.

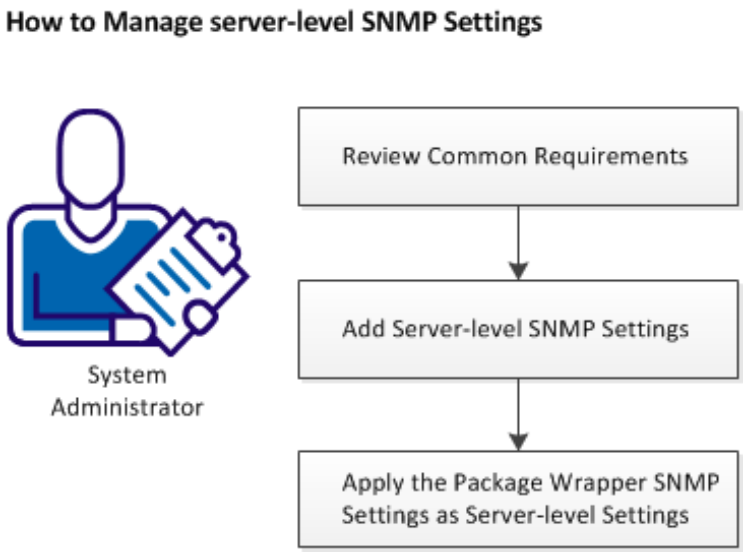

Follow these steps:

Review Requirements (Server-level) (see page 107)
Add Server-level SNMP Settings (see page 108)
Apply the Package Wrapper SNMP Settings as Server-level Settings (see page 109)

Review Requirements (Server-level)

Review the following requirements before you start managing the server-level SNMP settings on CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, and Windows Server operating systems.
- You have a basic understanding of CA SystemEDGE.
- You have read the How to Configure SNMPv1/v2 Settings and Access Control Lists scenario.
- You can access a CA Virtual Assurance manager installation that includes the Monitoring Agent (CA SystemEDGE).
- You can access the monitoring agents (CA SystemEDGE) on managed nodes.
- You can access the CA Virtual Assurance user interface.
- CA Virtual Assurance has discovered all relevant systems.
- SystemEDGE runs in managed mode on all systems that you want to configure.

Add Server-level SNMP Settings

CA Virtual Assurance collects performance metrics from SystemEDGE through SNMP requests. You can configure SNMP settings for individual servers.

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Expand the Data Center folder, then any subfolder, and select the server that you want to configure.

3. Right-click and select Policy.

   The Policy submenu appears.

4. Click Configure SNMP Settings.

   The Configure SNMP Settings dialog opens showing the server-level settings.

5. Do one of the following options:

   - Select the check box for an existing metric from the list, and click the tool icon (Edit) to modify an existing entry.
   - Click Add to create an SNMP entry at the server level.

   The New SNMP Settings dialog appears.

6. Complete the following fields, and click OK:

   Name

   Describes the SNMP credentials that are being defined.

   Port

   Defines the port that is configured for SystemEDGE on the system which you want to manage with these credentials.

   SNMP version

   Specifies the SNMP version being used. If you select SNMP v3 Trap, the panel for additional configuration parameters appears.

   Community String (for SNMP v1/v2)

   Specifies the SNMP community string.

   Security User (for SNMP v3)

   Specifies the SNMP security user for the SNMP credentials that are being defined.

   Access Type

   Specifies the access type. Valid options are Read-Only or Read-Write.

Timeout

Specifies how long in seconds to wait for a confirmation of notification delivery before timing out.

Default: 10 seconds

Retry limit

Specifies the number of times to retry sending a notification after a timeout.

Authentication (for SNMP v3)

Specifies the authentication protocol to use. Select MD5 or SHA from the Type drop-down list and specify a password.

Privacy (for SNMP v3)

Specifies the privacy protocol to use. Select DES, AES, or 3DES from the Type drop-down list and specify a password.

The SNMP settings are saved and appear in the Server Settings table.

Apply the Package Wrapper SNMP Settings as Server-level Settings

You can instruct CA Virtual Assurance to use the package wrapper SNMP settings as server-level SNMP settings after SystemEDGE registers with Policy Configuration. Otherwise, the package wrapper SNMP settings are only used until SystemEDGE registers with Policy Configuration.

Follow these steps:

1. Change to Administration, Configuration, Deployment & Configuration.

   The following option controls whether the SNMP settings in the package wrapper become server-level SNMP settings or not.

   - Create server-specific SNMP settings when a SystemEDGE Agent registers.

2. Enable or disable this option according to your requirements.

   If you disable this option, the package wrapper SNMP settings are not stored on the manager and not available for distribution.

   If you enable this option, the package wrapper SNMP settings are stored as server-level SNMP settings on the CA Virtual Assurance manager.

3. Change to Resources, Configure, and open the SystemEDGE policy that you want to apply to managed nodes.

   The policy pane appears.

4. Select the appropriate items under Traps & Communities, Server Communities.

   Server Read and Server Write represent the server-level SNMP settings which are available for the managed nodes to which you want to apply this policy.

5. Select the appropriate items under Default Communities.

   The Default Communities represent the global SNMP settings.

6. Change to Resources, Deployment, Packages and open the SystemEDGE package wrapper.

7. Set the SNMP credentials, select the policy that you want to apply to the SystemEDGE agents after installation on the managed node, and save the wrapper.

8. Create a Deployment Job and deploy the SystemEDGE package with its policy to the managed nodes.

How to Configure SNMPv3

The following diagram provides an overview of the required actions when you specify SNMPv3 settings for your environment. This use case describes an SNMPv3 configuration for CA Virtual Assurance.

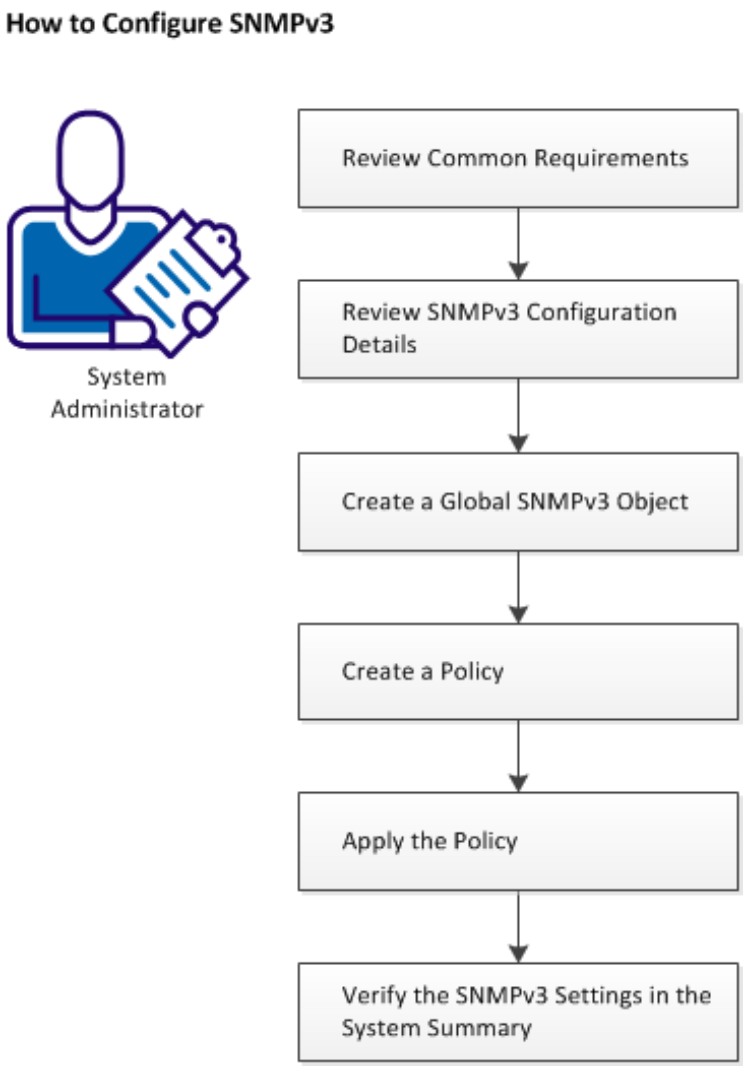

Follow these steps:

Review Common Requirements (SNMPv3) (see page 112)
Review SNMPv3 Configuration Details (see page 112)
Create a Global SNMPv3 Object (see page 113)
Create a Policy (see page 114)
Apply the Policy (see page 116)
Verify the SNMPv3 Settings in the System Summary (see page 116)

Review Common Requirements (SNMPv3)

Review the following requirements before you start configuring the SNMP settings on CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, and Windows Server operating systems.
- You have a basic understanding of CA SystemEDGE.
- You have read the How to Configure SNMPv1/v2 Settings and Access Control Lists scenario.
- You have read the How to Manage Server-level SNMP Settings scenario.
- You can access a CA Virtual Assurance manager installation that includes the Monitoring Agent (CA SystemEDGE).
- You can access the monitoring agents (CA SystemEDGE) on managed nodes.
- You can access the CA Virtual Assurance user interface.
- CA Virtual Assurance has discovered all relevant systems.
- SystemEDGE runs in managed mode on all systems that you want to configure.

Review SNMPv3 Configuration Details

Consider the following details when you intend to use SNMPv3 for the communication between the CA Virtual Assurance manager and managed nodes in your environment.

- SystemEDGE requires at least one SNMPv1 community for its installation. After CA Virtual Assurance has discovered the server, CA Virtual Assurance treats these SNMPv1 settings as server-specific SNMP settings.
- Verify, that your Infrastructure Managers support SNMPv3.
- Create global SNMPv3 credentials.
- Create a policy for applying SNMPv3 settings to remote servers.
- If you want a pure SNMPv3 configuration, prohibit Policy Configuration from applying server-specific SNMPv1 settings.

Create a Global SNMPv3 Object

You can create global SNMP settings or server-specific SNMP settings which are valid for one particular server. Global settings can be applied to server groups through policies.

Follow these steps:

1. Navigate to Administration, SNMP in the user interface.

   SNMP settings page for global objects appears.

2. Click Actions, New to create an SNMP Settings object.

   The New SNMP Settings dialog appears.

3. Set SNMP Version to SNMPv3.

   The SNMPv3-related fields appear in the dialog.

4. Complete the following fields, and click OK:

   Name

   Specifies a name for the SNMP credentials that are being defined.

   Port

   Defines the port that is configured for SystemEDGE on the systems which you want to manage with these credentials.

   SNMP version

   Specifies SNMPv3 (already set in the previous step).

   Security User

   Specifies the SNMP security user for the SNMP credentials that are being defined.

   Access Type

   Specifies the access type. Valid options are Read-Only or Read-Write.

   Timeout

   Specifies how long in seconds to wait for a confirmation of notification delivery before timing out.

   Default: 10 seconds

   Retry limit

   Specifies the number of times to retry sending a notification after a timeout.

   Authentication

   Specifies the authentication protocol to use. Select MD5 or SHA from the Type drop-down list and specify a password.

Privacy

Specifies the privacy protocol to use. Select DES, AES, or 3DES from the Type drop-down list and specify a password.

The SNMP settings are saved and appear in the Server Settings table.

Create a Policy

After you completed your global SNMPv3 settings, apply the SNMPv3 settings to policies.

Follow these steps:

1. Navigate to Resources, Configuration in the user interface.

   The Policy page appears.

2. Expand Policy, Policies, SystemEDGE in the navigation pane.

   The SystemEDGE page appears listing the available policies.

3. Click  to create a policy.

   The New SystemEDGE Policy dialog appears.

4. Specify the required data to create a policy and click OK.

5. Open the policy that you want to apply to one or more managed systems and click Traps & Communities.

   The Communities page appears with the table of SNMP settings and the following options:

   - Include only Server Communities
   - Include Server Communities and all Default Communities
   - Custom Selection

   Note: The only default (global) SNMP settings from the table that are included in the configuration are those settings with a port that matches the agent port.

6. Select Custom Selection.

   This option allows you to select global SNMPv3 objects only and clear any server-specific SNMP settings.

7. Select at least one SNMPv3 settings object with the appropriate port specified for each of your target systems and click Save Policy.

   The selected SNMPv3 objects are associated with the policy.

8. Click the Trap Destinations tab.

   The Trap Destinations page appears. You can configure SNMPv3 Trap Destinations.

9. From the Trap Type field, select SNMPv3 Trap Info or SNMPv3 Notification Info (also referred to as INFORM requests and confirmed traps).

   Depending on the selection, the following fields appear:

   Destination

   Specifies the host to which to send the trap. You can specify a host name or an IP address.

   Port

   Specifies the port number on the destination host that you want to send the trap.

   Username

   Specifies the SNMPv3 user with which to send the trap.

   Encoding

   Specifies the type of encoding to use when sending traps.

   Default: 000

   This encoding is similar to configuring the trap encoding in SNMPv1. See also *SystemEDGE User Guide,* Configure SNMPv1 Trap Destinations.

   Context

   * (asterisk) is the only supported value for this field. This value is mandatory.

   Timeout

   (Notifications only) Specifies how long in seconds to wait for a confirmation of notification delivery before timing out.

   Retries

   (Notifications only) Specifies the number of times to retry sending a notification after a timeout.

10. Fill out the fields and click Add.

    A new entry appears in the Trap Destinations table.

    You can repeat the last step to add more entries to the table.

11. Select one of the Trap Destinations and click Save Policy.

    The policy is saved with the appropriate trap destination.

You can distribute the policy to the appropriate server group.

Apply the Policy

After you completed your SNMPv3 settings, you can apply the policy to systems in the network.

Follow these steps:

1. Navigate to Resources, Configuration in the user interface.

   The Policy page appears.

2. Expand Policy, Policies, SystemEDGE in the navigation pane.

   The SystemEDGE page appears listing the available policies.

3. Select the policy that you have previously saved with appropriate SNMPv3 settings.

   The policy page opens.

4. Click Action, Apply.

   The Select Machines page appears.

5. Select all systems which you want to configure with that policy, and click Apply.

   You can view the delivery status or return to the Policy page.

   The new settings are applied to the target systems.

Alternative

If you want to deploy SystemEDGE to a remote system and use SNMPv3 credentials, you can apply the policy to the package wrapper and run the deployment job.

Verify the SNMPv3 Settings in the System Summary

To verify that the SNMPv3 settings have been applied to the target systems correctly, change to the Explore pane in the CA Virtual Assurance user interface.

Follow these steps:

1. Expand the components tree in the Explore pane.
2. Select a managed system to which you have applied the SNMPv3 settings and open Summary.

   The Machine Status Information appears.

3. Verify that the Active SNMP Credentials field shows the SNMPv3 global object.

Note: If you have applied a pure SNMPv3 configuration to a managed server and you open the SystemEDGE Control Panel (Windows only) on that server, then the community and trap fields are empty. The SystemEDGE Control Panel shows SNMPv1 information in these fields.

Configure CA Virtual Assurance to Forward Events

Configure the product to forward events to a CA or third-party SNMP Event Manager. The process consists of two parts:

1. Configure the event manager to receive CA Virtual Assurance traps or events.
2. Configure CA Virtual Assurance to forward the events.

The following procedure assumes that you have configured your Event Manager console to receive events.

Follow these steps:

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. Enter the management server name in the Server field.
7. Enter a different port number for SNMP or leave the default port 162, which is automatically populated.
8. Click OK.

   A confirmation message appears.
9. 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.

How to Deploy SystemEDGE and AIMs

This section explains how to set up and manage jobs to deploy your monitoring software successfully.

More information:

Overview

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 that define how software packages are deployed on target systems. These configurations are referred to as Package Wrappers.

Deployment Job Management

Allows creation, start, cancellation, and filter of deployment jobs, allowing the concurrent deployment of packages to multiple targets using multiple distribution servers.

Deployment Job Reporting

Allows querying 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:

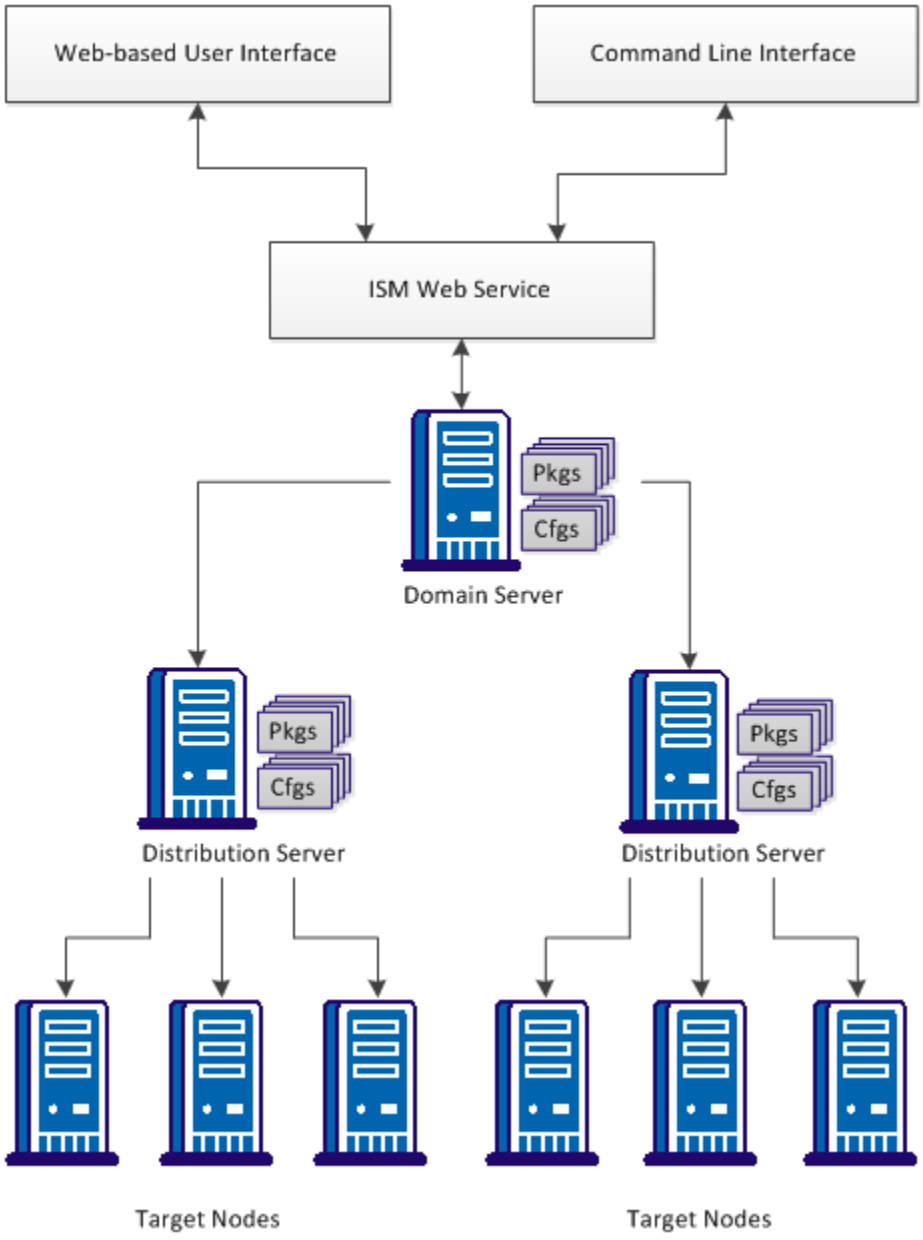

About Packages

Deployment packages provide the materials necessary to deploy monitoring software to systems across your enterprise. Deployment packages are broken into package wrappers, which are platform-specific. Package wrappers encapsulate the installation parameters required to install the agent software and are available for all platforms that support deployment.

Note: The default package wrapper name is not localized and reads 'default' in all supported languages. Custom package wrapper names are localized.

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.

Configuration

This section provides you details about Remote Deployment user interface configurations and Distribution Server connections.

More information:

Deployment Dashboard Views (see page 121)
Enhanced Search Functionality for Remote Deployment (see page 122)
Job Status Filter (see page 122)
Change the Domain Server a Distribution Server Connects To (see page 123)

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 deployments for each deployment package type.

Completed Deployment Jobs

Displays a list of deployments that is completed successfully. You can click the job ID to view details about the job.

Enhanced Search Functionality for Remote Deployment

The enhanced search functionality provides search results for keywords related to Remote Deployment and also provides quick access to the Remote Deployment operations from the search results. The benefits are as follows:

The benefits are as follows:

- Access the Remote Deployment components swiftly.
- Deploy the Remote Deployment software packages to the server and service.
- Manage Remote Deployment software packages and templates.
- Create and manage the deployment jobs quickly and give access to the available packages and wrappers.

Follow these steps:

1. Enter the keyword (or a partial value with a wildcard) in the Value field, and click Search.

   Example:

   Deploy or Remote or remote deployment.

   A list of Remote Deployment links appears.
2. Select the appropriate Remote Deployment operation.

   Perform the Remote Deployment operation.

Job Status Filter

The job status data is filtered to show only the relevant details of each job. You can sort and customize the columns, and filter by one or more columns.

Follow these steps:

1. Select Resources, Deploy.

   The Deployment pane displays the Packages, Templates, and Jobs folders.
2. Click Jobs.

   The job details appear in the right pane.
3. (Optional) Select/unselect the check box for Job status columns.

   The customized columns appear.
4. Select/unselect the Filters of the columns.

   Jobs appear per the filter selection.

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.

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.

Note: The initial deployment to a target installs IDPrimer, a small installation agent. Once IDPrimer is installed, subsequent deployments to the same target require less time.

Deployment recommendations:

- Verify that target servers typically meet the requirements for deploying software remotely.
- 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.

Note: For more information about CA Virtual Assurance scalability, see Scalability Best Practices.

Multiple Distribution Servers

Although the remote deployment solution lets you use a single central server (manager) for all your deployments, CA Technologies recommends 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 but need them managed centrally using a single manager.

  In such a case, CA Technologies recommends that each location have at least one distribution server connected 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 such a case, you can install many distribution servers split logically across subnets, and these distribution servers connect to the central domain server.

Deployment Packages

Deployment packages provide the materials necessary to deploy monitoring software to systems across your enterprise. Deployment packages are broken into platform-specific variants, package wrappers are available for all platforms that support deployment.

Important! The AIMs depend on SystemEDGE and Advanced Encryption packages. To deploy any of these packages, SystemEDGE and Advanced Encryption must either exist on the system already, or be included in the deployment job.

The following deployment packages are available:

Performance Agent (CA Systems Performance LiteAgent)

Provides a lightweight monitoring agent for monitoring and collecting performance metrics on Windows, UNIX, or Linux.

SystemEDGE

Provides the core SystemEDGE agent.

SystemEDGE ADES

Provides the AIM for Active Directory and Exchange Server.

SystemEDGE Advanced Encryption

Provides a FIPS 140 compliant encryption package for SystemEDGE.

SystemEDGE AIX LPAR

Provides the AIM for IBM PowerVM (LPAR).

SystemEDGE CXEN

Provides the AIM for Citrix XenServer.

SystemEDGE Citrix XenDesktop

Provides the AIM for Citrix XenDesktop.

SystemEDGE GalaX

Provides the AIM for Huawei GalaX8800.

SystemEDGE Hyper-V

Provides the AIM for Microsoft Hyper-V.

SystemEDGE IBM PowerHA

Provides the AIM for IBM PowerHA, formerly known as High Availability Cluster Multi-Processing.

SystemEDGE KVM

Provides the AIM for Red Hat Enterprise Virtualization (RHEV) based on KVM technology.

SystemEDGE MSCS

Provides the AIM for Microsoft Cluster Support (MSCS).

SystemEDGE RM

Provides the Remote Monitoring AIM.

SystemEDGE Solaris Zone

Provides the AIM for Oracle Solaris Zones.

SystemEDGE SRM

Provides the Service Response Monitor AIM.

SystemEDGE UCS

Provides the AIM for Cisco UCS.

SystemEDGE VC

Provides the AIM for VMware vCenter.

SystemEDGE VCLOUD

Provides the AIM for VMware vCloud Director.

Default Package Wrappers

Default package wrappers are provided for the software packages that are deployed using Remote Deployment. These package wrappers contain installer parameters with a set of default values for the chosen software package. If a package requires mandatory parameters, specify these parameters and save the settings before you deploy the package.

You do not have to edit the parameters again, unless there is a need to modify the installer parameter values for a package. If you proceed to deploy a package without specifying mandatory parameters, the deployment process stops. The package wrapper is not in a deployable state.

The available package wrappers provide the following parameters. Mandatory parameters are indicated in the user interface:

SystemEDGE

Global SNMP Settings that are specified under Administration, Configuration, SNMP populate the drop-down lists for the following fields in the SystemEDGE package wrapper:

- Port
- Read Community
- Read-Write Community

Alternatively, you can edit the fields inline.

The available community strings depend on the port setting. When you select the port number first, then you get automatically the valid community strings in the drop-down lists for that port.

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.

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 the default group of this user as an effective group.

Default: The agent operates using root account.

Start Agent check box

Specifies whether to start SystemEDGE at the end of the installation automatically.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

Note: The Suppress Reboot check box is available for Windows package only.

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 (1)
CASE_SNMP_TRAP_DESTINATION (1)
CASE_SNMP_TRAP_PORT (1)
CASE_SNMP_READ_COMMUNITY (1)
CASE_SNMP_WRITE_COMMUNITY (1)
CASE_SNMP_READ_ALLOWED_MANAGERS (1)
CASE_SNMP_WRITE_ALLOWED_MANAGERS (1)

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 parameters, see the *Installation and Deployment* chapter in the *SystemEDGE User Guide*.

CA SystemEDGE ADES

Windows Domain

Specifies the Windows Domain to monitor.

Domain User

Specifies the domain administrator user to connect to the Domain Server or Exchange Server.

Domain User Password

Specifies the password of the domain administrator user to connect to the Domain Server or Exchange Server.

Management 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.

Management Mode

Specifies the option for providing management.

0

Specifies the entire domain for monitoring.

1

Specifies a specific host of the domain for monitoring.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

SystemEDGE Advanced Encryption

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE AIX LPAR

LPAR Host

Specifies the hostname to use for connecting to the IBM LPAR server. Specify the name of the IBM LPAR host to deploy this package.

Username

Specifies the username to use for connecting to the IBM LPAR server. Specify the name of the IBM LPAR user to deploy this package.

Password

Specifies the password to use for connecting to the IBM LPAR server. Specify an IBM LPAR password to deploy this package.

CA SystemEDGE CXEN

CXEN Hostname

Specifies the hostname to use for Citrix XenServer integration.

CXEN Username

Specifies the username to use for Citrix XenServer integration.

CXEN Password

Specifies the password to use for Citrix XenServer integration.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE CXenDesktop

Hostname

Specifies the hostname for Citrix XenDesktop integration.

Username

Specifies the username for Citrix XenDesktop integration.

Password

Specifies the password for Citrix XenDesktop integration.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE GalaX

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE PowerHA

Hostname

Specifies the hostname to use for connecting to IBM PowerHA. Specify a PowerHA host name to deploy this package.

Username

Specifies the username to use for connecting to IBM PowerHA. Specify a PowerHA user name to deploy this package.

Password

Specifies the password to use for connecting to IBM PowerHA. Specify a PowerHA password to deploy this package.

Port

Defines the PowerHA port number.

Default: 22

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE Hyper-V

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE KVM (RHEV)

KVM Hostname

Specifies the hostname to connect to Red Hat Enterprise Virtualization (RHEV).

KVM Username

Specifies the username to connect to RHEV.

KVM Password

Specifies the password to connect to RHEV.

KVM Port

Specifies the port to connect to RHEV.

Default: 8443

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE MSCS

MSCS Hostname

Specifies the hostname to connect to the cluster.

MSCS Username

Specifies the username to connect to the cluster.

MSCS Password

Specifies the password to connect to the cluster.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE RM

Default WMI username

Defines the default username to use for connecting to remote machines. Specify a username to deploy this package.

Default WMI password

Defines the default password to use for connecting to remote machines. Specify a password to deploy this package.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA 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.

Allow untrusted SSL check box

Specifies whether to allow accessing an SSL site with unverified certificates.

Disable user TOS check box (Windows)

Specifies whether to disable applications from setting type of service bits in outgoing IP packets.

Suppress Reboot check box (Windows)

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE Solaris Zone

Zones Host

Specifies the hostname to use for connecting to the Solaris Zone server. Specify a Solaris Zone hostname to deploy this package.

Username

Specifies the username to use for connecting to the Solaris Zone server. Specify a Solaris Zone username to deploy this package.

Password

Specifies the password to use for connecting to the Solaris Zone server. Specify a Solaris Zone password to deploy this package.

CA SystemEDGE UCS

UCS hostname

Specifies the hostname to use for connecting to UCS. Specify a UCS host name to deploy this package.

UCS username

Specifies the username to use for connecting to UCS. Specify a UCS user name to deploy this package.

UCS password

Specifies the password to use for connecting to UCS. 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.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE VC

Hostname

Specifies the hostname to use for connecting to vCenter. Specify a vCenter host name to deploy this package.

Username

Specifies the username to use for connecting to vCenter. Specify a vCenter user name to deploy this package.

Password

Specifies the password to use for connecting to vCenter. Specify a vCenter password to deploy this package.

Port

Defines the vCenter port number.

Default: 443

Protocol

Specifies what protocol to use, HTTP or HTTPS.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA SystemEDGE VCLOUD

VCLOUD hostname

Specifies the hostname to use for connecting to vCloud.

VCLOUD username

Specifies the username to use for connecting to vCloud.

VCLOUD password

Specifies the password to use for connecting to vCloud.

VCLOUD port

Defines the vCloud port number.

Default: 443

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

CA Systems Performance LiteAgent

Shared Path

Defines the root installation directory to use for CA Shared Components.

Install Path

Defines the root installation directory for the package.

Suppress Reboot check box

Specifies whether to suppress an automatic reboot at the end of the installation.

More Information

Create a New Package Wrapper (see page 142)
Modify a Package Wrapper (see page 142)

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 139) section.

The package library is located in the following directory:

%AllUsersProfile%\CA\SM\domainserver\Deployment\Packages\SM

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
    |   └──5.8.0
    |       └──ENU
    |           ├──AIX_aix
    |           ├──HPUX_hp
    |           ├──HPUX_ia64
    |           ├──Linux_ia64
    |           ├──Linux_ppc
    |           ├──Linux_x86
    |           ├──Solaris_sparc
    |           ├──Solaris_x86
    |           ├──Windows_ia64
    |           ├──Windows_x64
    |           └──Windows_x86
    ├──CA_SystemEDGE_CXEN
    ├──CA_SystemEDGE_CXENDESKTOP
    ├──CA_SystemEDGE_ESAD
    ├──CA_SystemEDGE_GALAX
    ├──CA_SystemEDGE_HACMP
    ├──CA_SystemEDGE_HyperV
    ├──CA_SystemEDGE_KVM
    ├──CA_SystemEDGE_LPAR
    ├──CA_SystemEDGE_MSCS
    ├──CA_SystemEDGE_RM
    ├──CA_SystemEDGE_SRM
    ├──CA_SystemEDGE_UCS
    ├──CA_SystemEDGE_VCLOUD
    ├──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.

Language

The installation package language.

Example: ENU

Architecture

The architecture-specific installation materials, for example, Windows_ia64, Solaris_x86.

Note: 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.

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 an 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 software package installation materials, each deployable package is referenced by an additional package configuration file, pkginfo_PLATFORM.xml. The package configuration file describes installation packages and configuration process.

The configuration files provide the following:

- A localizable description of the installation package.
- A mechanism by which package dependencies are 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 are provided optionally using either side-by-side locale-specific files or embedded within a single file. The file name that matches with pkginfo_PLATFORM.xml is loaded to obtain localized message data.

The deployment system requires the package configuration files to be located parallel to the platform-specific subdirectory in the packaging tree. For example, see the following directory:

%AllUsersProfile%\CA\SM\domainserver\Deployment\Packages\SM\CA_SystemEDGE_Core\5.7.1\ENU

```
pkginfo_AIX.xml
pkginfo_HPUX.xml
pkginfo_Linux.xml
pkginfo_LinuxPPC.xml
pkginfo_Solaris_sparc.xml
pkginfo_Solaris_x86.xml
pkginfo_Windows.xml
```

Using Remote Deployment

You can deploy monitoring agents to multiple systems in one operation using centralized remote deployment from the CA Virtual Assurance user interface. Package deployment through CA Virtual Assurance is a secure, reliable solution that lets you configure the monitoring software that is installed across your enterprise from a central interface.

Deployment Restrictions

Consider the following restrictions before performing a deployment:

- If you want to install an agent on the CA Virtual Assurance Manager system, 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 to gain remote access to target systems. When these services are not available on the target nodes, it is necessary to install the IDPrimer client package and a corresponding key on target systems.

  Note: For more information about installation, see the section *Manual Installation of the Remote Deployment Primer Software*.
- Deployment is supported to most, but not all, supported agent platforms.

Deployment Credential Restrictions

The UI limits the entries for both username and password fields to 64 characters.

Audit Trail

Jobs and tasks are the two fundamental concepts of the deployment system. A deployment job specifies one or more packages to be available on one or more target systems. A deployment task represents each individual deployment of a software package on a target system. Deployment job reporting allows you to query the status of deployment jobs.

You can create, control, and can inquire the state of deployment jobs. After 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 identify and address any problems.

Remote deployment is able to provide the following information:

- Which deployment jobs are currently:
  - Inactive (not yet started)
  - Active
  - Completed, 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 the /tmp file system mounted with the noexec flag.

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.

   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, Deploy.

   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.

Copy a Package Wrapper

You can copy a package wrapper to edit the properties according to your needs.

Follow these steps:

1. Select Resources, Deploy.

   The Deployment pane displays the Packages, Templates, and Jobs.

2. Expand Packages, the specific package type, and platform.

3. Select the wrapper.

   The wrapper details appear in the right pane.

4. Right-click a wrapper name. Select Copy. You can also select Copy from the Actions drop-down menu.

   The Copy dialog appears.

5. Enter a new name for the package wrapper, an optional description, and click Ok.

   The new package wrapper appears in the deployment pane.

6. Edit the properties according to your needs and click Save.

   The new package wrapper appears in the left pane.

Delete a Package Wrapper

You can delete a package wrapper that you no longer need.

Follow these steps:

1. Select Resources, Deploy.

   The Deployment pane displays the Packages, Templates, and Jobs.

2. Expand Packages, the specific package type, and platform.

3. Select the wrapper.

   The wrapper details appear in the right pane.

4. Right-click a wrapper name. Select Delete. You can also select Delete from the Actions drop-down menu.

   A warning message appears.

5. Click Yes to confirm the deletion.

   The package wrapper is deleted.

Rename a Package Wrapper

You can rename a package wrapper.

Follow these steps:

1. Select Resources, Deploy.

   The Deployment pane displays the Packages, Templates, and Jobs.
2. Expand Packages, the specific package type, and platform.
3. Select the wrapper.

   The wrapper details appear in the right pane.
4. Right-click a wrapper name. Select Rename.

   The Rename dialog appears.
5. Enter a new name and click OK.

   The package wrapper is renamed.

Create a Deployment Job

To deploy agents to systems, 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.

Follow these steps:

1. Select Resources, Deploy.

   The Deployment pane displays the Packages, Templates, and Jobs.
2. Right-click the Jobs folder in the Manage Resource pane and select Create New Job.

   The Jobs Setup page appears.
3. Enter a name in the Job Name pane and optionally base the job on an existing template, and click Next.

   The Package Selection page appears.
4. Select a platform and the packages you want to deploy.
5. (Optional) Click the Details tab.

   The Package Wrapper Details dialog appears and lets you edit the package properties in-line. If the package wrappers are in an incomplete or invalid state, and the fields can be modified through in-line editing.

   a. Click Edit and modify the package wrapper properties.

   b. Click Save, and then click OK.

   c. The package wrapper properties are updated.

6. Click the down arrow to add the package wrappers to the job, and click Next.

   The Machine Selection page appears.

7. Select the systems to deploy to and click Next. If you have many servers in your environment, multiple pages with some entries can be required to list all servers. When you select servers on a page and scroll to the next page, any selections that are made on previous pages remain valid.

   The Machines Selected page appears.

8. Click Set Credentials, set the system credentials that are required to establish a connection and click Next.

   Note: Deployment to Windows target systems using domain credentials must be in the form of DOMAIN\username.

   The Advanced page appears.

9. (Optional) Set the distribution server to manage the deployment. If not set, it is automatically chosen.

10. Select the scheduling options for the job:

    Immediate Delivery

    Starts the job immediately after creating new deployment job. The immediate delivery is the default option.

    Staggered Delivery

    Delivers the packages over a specific time period.

    Scheduled Delivery

    Schedules the deployment for a specific time in the future.

11. (Optional) If a package has previously been successfully deployed to a system using this deployment infrastructure, you can force it to run again.

12. Click Next.

    The Summary page appears.

13. Review the details of the job and click Deploy.

    The deployment job is created.

Note: You can 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.

Specify Read-Write Community Prior To Deployment

To get full SystemEDGE monitoring and management functionality you must specify a valid read-write SNMP community for the SystemEDGE agent. You can configure the read/write community string in the remote deployment package wrapper for SystemEDGE prior to deployment.

Specify read-write community prior to deployment

1. Select Resources, Deployment.
2. Open the Deployment pane.

   Available deployment groups appear.
3. Expand Packages, the specific package type, and the wrapper platform, and select the wrapper to modify.

   The wrapper details appear in the right pane.
4. Specify the read-write parameters in the Read-Write Community field and click Save.

   Note: The agent does not function correctly, if you specify community strings with space characters or semicolon (;) in the user interface.

More Information:

Add Server-level SNMP Settings (see page 108)
Apply Policy to Machines (see page 263)

Specify Read-Write Community Post-Install

To get full SystemEDGE monitoring and management functionality you must specify a valid read-write SNMP community for the SystemEDGE agent. If you have already deployed the SystemEDGE agent, you can add the read/write community string post-installation. You can do this 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.

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 'Include only Server-specific SNMP settings'.

5. Click Save Policy.

   The policy is saved.

To create a global SNMP entry

1. Click Administration.

   The Administration page appears.

2. Click Configuration.

   The Configuration page appears.

3. Click SNMP.

4. Select the check box for the SNMP settings you want to edit from the list and click the tool icon (Edit).

   The Edit SNMP Settings dialog appears.

5. Select Read-Write from the Access Type drop-down, specify the parameters in the Community String field, and click OK.

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

   The Available Policies page appears.

7. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

8. Click the Traps and Communities tab.

   The Communities page appears.

9. Select ‘Include Server-specific SNMP settings and selected Default Settings’.

10. Click Save Policy.

    The policy is saved.

Note: For more information, see the section "Apply Policy to Machines".

More Information:

Add Server-level SNMP Settings (see page 108)

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 tab displays a table of all created deployment jobs that lists the job name, included packages, 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.

Note: In the Job Status pane, you can filter out particular job tasks by selecting the filters.

Follow these steps:

1. Select Resources, Deploy.

   The Deployment pane displays the Packages, Templates, and Jobs.

2. Click the Jobs folder.

   The Job Status page appears.

3. Click the job that 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.

5. Click Status Extended to view extended information about a task.

   The Extended Status information dialog appears with details about the task:

   Information

   Displays general information about the task.

   Message

   Displays a message about the task, for example, Package delivery failed.

   Reason

   Displays the reason for the failure.

   Examples:

   - Lack of machine availability
   - Invalid system credentials
   - Inability to resolve the system host name
   - Unfulfilled package dependency

   Action

   Displays what action to take to correct the problem.

Resubmit a Deployment Job

You can resubmit a failed or partially failed deployment job.

Follow these steps:

1. Select Resources, Deploy.

   The Deployment pane displays the Packages, Templates, and Jobs.

2. Click the Jobs folder.

   The Job Status page appears.

3. Click the job that you want to resubmit.

   The Job Information page appears.

4. Click Actions and select Resubmit.

   The Deployment wizard appears at the Package Selection screen.

5. Select the package wrappers to be deployed as desired and click Next.

   The Machines Selected page appears.

   (Optional) Delete the machines that you do not want to deploy to.

   Note: The machines where all packages were previously successfully deployed are not selected.

   Click Set Credentials, modify the credentials for all selected machines as required, optionally remove machines from the restarted job, then click Next.

   The Advanced page appears.

6. (Optional) Modify the Scheduling options for the job.

   Immediate Delivery

   Starts the job immediately after creating new deployment job. The immediate delivery is the default option.

   Staggered Delivery

   Delivers the packages over a specific time period.

   Scheduled Delivery

   Schedules the deployment for a specific time in the future.

7. Select Redeploy previously deployed packages to force all packages, including packages previously successfully deployed to be deployed, and click Next.

   The Summary page appears.

8. Review the details of the job and click Deploy.

   The job is resubmitted.

View Deployed Packages

The Monitoring Software page lets you view a list of packages deployed to a single machine or group of machines.

To view deployed packages

1. Select Resources.

   The Explore pane appears.

2. Select a system or a service.

   The Summary page appears.

3. Select Monitoring Software, Deployment.

   The Deployment History page appears, with a list of all deployment jobs for the machine. The table displays details of all deployment jobs for the selected system:

   - Task ID
   - Job ID
   - Target
   - Package
   - Platform
   - Wrapper
   - Wrapper Version
   - Stated By
   - Start Time
   - End Time
   - Status
   - Status Extended

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, included packages, 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 cases:

- Invalid system credentials
- Inability to resolve the system host name
- Unfulfilled package dependency

Note: You can resubmit a job from this pane to correct the reason for failure and redeploy. Click the Jobs folder to access this view.

Specific Remote Deployment Use Cases

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.

Reconfigure the SystemEDGE Agent Port

You can reconfigure the SystemEDGE agent port from the standard (default) port 161 to 1691. For example, you can install the Microsoft SNMP Service which implements a MIB-II agent on default port 161. Editing the sysedge.cf file is not a supported way of reconfiguring the agent port. Changing the port should be done by reinstalling the agent. This can be done by redeploying the agent using Remote Deployment specifying a different port. On the Windows systems, you can also reconfigure the agent using the SystemEDGE control panel applet.

Reconfigure the SystemEDGE agent using the control panel

1. Click Start, Control Panel, select Add or Remove programs, select SystemEDGE Core in the list, click Change.

   The SystemEDGE Setup Wizard opens.

2. Click Next.

   The Reinstallation Type page opens.

3. Select Reinstall and click Next.

   The Application Configuration page opens and allows you to change the Install documentation setting.

4. Click Next.

   The SystemEDGE SNMP Port Number page opens.

5. Specify the SystemEDGE port number 1691, and click Next.

6. Review the settings, and click Reinstall.

   The SystemEDGE agent is reconfigured to use port number 1691.

Reconfigure the SystemEDGE agent port using remote deployment/policy configuration

1. Follow the steps in the chapter Create a Deployment Job (see page 144). In Step 5 of the wizard, select "Redeploy previously deployed packages".

   Note: On reinstall, all supplied installation parameters except the port number are ignored.

   After you reinstall the agent to change the port, some manual steps are required on the manager to ensure that the agent is configured with the correct community strings.

2. Do *one* of the following:

   Create server-specific SNMP entries for the server:

   1. In the CA Virtual Assurance UI, click the Resources tab, open the Explore pane, select the Machine_Name.

      The Machine_Name is selected.

   2. Right-click the Machine_Name and select Policy, Configure SNMP Settings.

      The SNMP Settings page appears.

   3. Click Add to create a new entry for the required port.

      The New SNMP Settings page appears.

   4. Enter the required details, and click Ok.

      The server-specific SNMP entries are created for the server.

   Ensure that global SNMP settings exist and update the policy:

   In the CA Virtual Assurance UI, navigate to Administration, SNMP and add a new entry for the required port. When the settings are correct, edit the policy by navigating to the Resources tab, open the Configure pane, and select the policy. Then click Traps & Communities, Communities and select the middle option, "Include server-specific SNMP settings and all Default settings". Save the policy by clicking Save Policy. You should now apply the policy to the system (see page 263). You can check the community strings in use by the agent using the SystemEDGE control panel applet on the agent machine.

Note: For more information, see the chapter "Deploying/Installing SystemEDGE Agents Using Custom Ports" in the *Administration Guide.*

More Information:

Add Server-level SNMP Settings (see page 108)

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 that is 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 that is 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 *
  ```

  Note: Replace $username by the actual username that is used for Remote Deployment. If the paths for "id" and "sh are different from /usr/bin/id or /bin/sh, adjust the path in the configuration entry appropriately.

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 the file /etc/user_attr:

  ```
  user::::type=normal;profiles=Primary Administrator
  ```

Agent Configuration Without Write Community

Although it is not mandatory to provide a write community for SystemEDGE package wrappers, consider the following information:

- 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.

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

Deployment Jobs

To deploy agents to target systems, create a deployment job first. Deployment jobs contain the details that are required for CA Virtual Assurance to deliver the deployment packages to the appropriate systems at the scheduled time. You can create new jobs using the remote deployment job wizard accessible from several locations. Choose one of the following ways:

- Using the Deploy Job link in the Dashboard quick launch panel
- From the Jobs panel in the Resources, Deployment tab using the + (New) button
- From the context menu of a managed node in the Resources, Explore tab
- From the Monitoring Software, Deployment tab of the currently selected managed node in the Resources, Explore tab using the + (New) button

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 that are required to establish a connection.

IP Address

Specifies the IP address of the interface where you deploy the job. If a system has multiple IP addresses, the IP address with the management property is set as the default.

Note: You cannot select the IP address for which the management property is not enabled.

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.

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.

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 functionalities of the underlying operating systems on source and target computers. The installation is subject to any restrictions resulting from an enterprise network configuration.

The initial step when deploying software is to install a small primer application remotely, the IDPrimer, onto the target computer. *The 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 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 that is used is ADMIN$. IDManager configuration option controls the default share name. This mechanism is available only from deployment managers running on a Windows-based environment. 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 the following conditions are met:

  - the SSH port (22) is enabled to permit connection from the deployment manager
  - 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).

  Note: Most SSH servers support this configuration by default, but if they do not, consult your SSH server documentation for further instructions.

  To 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 154)
  - Verify that SFTP subsystem is enabled

  Remote Deployment supports deploying software to systems with the /tmp file system mounted with the noexec flag.

  When deploying to some IBM AIX systems that are running both an IPv4 and IPv6 stack, using an IPv6 address, configure the target computer SSH server to use port 22 for IPv4. To configure SSH, 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, 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, the deployment fails with a status of No Primer Transport. In such cases, install the software components in other ways, for example, using physical distribution media such as DVD.

Alternatively, you can preinstall or provision machines with the IDPrimer software. This process allows deployment without having to rely on facilities offered by the underlying operating systems. In cases where no authentication has been carried out, supply valid credentials 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, map a share from your deployment manager host computer to each deployment target computer. Use the target user credentials supplied in the deployment request.

  Default share: ADMIN$

- 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 154)

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 159).

Manual Installation of the Infrastructure Deployment Primer Software

If automatic deployment to target computers is not possible for some reason, you can still deploy software by installing the primer software on the target computer manually. This can be done by installing the primer package physically or running the installation using login scripts.

In addition to installing the primer software, you must install a security key that is generated by the deployment manager that 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:

  Valid on 32-bit Windows

  %PROGRAMFILES%\CA\SC\IDMgrApi\packages\private\idprimer\Windows_x86

  Valid on 64-bit Windows

  %PROGRAMFILES(X86)%\CA\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 32-bit libraries dependencies are present. These 32-bit libraries must be present on Linux hosts before installing IDPrimer.

Most 32-bit Linux distributions have them installed by default already. Dependencies on 64-bit Linux can be satisfied by issuing the following command:

- Valid for RedHat, CentOS, SuSE (32-bit and 64-bit OS):

  ```
  yum install libstdc++.i686
  ```

  This command installs in total 4 RPM packages: glibc, libstdc++, nss-softokn-freebl and libgcc.

- Valid for Debian (64-bit):

  ```
  apt-get install ia32-libs
  ```

  This command installs the following required 32-bit libraries: libc, libstd++, libgcc.

Note: For more information about required compatibility libraries and additional system packages, visit the support web site of your Linux supplier.

More Information:

Infrastructure Deployment Process (see page 158)

How to Configure SystemEDGE and Service Response Monitor Through Policies and Templates

This section explains how to manage your monitoring software configurations in your environment from CA Virtual Assurance, a central point of control.

More information:

Configuration Overview (see page 165)
How to Apply Policy and Layered Templates to Servers (see page 168)
How to Create and Apply an Autowatcher to a System (see page 202)
How to Monitor User-specific Metrics (MIB Extensions) (see page 209)
How to Monitor a Specific Windows Performance Registry Metric (see page 211)
How to Create SRM Policy (see page 214)
Discovering the Agents (see page 215)
Common Usage of Policy Configuration Functions (see page 215)

Configuration Overview

You can configure managed agents and apply the configuration to multiple systems in one operation using centralized Policy Configuration 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, 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 215)
Apply Policy to Machines (see page 263)
Agent Policy Dashboard Views (see page 166)
How to Create SRM Policy (see page 214)
Configure and View Applied Policies (see page 265)
Review Policy Application Progress (see page 264)
How to Monitor User-specific Metrics (MIB Extensions) (see page 209)
How to Monitor a Specific Windows Performance Registry Metric (see page 211)

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 policy configuration cannot manage.

Policy Breakdown

Displays a pie chart and list showing all policies and how many systems contain each policy.

Machines with Non-Standard Policies

Displays systems that contain nonstandard changes to an applied policy.

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:

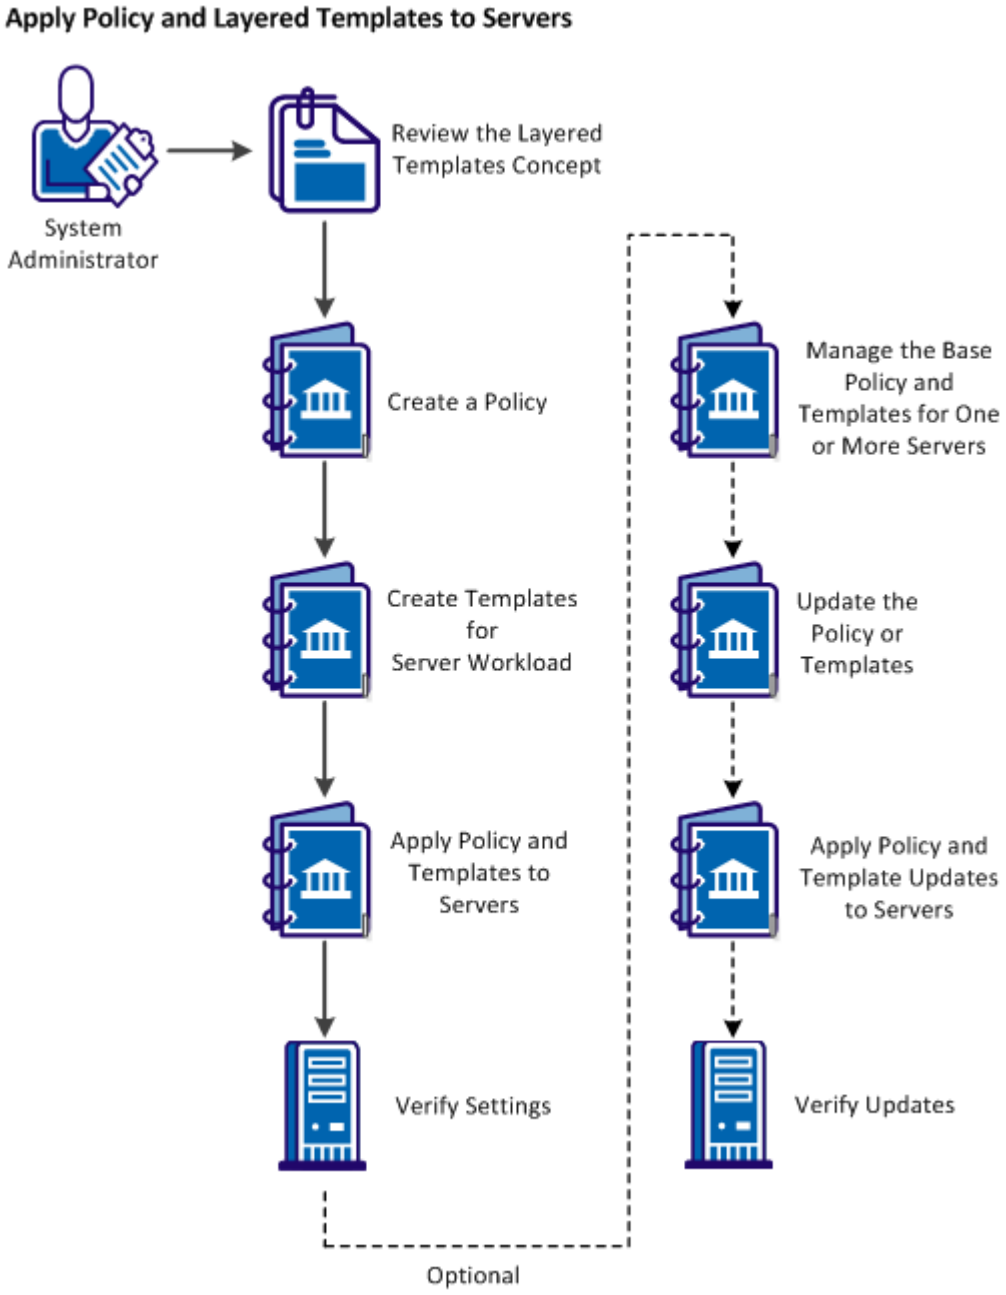

Follow these steps:

Create a Policy (see page 171)
Create Templates for Server Workload (see page 181)
(Optional) Apply Policy and Template Updates to Servers and Verify Updates (see page 201)
Apply Policy and Templates to Servers and Verify Settings (see page 198)
(Optional) Update the Policy or Templates (see page 197)
Layered Templates Concept (see page 169)
(Optional) Manage the Base Policy and Templates for One or More Servers (see page 199)

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

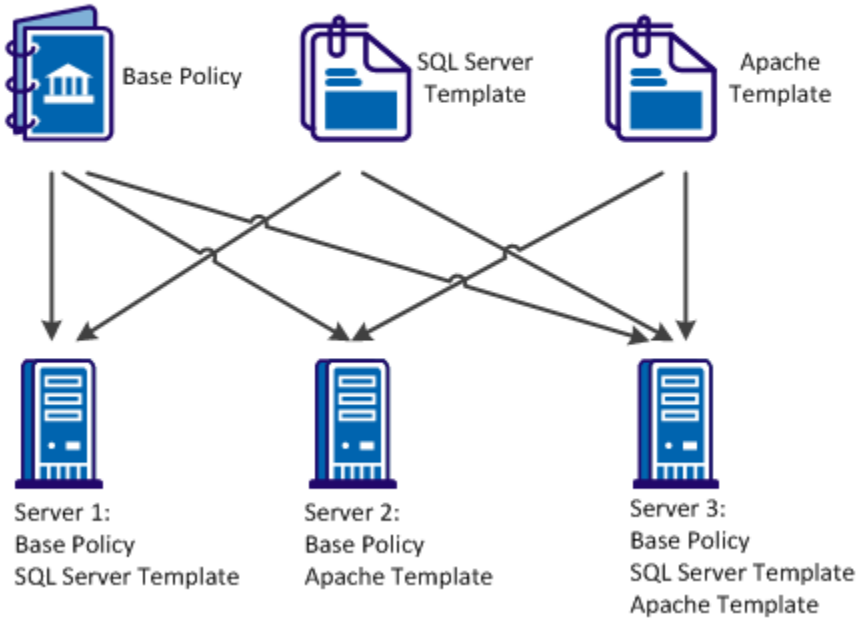

You can use Layered Templates in the following scenarios:

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.

Out of the box templates

You can apply out of the box templates to managed nodes. Configure the policy with the template configuration on the managed nodes. Templates are available for the following operating systems:

For All Operating Systems:

CPU Utilization - Autowatch

Swap Capacity

For Windows:

App Monitoring - CA eTrust Antivirus

Process Crash

System Errors

System Processes

User Activity

Windows Services - Autowatch

For UNIX (AIX, HPUnix, Linux, Solaris):

System Messages

System Processes

User Activity

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.

More Information:

Copy SystemEDGE Policy (see page 216)
Rename SystemEDGE Policy (see page 216)
Delete SystemEDGE Policy (see page 217)

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 177).

   The control settings are defined.

9. Click Save Policy.

   The policy is saved.

More Information:

Configure Object Aggregation (see page 223)

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 218)

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, click Action, and 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 182). To add MIB extensions to the template, see section Define MIB Extensions (see page 194).
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 183)
   - Create a Process Monitor (see page 185)
   - Create a Log File Monitor (see page 187)
   - Create a Windows Event Monitor (see page 189)
   - Create a History Monitor (see page 190)
   - Create a Process Group Monitor (see page 192)

4. (Optional) Repeat the process for any additional monitors.

5. Click Save.

   The monitor is loaded to the policy or the template.

More Information:

Define a Threshold Monitor (see page 233)
Define a Process Monitor (see page 235)
Define a Log File Monitor (see page 237)
Define a Windows Event Monitor (see page 239)
Define a History Monitor (see page 240)
Define a Process Group Monitor (see page 242)

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 Widows 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 194). 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.

(Optional) Reindex Monitors from Templates or a Policy

You can reindex the monitors on the Threshold, Process, Log File, Windows Event, History and Process Group tabs. Reindexing assigns a sequential value to the existing index.

Note: Once you reindex the monitors, this functionality ensures that the future indexes start at the next logical base index.

To reindex the monitors, consider the following:

- Verify that the monitors exist.

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.
4. Click the appropriate monitor tab, click Action, and select Reindex.

   The new base index dialog appears.
5. Enter a numeric value as the base index

   Example:1000
6. Select Make indexes contiguous

   Make indexes contiguous

   Select the Make indexes contiguous option to make the existing indexes sequential.

   Example: 1001, 1002, 1003, 1004 and so on.

   Note: If this option is not selected then the gaps between indexes are retained.

   Example: 1001, 1010, 1020, 1030 and so on.
7. Click Ok to confirm the reindex.

   The monitors are reindexed.

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 one or multiple monitors 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.

(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 182). To delete monitors from the Policy, see Delete Monitors from Templates or a Policy (see page 197).
2. Define MIB Extensions (see page 194).
3. Define SystemEDGE Policy Control Settings (see page 171).

   The policy or templates are updated.

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.

   Update machines running this policy/template

   Lets you apply the policy to systems that are already running the policy or template.

   Apply to machines not running this policy/templates

   Lets you apply policy or a template to systems.

4. (Options for policies) Do one of the following options 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:
     - Machines running an out-of-date version of the policy
     - Machines where policy exceptions have been applied
     - Machines running current version of the policy
     - Machines with configuration errors for this policy

     Policy exceptions occur when a user applies a point configuration change to an agent that is not represented in the applied policy.
   - Select 'Advanced (manually select machines)' to add the machines manually in the Select Machines pane to which you want to reapply the policy.

5. (Options for templates) Select one of the following options from the 'Update machines running this template' tab:

   Under Existing Machines, select one of the following options:

   - Update all machines with this template applied.
   - Update only those machines that do not have the latest changes of this template applied.
   - Update only those machines where the template has not been successfully applied.
   - Advanced (manually select machines)
   - Remove this template from machines.

6. (Optional) Select systems from the 'Apply to Machines not running this policy/template' tab for applying the policy or template.
7. Click Apply Policy or Apply Template.

   The application is initiated.

8. 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 server 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 where 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 182). To delete monitors from the Policy, see Delete Monitors from Templates or a Policy (see page 197).
2. Define MIB Extensions (see page 194).
3. Define SystemEDGE Policy Control Settings (see page 171).

   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.

How to Create and Apply an Autowatcher to a System

This scenario describes how a system administrator can use Autowatchers to dynamically monitor the resources on a managed system.

You can use Autowatchers to discover the resources that are added or removed on a managed system. If a resource is added, Autowatchers create a corresponding monitor. If a resource is removed, Autowatchers perform a 'Loss Action'.

The following diagram provides an overview of how to create and apply an Autowatcher to a managed system.

How to Create and Apply an Autowatcher to a System

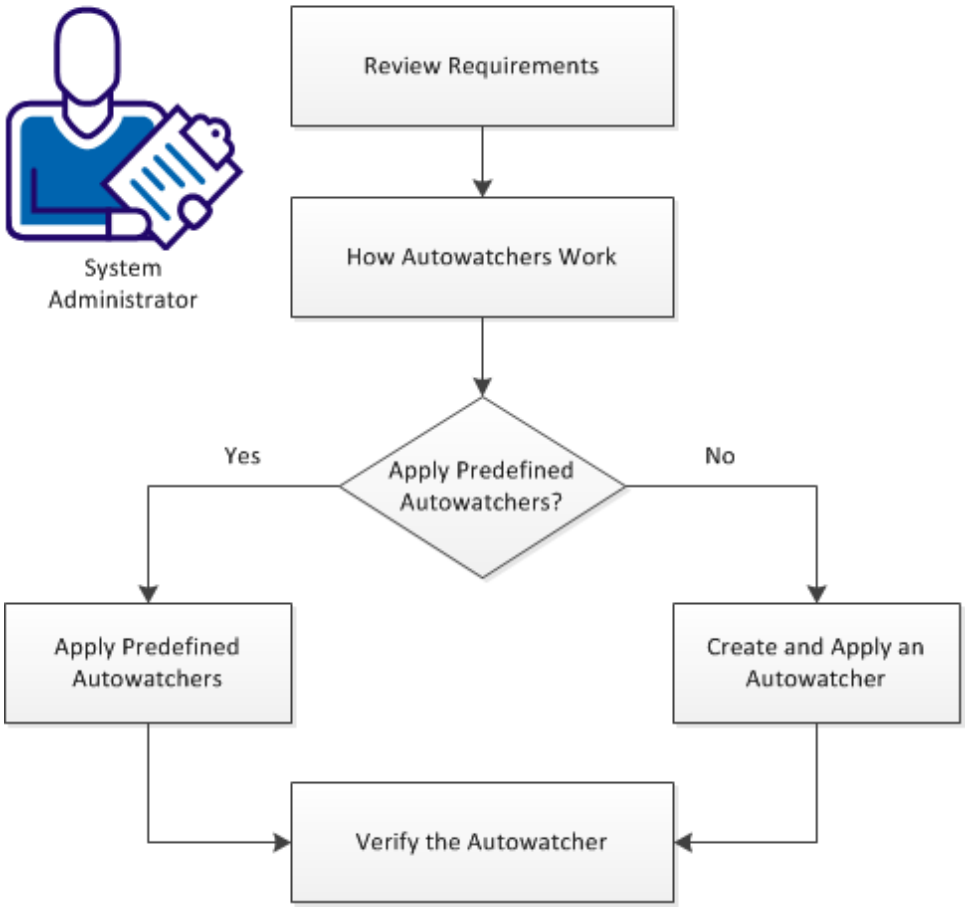

Follow these steps:

Review Requirements

Review the following requirements before you create an Autowatcher for SystemEDGE:

- You are familiar with TCP/IP and SNMP.
- You have a basic understanding of CA Virtual Assurance and SystemEDGE.
- You can access the CA Virtual Assurance user interface.
- Verify that the affected SystemEDGE agents are running in managed mode.

How Autowatchers Work

Autowatchers run periodic discovery processes using regular expressions as patterns to match the names of resources for which the Autowatchers create monitors. Autowatchers enable SystemEDGE to create automatically monitors for new resources when they come online. Autowatchers create monitors in a reserved range of indexes (1000000 - 1999999).

SystemEDGE sends traps to CA Virtual Assurance when resources disappear and applies the ‘Loss Actions’ that you have configured in the Autowatcher. In case of the loss of the monitored resource, the 'Loss Action' can remove the monitor or can set the status of the resource to a specific status:

OK, Warning, Minor, Major, Critical, Fatal, Up, or Down

Autowatchers enable you to create flexible Policies or Layered Templates without knowing what resources exist on a managed system. A resource can be a device, a service, or a process running on a managed system.

You can use the following Autowatcher Types:

- Generic Autowatchers – Creates monitors for various resources on a managed system, for example, for devices, interfaces, filesystems, or files.
- Process and Service Autowatchers – Creates monitors for processes and services running on a managed system.

Process Workflow of an Autowatcher

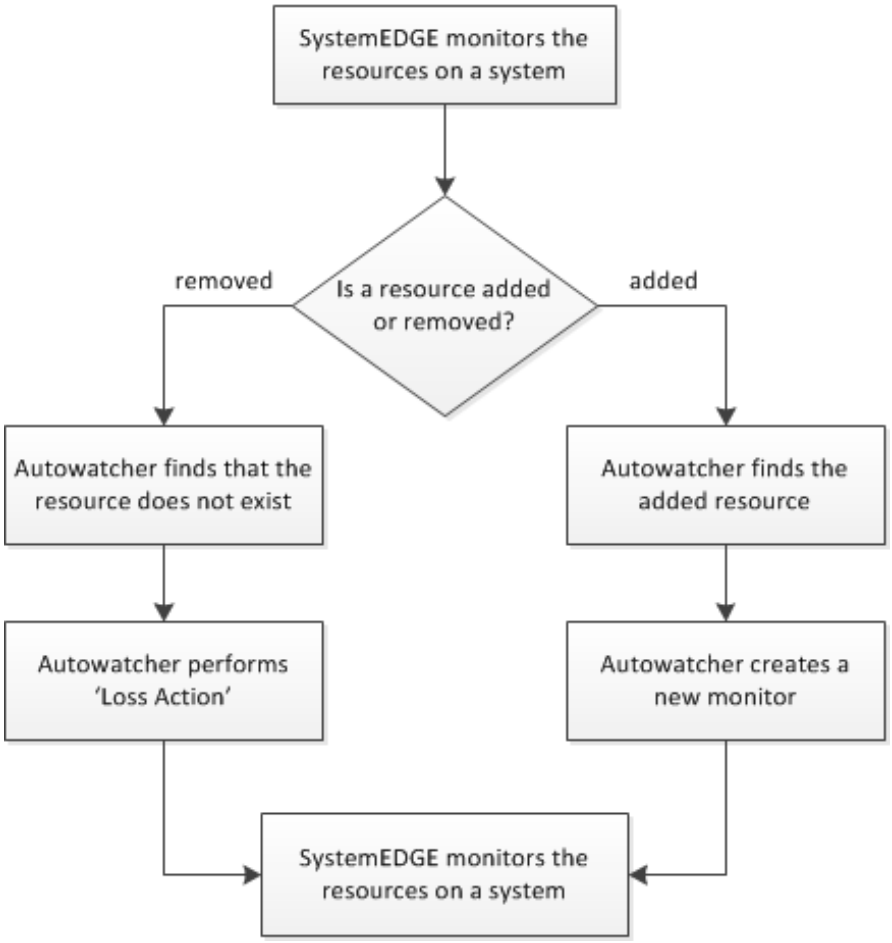

You can use the following guidelines when you configure Loss Actions:

- If the lost resource affects the health of a system, you can configure the Loss Action to change the status of the corresponding resource to a critical state.
- If the lost resource does not affect the health of a system, you can configure the Loss Action to remove the corresponding monitor.

More information:

Generic Autowatchers (see page 205)
Process and Service Autowatchers (see page 205)

Generic Autowatchers

Generic Autowatchers can create monitors for various resources on a managed system, for example, for devices, interfaces, filesystems, or files.

The following list provides you some examples of Generic Autowatchers:

- Capacity of all discovered devices
- Disk service time on all discovered disks
- Resident set size on all cmd processes
- Operating status of all tunneled network interfaces
- Device status of all devices

More information:

How Autowatchers Work (see page 203)

Process and Service Autowatchers

Use Process and Service Autowatchers to create process and service monitors dynamically.

A Service Autowatcher creates multiple service monitors in the process table whenever a service matches an Autowatcher criteria (service name, start type, and so on). For example, you can monitor all installed the SQL services with a start type as 'automatic'.

A Process Autowatchers create process monitors in two ways:

- Using a process name (default) - When a process name matches the Autowatcher criteria.

  For example, a process monitor is created when a process matches a criteria of process name of 'sql' or 'svchost'. Autowatcher-created process monitors track a matching process that currently runs on a managed system, regardless of PID.

  - Autowatcher-created process monitors have the same semantics as manually created process monitors.
  - You can individually monitor a set of processes with the same name, but different arguments. For example, "java.exe".
  - You can create monitors for a set of related processes.

- Using PID - When a PID matches the Autowatcher Criteria. The Autowatcher enables the Monitor Process using the PID flag in the user interface or specifies the watch flag 0x1000 in the sysedge.cf file.

  Each Autowatcher-created monitor tracks all matching instances of a process.

  - Creates monitors for particular instances of processes.
  - Monitors multiple instances of a process with no distinguishing arguments.

More information:

How Autowatchers Work (see page 203)

Apply Predefined Autowatchers

Policy Configuration provides the following predefined Autowatchers in templates and the SystemEDGE default policy:

- CPU Utilization (OS Independent Template)
- CA ARCserve (Windows Template)
- Windows Services (Windows Template)
- Microsoft Exchange (Windows Template)
- All Filesystems (SystemEDGE Default Policy)
- All Disks (SystemEDGE Default Policy)

Verify, if you can use predefined Autowatchers.

Follow these steps:

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

   The Available Policies pane or Template List opens displaying the predefined Autowatchers.

2. In the Available Policies pane or Template List, click the predefined Autowatchers.

   The Autowatchers Details pane opens.

3. Click Action, Apply.

   The machine selection page opens.

4. Select the appropriate systems and click Apply.

   The Autowatcher is added to the SystemEDGE configurations of the selected systems.

   SystemEDGE creates monitors automatically based on the Autowatcher settings.

   Note: For SystemEDGE in unmanaged mode, specify the Autowatchers in the sysedge.cf file. When SystemEDGE is changed to managed mode, Autowatchers that are defined before SystemEDGE registers with CA Virtual Assurance can be imported into a policy.

More information:

Create and Apply an Autowatcher to a System (see page 207)

Create and Apply an Autowatcher to a System

For a SystemEDGE in managed mode, you can specify an Autowatcher in a policy or in a template. The centralized configuration provides consistent monitoring across all servers. Configure an Autowatcher in a policy or a template and apply the Autowatcher to monitor resources on managed systems.

Follow these steps:

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

   The Available Policies pane or Template List opens.

2. Open a policy or template and click Autowatchers.

   The Generic Autowatchers pane opens.

3. Select the Process/Service tab if you want to add Process or Service Autowatchers.

4. Click  (Add) on the toolbar.

   The Autowatcher Details pane opens.

5. Specify the required values and click Save.

   The Autowatcher is saved.

6. Click Action, Apply.

   The machine selection page opens.

7. Select the appropriate systems and click Apply.

   The Autowatcher is added to the SystemEDGE configurations of the selected systems.

   SystemEDGE creates monitors automatically based on the Autowatcher settings.

   Note: For SystemEDGE in unmanaged mode, specify the Autowatchers in the sysedge.cf file. When SystemEDGE is changed to managed mode, Autowatchers that are defined before SystemEDGE registers with CA Virtual Assurance can be imported into a policy.

More information:

Verify the Autowatcher (see page 208)

Verify the Autowatcher

In the CA Virtual Assurance user interface, you can verify if the Autowatcher has created the corresponding monitors for resources. Autowatchers create monitors in the reserved range of indexes (1000000 - 1999999).

Follow these steps:

1. Click the Resources tab.

   The Resources page appears.

2. Expand the Data Center folder and the CA Virtual Assurance Services folder in the Explore pane.

   The discovered and managed resources in the data center appear.

3. Select the resource for which you want to verify the corresponding monitor.

   The Quick Start tasks for the selected resource appears.

4. Click the Configuration tab.

   The Self Monitor page appears and displays the monitors that Autowatcher creates in the reserved range of indexes.

More information:

How Autowatchers Work (see page 203)

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.

How to Monitor a Specific Windows Performance Registry Metric

The following example describes how to monitor a user-specific metric. The names used in the Windows Performance object and counter must match the names in perfmon.exe

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 the Resources tab, open the Configuration pane, expand the Policy tree, and click an appropriate subcategory.

       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 data into the fields:
       Example:

       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.

       Note: You can specify custom entries for Object and Counter while creating a policy. The same metrics are saved for future use while creating another policy.

    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. Click + (New) to create a monitor.

       The Threshold Monitor Details: New dialog appears.

c. 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 as 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.

d. Click Save.

The monitor is added to the policy.

3. Click Action and then Apply, to apply the policy to a computer.

The Selected Machines pane appears.

a. Verify that the selected machines are correct and click Apply.

The policy with the MIB extension is applied to the selected computers.

b. 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 under Explore, Summary.

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.
5. Define test thresholds.
6. Define control settings (see page 255).
7. Click Save Policy.

   The policy is saved.

Discovering the Agents

When an agent has multiple NICs (network interface controller), Policy Configuration discovers all the name or addresses for that agent. To avoid discovering the unwanted names and addresses, Policy Configuration supports to discover the agents with management names or addresses to deploy a job.

Note: The system refreshes the discover agents list for every 30 minutes.

Follow these steps:

1. Log in to the CA Virtual Assurance application and click the Resource tab.
2. From the Explorer tab, right-click a Domain Server and select Policy, SystemEDGE, Discover Agents.

   A confirmation dialog opens.
3. Click OK.

   Note: To view the list, click Monitoring Software tab and then click the Policy tab. The list of available agents with management names or addresses is displayed.

Common Usage of Policy Configuration Functions

This section describes the common Policy Configuration functions.

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.
5. Define control settings (see page 218).
6. Define SNMP settings (see page 178).
7. Define MIB extensions.
8. Click Save Policy.

   The policy is saved.

Copy SystemEDGE Policy

You can copy an existing SystemEDGE policy.

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 you want to copy in the Available Policies table, click Actions and select Copy. You can also right-click the policy in the Configure pane and select Copy.

   The Copy dialog appears.
3. Enter a new name for the policy and click Ok.

   The policy is copied and a configuration screen appears in the right pane.
4. Click Save Policy.

   The policy is saved.

Rename SystemEDGE Policy

You can rename an existing SystemEDGE policy.

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 you want to rename in the Available Policies table, click Actions and select Rename. You can also right-click the policy in the Configure pane and select Rename.

   The Rename dialog appears.

   Note: If the policy is in use, an error message is displayed indicating that the policy cannot be renamed.

3. Enter a new name for the policy and click Ok.

   A confirmation message appears notifying you that the policy is renamed.

4. Click Save Policy.

   The policy is saved.

Delete SystemEDGE Policy

You can delete an existing SystemEDGE policy.

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 you want to delete in the Available Policies table, click Actions, and select Delete. You can also right-click the policy in the Configure pane and select Delete.

   Note: If the policy is in use, an error message appears indicating that the policy cannot be deleted.

   A warning message appears.

3. Click Ok to confirm the deletion.

   A confirmation message appears. The policy is deleted.

Import a SystemEDGE Configuration to a Policy

After upgrading SystemEDGE to the current version, import the previous SystemEDGE configuration, and convert it to a SystemEDGE 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. Click Import.

   The SystemEDGE Agent Machines window appears.

4. Select the computer you want to import a SystemEDGE configuration from, and click OK.

   Note: The machine list displays all computers that are upgraded from original configuration file, with monitors defined. The computer appears in the list once SystemEDGE 5.x is discovered and is registered with Policy Configuration. If a computer is not listed, verify if it has monitors defined at previous SystemEDGE version levels, and is configured with Policy Configuration.

5. Enter a name and an optional description to the New SystemEDGE Policy dialog, and click OK to complete the import process.

6. Click Save Policy.

   The policy is saved.

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 223).

   The control settings are defined.

9. Click Save Policy.

   The policy is saved.

More Information:

Configure Object Aggregation (see page 223)

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 218)

Define New SystemEDGE Monitoring Template

You can configure the SystemEDGE with different policies. Monitoring Templates lets you configure and deliver multiple policies to the same agent on the shared server.

The Monitoring Templates page lets you view and update the policies applied to a specific server or server group. You can create SystemEDGE monitoring templates (layered templates) and imported into a policy. This lets you reuse monitors across multiple policies without the need to set up monitors multiple times.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates, then click SystemEDGE.

   The SystemEDGE 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. To add a monitor to the template, see the section Add a Monitor To SystemEDGE Policy. (see page 232)

4. Click Save Template.

   The template is saved.

More Information:

Layered Templates (see page 226)
Import a Monitoring Template to SystemEDGE Policy (see page 227)
Copy SystemEDGE Monitoring Template (see page 228)
Modify SystemEDGE Monitoring Template (see page 228)
Rename SystemEDGE Monitoring Template (see page 229)
Delete SystemEDGE Monitoring Template (see page 229)
Review Monitoring Template Application Progress (see page 230)
Apply Templates to Machines (see page 230)
Rename SystemEDGE Monitoring Template (see page 229)
Modify SystemEDGE Monitoring Template (see page 228)
Apply Templates to Machines (see page 230)
Review Monitoring Template Application Progress (see page 230)
Layered Templates (see page 226)

Layered Templates

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

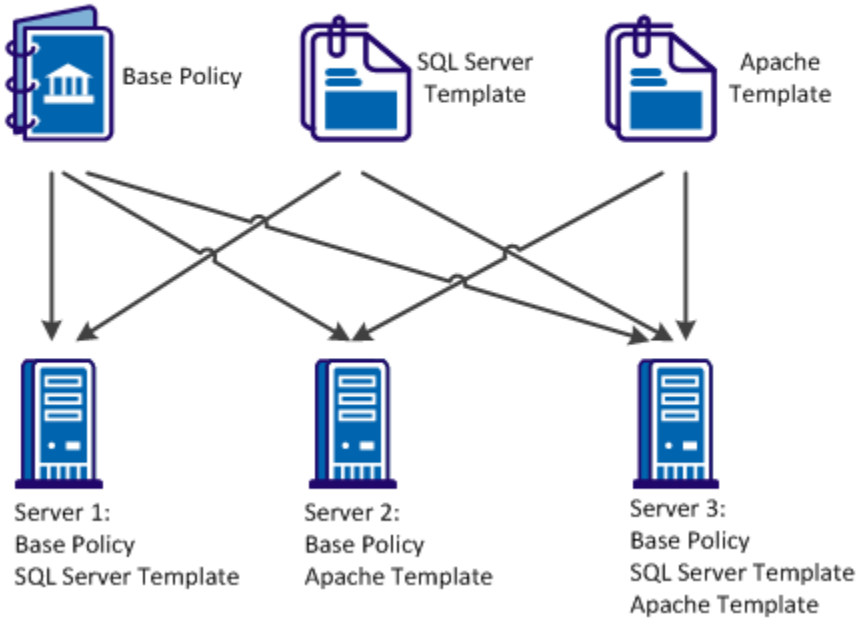

You can use Layered Templates in the following scenarios:

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.

Out of the box templates

You can apply out of the box templates to managed nodes. Configure the policy with the template configuration on the managed nodes. Templates are available for the following operating systems:

For All Operating Systems:

CPU Utilization - Autowatch

Swap Capacity

For Windows:

App Monitoring - CA eTrust Antivirus

Process Crash

System Errors

System Processes

User Activity

Windows Services - Autowatch

For UNIX (AIX, HPUnix, Linux, Solaris):

System Messages

System Processes

User Activity

Import a Monitoring Template to SystemEDGE Policy

You can import a monitoring template into SystemEDGE policy. This replaces the existing policy of all systems with a consistent policy at one operation.

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 Monitors tab.

   The Summary page appears with a list of monitors managed by the policy.

4. Click Action and select Import.

   The Import Template Wizard appears.

5. Select the System Type and the monitoring template you want to import from the drop-down lists.

6. (Optional) Define a new base index for each of the imported monitors.

7. Select a Conflict Resolution Option from the drop-down list and click Next.

   The Resolve Conflict page appears.

8. Review any monitor conflicts and make adjustments to the indexes, then click Next.

   The Summary page appears.

9. Review the monitors that will be imported, then click Finish to complete the import process.

10. Click Save Policy.

    The policy is saved.

Copy SystemEDGE Monitoring Template

You can copy an existing SystemEDGE monitoring template.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates, then click SystemEDGE.

   The summary page displays a list of the SystemEDGE Monitoring Templates.

2. Select the monitoring template you want to copy, click Action and select Copy. You can also right-click the monitoring template in the Configure pane and select Copy.

   The Copy dialog appears.

3. Enter a new name for the monitoring template and click Ok.

   The monitoring template is copied and a configuration screen appears in the right pane.

Modify SystemEDGE Monitoring Template

You can modify the SystemEDGE monitoring template.

Follow these steps:

1. Click Resources tab, open the Configure pane, expand Monitoring Templates, and then click 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 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 modify.

   The Edit dialog appears.

5. Modify the settings according to your needs and click Save.

6. (Optional) Repeat the process for any additional monitors.

7. Click Save.

   The Monitoring Template is saved.

Rename SystemEDGE Monitoring Template

You can rename an existing SystemEDGE monitoring template.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates, then click SystemEDGE.

   The summary page displays a list of the SystemEDGE Monitoring Templates.

2. Select the monitoring template you want to rename, click Action and select Rename. You can also right-click the monitoring template in the Configure pane and select Rename.

   The Rename dialog appears.

3. Enter a new name for the monitoring template and click Ok.

   The monitoring template is renamed and a configuration screen appears in the right pane.

Delete SystemEDGE Monitoring Template

You can delete an existing SystemEDGE monitoring template that you no longer need.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates, then click SystemEDGE.

   The summary page displays a list of the SystemEDGE Monitoring Templates.

2. Click the Managed Machines tab

   The Summary page appears with a list of managed machines applied to the template.

3. Select the monitoring template you want to delete, click Delete icon.

   A confirmation message appears.

4. Click Ok to confirm the deletion.

   The monitoring template is deleted.

Review Monitoring Template Application Progress

You can review the progress of monitoring template application operations at a detailed level for each individual template.

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 the Managed Machines tab.

   The Managed Machines page appears with a list of machines currently running the monitoring template that lets you view the configuration status.

4. (Optional) Click View Configuration.

   The SystemEDGE Configuration pane appears and lets you view the Policies and Templates, and Configuration file delivered for the agent.

Apply Templates to Machines

After you update monitoring templates, you can apply it to machines across the enterprise.

Follow these steps:

1. Click the Resources pane, 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 one of the following options:

   - Update all machines with this template applied.
   - Update only those machines that do not have the latest changes of this template applied.
   - Update only those machines where the template has not been successfully applied.
   - Advanced (manually select machines)
   - Remove this template from machines.

5. (Optional) Under Selected Machines, select the machines to which the template is reapplied.

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.

Import a SystemEDGE Configuration to a Template

After upgrading SystemEDGE to the current version, import the previous SystemEDGE configuration, and convert it to a SystemEDGE monitoring template.

Follow these steps:

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

   The Available SystemEDGE Monitoring Templates page appears.

2. Click + (New) on the Available SystemEDGE Monitoring Templates toolbar.

   The New SystemEDGE Monitoring Template dialog appears.

3. Click Import.

   The SystemEDGE Agent Machines window appears.

4. Select the computer you want to import a SystemEDGE configuration from, and click OK.

   Note: The machine list displays all computers that are upgraded from original configuration file, with monitors defined. The computer appears in the list once SystemEDGE 5.x is discovered and is registered with Policy Configuration. If a computer is not listed, verify if it has monitors defined at previous SystemEDGE version levels, and is configured with Policy Configuration.

5. Enter a name and an optional description to the New SystemEDGE Monitoring Template dialog, and click OK to complete the import process.

6. Click Save Template.

   The template is saved.

Add a Monitor To SystemEDGE Policy

You can add a monitor to a SystemEDGE policy.

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 Monitors and select the monitor you want to add.

   - Define a Threshold Monitor (see page 233)
   - Define a Process Monitor (see page 235)
   - Define a Log File Monitor (see page 237)
   - Define a Windows Event Monitor (see page 239)
   - Define a History Monitor (see page 240)
   - Define a Process Group Monitor (see page 242)

4. (Optional) Repeat the process for any additional monitors

5. Click Save Policy.

   The monitor is loaded to the policy and the policy is saved.

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

More Information:

Define a Threshold Monitor (see page 233)
Define a Process Monitor (see page 235)
Define a Log File Monitor (see page 237)
Define a Windows Event Monitor (see page 239)
Define a History Monitor (see page 240)
Define a Process Group Monitor (see page 242)

Define a Threshold Monitor

You can define threshold settings for the SystemEDGE policy.

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 Monitors.

   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 to be used.

   Platform

   Specifies the platform.

Description

Defines an optional description.

Object Class

Specifies the object class to monitor. The values in the drop-down list refer to the available MIB tables.

Object Class Name

Defines the object class name to use for the object state model. This is an arbitrary string, for example, FileSystems.

Object Attribute

Specifies the object attribute to monitor. The values in the drop-down list 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. This 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.

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

7. Click Save

   The Threshold Monitor settings are saved.

8. Click Save Policy.

   The Threshold Monitor is loaded to the Policy.

Define a Process Monitor

You can define process settings for the SystemEDGE policy.

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 Monitors.

   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 to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class Name

Specifies the object class name to use for the object state model. This is an arbitrary string, for example, Process.

Object Attribute

Specifies the object attribute to monitor. The values in the drop-down list define the available attributes for process monitoring.

Object Attribute Name

Defines the object attribute name to use for the object state model. This 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. This 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.

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

7. Click Save

   The Process Monitor settings are saved.

8. Click Save Policy.

   The Process Monitor is loaded to the Policy.

Define a Log File Monitor

You can define log file settings for the SystemEDGE policy.

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 Monitors.

   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 to be used.

   Monitor Type

   Specifies the monitor type to be used.

   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.

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

7. Click Save

   The Log File Monitor settings are saved.

8. Click Save Policy.

   The Log File Monitor is loaded to the Policy.

Define a Windows Event Monitor

You can define Windows event settings for the SystemEDGE policy.

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 Monitors.

   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 Widows Event Monitors toolbar.

   The Windows Event Details: New dialog appears.

6. Configure the following process settings:

   Index

   Defines the table index to be used.

   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.

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

7. Click Save

   The Windows Event Monitor settings are saved.

8. Click Save Policy.

   The Windows Event Monitor is loaded to the Policy.

Define a History Monitor

You can define history settings for the SystemEDGE policy.

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 Monitors.

   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 to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class

Specifies the object to monitor. The values in the drop-down list refer to the available MIB table.

Object Attribute

Specifies the object attribute to monitor. The values in the drop-down list 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, .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.

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

7. Click Save

   The History Monitor settings are saved.

8. Click Save Policy.

   The HIstory Monitor is loaded to the Policy.

Define a Process Group Monitor

You can define process group settings for the SystemEDGE policy.

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 Monitors.

   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 to be used.

   Platform

   Specifies the platform.

   Description

   Defines an optional description.

Process Name

Defines the process name. This 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.

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

7. Click Save

   The Process Group Monitor settings are saved.

8. Click Save Policy.

   The Process Group Monitor is loaded to the Policy.

View Monitors Within a SystemEDGE Policy

You can view the monitors contained within a SystemEDGE policy.

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 Monitors tab.

   The Summary page appears with a list of monitors managed by the policy. You can click the monitoring classes subtabs to view the different monitors that are contained within the policy.

More Information

Delete a Monitor from SystemEDGE Policy (see page 245)
Modify a Monitor Within SystemEDGE Policy (see page 245)

Copy a Monitor Within SystemEDGE Policy

You can copy a monitor within a SystemEDGE policy.

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 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 copy.
5. Click Action and select Copy.

   The Edit dialog appears.

6. Modify the settings according to your needs and click Save.
7. (Optional) Repeat the process for any additional monitors.
8. Click Save Policy.

   The policy is saved.

Modify a Monitor Within SystemEDGE Policy

You can modify a monitor within a SystemEDGE policy.

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 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 modify.
5. Click Action and select Modify.

   The Edit dialog appears.

6. Modify the settings according to your needs and click Save.
7. (Optional) Repeat the process for any additional monitors.
8. Click Save Policy.

   The policy is saved.

Delete a Monitor from SystemEDGE Policy

You can delete a monitor from a SystemEDGE policy.

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 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 policy is saved.

Modify Existing Template in SystemEDGE Policy

You can modify existing monitoring template and import into SystemEDGE policy.

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 Monitors tab.

   The Summary page appears with a list of monitors managed by the policy.
4. Click Action and select Import.

   The Import Template Wizard appears.
5. Update the monitoring template with the required information.
6. Select the System Type and the updated monitoring template to import from the drop-down lists.
7. (Optional) Define a new base index for each of the imported monitors.
8. Select the Conflict Resolution Option as "Replace existing monitors with imported entities" and click Next.

   The Resolve Conflict page appears.
9. Review any monitor conflicts and make adjustments to the indexes, then click Next.

   The Summary page appears.
10. Review the monitors that will be imported, then click Finish to complete the import process.
11. Click Save Policy.

    The policy is saved.
12. Select Apply from Action drop-down list.

    The saved policy will be applied to desired machines.

Define New SRM Policy

You can 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.

Follow these steps:

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

   The Available Policies page appears.

2. Click + (New) on the Available Policies toolbar.

   The New Service Response Monitoring 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 saved.

More Information:

Copy SRM Policy (see page 247)
Rename SRM Policy (see page 248)
Delete SRM Policy (see page 248)

Copy SRM Policy

You can copy an existing SRM 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 you want to copy in the Available Policies table, click Actions and select Copy. You can also right-click the policy in the Configure pane and select Copy.

   The Copy dialog appears.

3. Enter a new name for the policy and click OK.

   The policy is copied and a configuration screen appears in the right pane.

4. Click Save Policy.

   The policy is saved.

Rename SRM Policy

You can rename an existing SRM 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 you want to rename in the Available Policies table, click Actions and select Rename. You can also right-click the policy in the Configure pane and select Rename.

   The Rename dialog appears.

   Note: If the policy is in use, an error message is displayed indicating that the policy cannot be renamed.

3. Enter a new name for the policy and click OK.

   A confirmation message appears notifying you that the policy is renamed.

4. Click Save Policy.

   The policy is saved.

Delete SRM Policy

You can delete an existing SRM 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 you want to delete in the Available Policies table, click Actions and select Delete. You can also right-click the policy in the Configure pane and select Delete.

   Note: If the policy is in use, an error message is displayed indicating that the policy cannot be deleted.

3. A warning message appears. Click OK to confirm the deletion.

4. Click Save Policy.

   The policy is saved.

Add a Test to SRM Policy

You can add a test to an SRM 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 Test and click + (New) on the Test Monitors toolbar.

   The New dialog appears.

4. Specify a unique name for the test in the Test name field. The name must be 64 characters or less. Test names are case-sensitive.

5. (Optional) Enter a Description for the test.

6. (Optional) Define a test class.

7. Click the test type you want in the Test Type list:

   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 that 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.

Note: For more information and definitions of each test type, see the *SRM User Guide*.

8. Specify the interval (in seconds) between tests in the Test Interval field. The interval must be a multiple of 30 seconds. Use this option for tuning the performance of your tests.
9. In the Test Timeout field, specify the time (in seconds) after which the test should time out. Select a number that is less than the interval but greater than the amount of time that the test requires to execute.
10. Set the polling interval by selecting one of the following from the Polling Interval list:
    - Normal
    - Off
    - Slow

    Note: For more information, see the *SRM User Guide*
11. (Optional) Select the Keep Historical Data check box
12. Click Save to add the test to the policy.

    The test is saved.
13. (Optional) Repeat the process for any additional tests.
14. Click Save Policy.

    The policy is saved.

More Information:

Define SRM Control Settings (see page 255)
Add a Threshold Definition To SRM Policy (see page 253)
Modify SRM Test (see page 251)
Modify SRM Threshold Definition (see page 254)

Modify SRM Test

You can modify an existing SRM test.

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 containing the test you want to modify in the Available Policies table

   The Summary page appears.

3. Click the Test tab.

   The Test Monitors page appears.

4. Select the test you want to modify, click Actions, and select Modify.

   The Edit dialog appears.

5. Modify the test according to your needs and click Save.

   The test is updated.

6. Click Save Policy.

   The policy is saved.

Copy SRM Test

You can copy an existing SRM test.

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 containing the test you want to copy in the Available Policies table

   The Summary page appears.

3. Click the Test tab.

   The Test Monitors page appears.

4. Select the test you want to copy, click Actions, and select Copy.

   A copy dialog appears.

5. Enter the SRM test name.

   The SRM test is copied.

Delete SRM Test

You can delete an existing SRM test.

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 containing the test you want to delete in the Available Policies table

   The Summary page appears.

3. Click the Test tab.

   The Test Monitors page appears.

4. Select the test you want to modify, click Actions, and select Delete.
5. Confirm your action.

   The SRM test is deleted.

Add a Threshold Definition To SRM Policy

You can add a threshold definition to an SRM 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 Threshold tab and Click + (New) on the Threshold Monitors toolbar.

   The Threshold Monitor Details dialog appears.

4. Configure the following threshold monitor settings:

   Name

   Defines the threshold monitor name.

   Attribute

   Specifies the attribute to use.

   Operator

   Specifies the operator to use.

   Warning Value

   Defines the warning value to use.

   Minor Value

   Defines the minor value to use.

   Major Value

   Defines the major value to use.

Critical Value

Defines the critical value to use.

Fatal Value

Defines the fatal value to use.

5. Click Save to add the threshold definition to the policy.

   The threshold definition is saved.

6. Click Save Policy.

   The policy is saved.

Modify SRM Threshold Definition

You can modify an existing SRM threshold definition.

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 containing the threshold definition you want to modify in the Available Policies table

   The Summary page appears.

3. Click the Threshold tab.

   The Threshold Monitors page appears.

4. Select the threshold definition you want to modify, click Actions, and select Modify.

   The Edit dialog appears.

5. Modify the threshold definition according to your needs and click Save.

   The threshold definition is updated.

6. Click Save Policy.

   The policy is saved.

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.

Define New SRM Test Definition Template

You can create an SRM test definition template that can be imported into a policy. This lets you reuse tests across multiple policies without the need for setting up tests multiple times.

To define a new SRM test definition template

1. Click Resources tab, open the Configure pane, expand Monitoring Templates, then click Service Response.

   The Service Response page appears

2. Click + (New) on the Test Templates List toolbar.

   The New Service Response Test Definition Template dialog appears.

3. Enter a name and an optional description for the test definition template, and whether to base it on an existing template and click OK.

   The test definition template is created, and the Summary page appears. To add a test to the template, see the section Add a Test to SRM Policy (see page 249).

4. Click Save Template.

   The template is saved.

More Information:

Modify SRM Test Definition Template (see page 258)
Copy SRM Test Definition Template (see page 258)
Rename SRM Test Definition Template (see page 258)
Delete SRM Test Definition Template (see page 259)
Import a Test Definition Template into SRM Policy (see page 257)

Import a Test Definition Template into SRM Policy

You can import a test definition template into an SRM policy.

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 Test tab.

   The Summary page appears with a list of test monitors managed by the policy.

4. Click Action and select Import.

   The Import Template Wizard appears.

5. Select the test template you want to import from the drop-down lists.
6. (Optional) Define a new base index for each of the imported test definitions.
7. Select a Conflict Resolution Option from the drop-down list and click Next.

   The Resolve Conflict page appears.

8. Review any test definition conflicts and make adjustments to the indexes, uncheck any test definitions that should not be imported, then click Next.

   The Summary page appears.

9. Review the test definitions that will be imported, then click Finish to complete the import process.
10. Click Save Policy.

    The policy is saved.

Modify SRM Test Definition Template

You can modify an SRM test definition template.

To modify an SRM test definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Test Definition Templates.

   The Test Template List appears with a list of test templates.

2. Select the Service Response test template you want to modify.

   The Summary page appears for the test template.

3. Click the Tests tab, select the test monitor you want to modify, click Action, and select Modify.

   The Edit dialog appears.

4. Modify the settings according to your needs and click Save.

5. Click Save Template.

   The template is saved.

Copy SRM Test Definition Template

You can copy an SRM test definition template.

To copy an SRM test definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Test Definition Templates.

   The Test Template List appears with a list of test templates.

2. Select the Service Response test template you want to copy.

   The Summary page appears for the test template.

3. Click the Tests tab, select the test you want to copy, click Action, and select Copy. You can also right-click the test template in the Configuration pane and select Copy.

   The Copy dialog appears.

4. Enter a new name for the test definition template and click Ok.

   The test definition template is copied and appears in the Test Templates list.

Rename SRM Test Definition Template

You can rename an SRM test definition template.

To rename an SRM test definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Test Definition Templates.

   The Test Template List appears with a list of test templates.

2. Select the Service Response test template you want to rename.

   The Summary page appears for the test template.

3. Click the Tests tab, select the test you want to rename, click Action, and select Rename. You can also right-click the test template in the Configure pane and select Rename.

   The Rename dialog appears.

4. Enter a new name for the test definition template and click Ok.

   A confirmation message appears notifying you that the test definition template is renamed.

Delete SRM Test Definition Template

You can delete an SRM test definition template.

To delete an SRM test definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Test Definition Templates.

   The Test Template List appears with a list of test templates.

2. Select the Service Response test template you want to delete.

   The Summary page appears for the test template.

3. Click the Tests tab, select the test you want to delete, click Action, and select Delete. You can also right-click the test template in the Configure pane and select Delete.

   A warning message appears.

4. Click Ok to confirm the deletion.

   A confirmation message appears. The test template is deleted.

Define New SRM Threshold Definition Template

You can create an SRM threshold definition template that can be imported into a policy. This lets you reuse thresholds across multiple policies without the need for setting up thresholds multiple times.

To define a new SRM threshold definition template

1. Open the Configure pane, expand Monitoring Templates, then click Service Response.

   The Service Response page appears

2. Click + (New) on the Threshold Templates List toolbar.

   The New Service Response Threshold Definition Template dialog appears.

3. Enter a name and an optional description for the threshold definition template, and whether to base it on an existing template and click OK.

   The threshold definition template is created, and the Summary page appears. To add a threshold definition to the template, see the section Add a Threshold definition To SRM Policy. (see page 253)

4. Click Save Template.

   The template is saved.

Import a Threshold Definition Template into SRM Policy

You can import a threshold definition template into an SRM policy.

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 Threshold tab.

   The Summary page appears with a list of threshold monitors managed by the policy.

4. Click Action and select Import.

   The Import Template Wizard appears.

5. Select the threshold template you want to import from the drop-down lists.

6. Select a how to handle index conflicts and click Next.

   The Resolve Conflict page appears.

7. Review any threshold definition conflicts, make adjustments to the threshold definition name and uncheck any threshold definitions that should not be imported, then click Next.

   The Summary page appears.

8. Review the threshold definitions that will be imported, then click Finish to complete the import process.

9. Click Save Policy.

   The policy is saved.

Modify SRM Threshold Definition Template

You can modify an SRM threshold definition template.

To modify an SRM threshold definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Threshold Definition Templates.

   The Threshold Template List appears with a list of test templates.

2. Select the Service Response threshold template you want to modify.

   The Summary page appears for the test template.

3. Click Thresholds, select the threshold monitor you want to modify, click Action, and select Modify.

   The Threshold Monitor Details dialog appears.

4. Modify the settings according to your needs and click Save.

5. Click Save Template.

   The template is saved.

Copy SRM Threshold Definition Template

You can copy an SRM threshold definition template.

To copy an SRM threshold definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Threshold Definition Templates.

   The Threshold Template List appears with a list of test templates.

2. Select the Service Response threshold template you want to copy.

   The Summary page appears for the threshold template.

3. Click the Threshold tab, select the threshold monitor you want to copy, click Action, and select Copy. You can also right-click the threshold template in the Configure pane and select Copy.

   The Copy dialog appears.

4. Enter a new name for the threshold definition template and click Ok.

   The threshold definition template is copied and appears in the Threshold Template list.

Rename SRM Threshold Definition Template

You can rename an SRM threshold definition template.

To rename an SRM threshold definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Threshold Definition Templates.

   The Threshold Template List appears with a list of threshold templates.

2. Select the Service Response threshold template you want to rename.

   The Summary page appears for the threshold template.

3. Click the Threshold tab, select the threshold monitor you want to rename, click Action, and select Rename. You can also right-click the test template in the Configure pane and select Rename.

   The Rename dialog appears.

4. Enter a new name for the threshold definition template and click Ok.

   A confirmation message appears notifying you that the threshold definition template is renamed.

Delete SRM Threshold Definition Template

You can delete an SRM threshold definition template.

To delete an SRM threshold definition template

1. Open the Configure pane, expand Monitoring Templates, click Service Response, and expand Threshold Definition Templates.

   The Threshold Template List appears with a list of test templates.

2. Select the Service Response threshold template you want to delete.

   The Summary page appears for the threshold template.

3. Click the Threshold tab, select the threshold monitor you want to delete, click Action, and select Delete. You can also right-click the threshold template in the Configure pane and select Delete.

   A warning message appears.

4. Click Ok to confirm the deletion.

   A confirmation message appears. The threshold template is deleted.

Import an Existing SRM Configuration

After upgrading an existing Service Availability (SA) 2.0 AIM to SRM 3.1.0, you can import the previous SA 2.0 configuration and convert it to an SRM 3.1.0 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. Click + (New) on the Available Policies toolbar.

   The New service Response Policy dialog appears.

3. Click Import, select the machine you want to import the policy from in the list, and click Ok.

4. Enter a name and an optional description, and click Ok to complete the import process.

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. The new policy is implemented after an automatic agent warm start.

If one of the following cases occurs, you can reapply the policy to machines:

- 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 that 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 policies to machines without any policy or using a different policy.

5. (Optional) Do one of the following options 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:
     - Machines running an out-of-date version of the policy
     - Machines where policy exceptions have been applied
     - Machines running current version of the policy
     - Machines with configuration errors for this 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 (manually select machines)' to add the machines manually in the Select Machines pane to which you want to reapply the policy.

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 since the last policy or template application.
- 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 displays a list of layered templates in Pending, Delivered (successful), Configured, and Failed states.

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.
   - Bulk remove Templates

   Note: If the policy is being applied to a single server, you are prompted for the policy name.

   Bulk update for Policies:

   When the selected policy is applied to a service group, you have an option to select the machine 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:

   Replace existing configurations with the selected templates

   Removes the existing templates that are applied to all machines and applies the selected templates to all machines.

   Append the selected templates to existing configurations

   Adds the selected templates. If any of the templates selected are already applied as part of a machine configuration, then those templates are reapplied.

   Bulk remove Templates:

   Removes the existing templates that are applied to machines.

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.

Revert a Policy Back To an Earlier Version

You can revert a policy back to an earlier version.

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. Click the Versions tab.

   The Versions page appears.

4. Select the version you want to restore to in the table and click Make Current.

   A message appears. Click Ok. A new version of the policy is created and the summary page appears.

5. (Optional) You can make a new copy of the version. Select the version in the Available Policies table and click Copy.

   The Copy dialog appears.

6. Enter a new name for the policy and click Ok.

   The policy is copied and added to the policy tree in the Configure pane. The summary page for the new copy appears.

7. Click Save Policy.

   The policy is saved.

Specify Default Policy for New Instances

You can set a single default policy for all new discovered instances. The policy will be delivered if a policy was not specified during installation or deployment of SystemEDGE or SRM, or if the specified policy is not available.

To specify default policy

1. Open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

   The Available Policies page appears.

2. In the Default Policy section, select the policy you want to use from the Default Policy drop-down list and click Apply

   The default policy is applied.

How to Change the Configuration Mode for SystemEDGE

In some cases, it can be necessary to change the configuration mode of SystemEDGE. The following diagram provides an overview of the required actions to change the configuration mode.

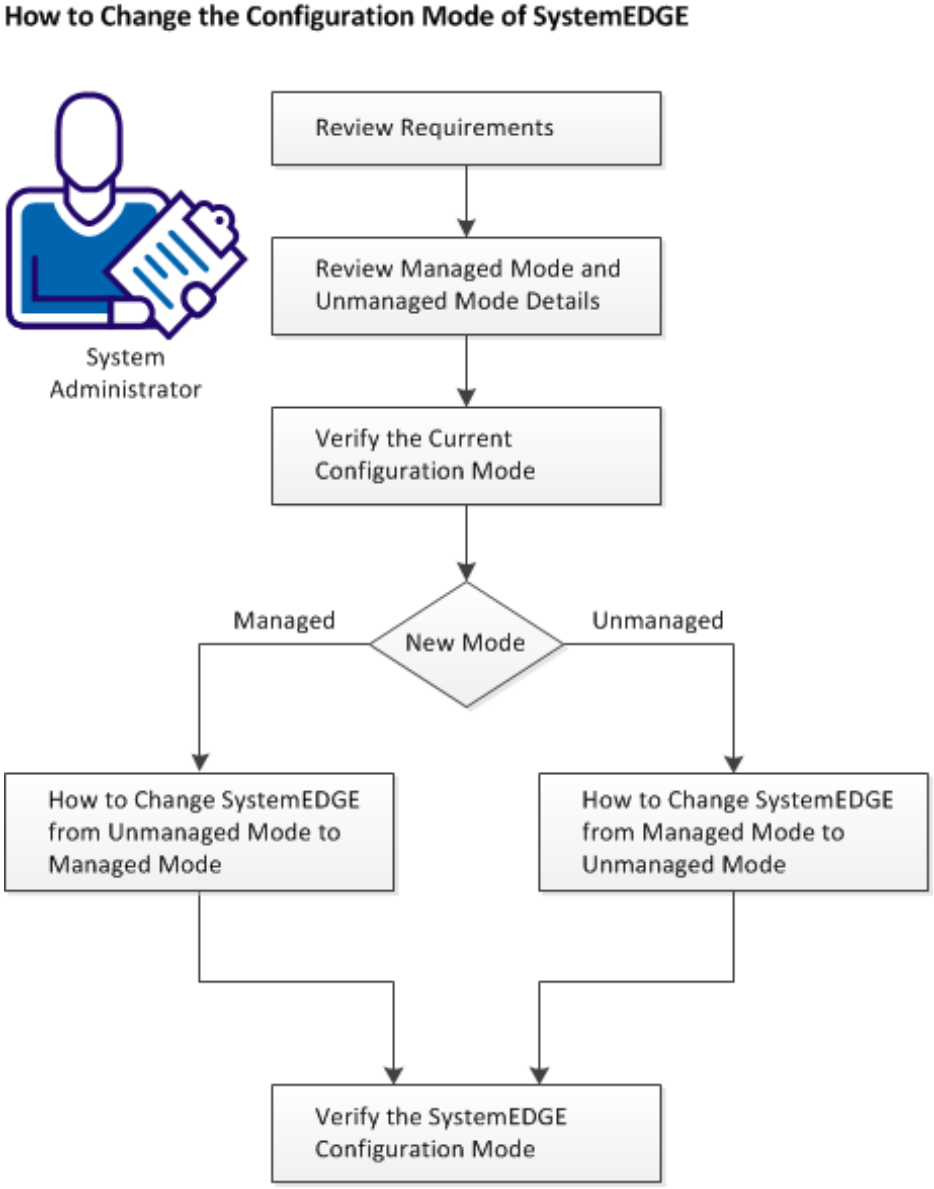

Follow these steps:

Review Requirements

Review the following requirements before you change the configuration mode for SystemEDGE.

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You have a basic understanding of CA SystemEDGE.
- You can access a CA Virtual Assurance manager installation that includes the Monitoring Agent (CA SystemEDGE).
- You can access the monitoring agents (CA SystemEDGE) on managed nodes.
- You can access the CA Virtual Assurance user interface.
- CA Virtual Assurance has discovered all relevant systems.

Review Managed Mode and Unmanaged Mode Details

Consider the following terminology: In this scenario, the terms unmanaged mode and managed mode are used in the context of SystemEDGE configuration.

Unmanaged mode

The SystemEDGE configuration on a particular server is not managed through CA Virtual Assurance Policy Configuration. You can edit the sysedge.cf file to modify the configuration.

Managed mode

The SystemEDGE configuration on a particular server is managed through CA Virtual Assurance Policy Configuration. You specify the SystemEDGE configuration in Policy Configuration on the CA Virtual Assurance manager and distribute it to the appropriate servers in the network. If you edit the sysedge.cf file locally, CA Virtual Assurance overwrites your changes with the next policy distribution.

Consider the following cases which have an influence on the configuration mode of SystemEDGE:

- If you run a typical SystemEDGE installation from the product media, SystemEDGE is configured to run in unmanaged mode after the installation.
- If you run a custom SystemEDGE installation from the product media, you can specify a manager system to use for the managed mode. If you have specified a manager system and CA Virtual Assurance discovers SystemEDGE after the installation, SystemEDGE registers with Policy configuration automatically and SystemEDGE runs in managed mode.
- If you use Remote Deployment to install SystemEDGE on remote systems, you can specify the configuration mode for SystemEDGE in the deployment job. The default value is managed mode.

Important! The Explore pane shows the Managed and Unmanaged folders which list the discovered servers that are polled (managed) or not polled (unmanaged) by CA Virtual Assurance. This property is different from the managed or unmanaged mode of the SystemEDGE configuration. The managed or unmanaged status of a server in the Explore pane has no influence on the configuration mode of SystemEDGE. Specific entries in the SystemEDGE configuration file indicate the configuration mode of SystemEDGE.

Verify the Current Configuration Mode of SystemEDGE

The following steps describe a method to determine the configuration mode of SystemEDGE.

The following terminology is used throughout these use cases:

Static sysedge.cf file

Identifies the file laid down by the installer, and is located in the *Installed_Dir*\SystemEDGE\config directory.

Default:

Windows: C:\Program Files\CA\SystemEDGE\config

UNIX/Linux: /opt/CA/SystemEDGE/config

Dynamic sysedge.cf file

Identifies the ongoing SystemEDGE configuration file, and is located under the *Data_Dir*\port<*number*> directory.

Default:

Windows: C:\Users\Public\CA\SystemEDGE\port161

UNIX/Linux: /opt/CA/SystemEDGE/config/port161

Follow these steps:

1. Log in the server on which SystemEDGE runs for which you want to determine the configuration mode.
2. Change to the 'data' directory of SystemEDGE and open the port<*number*> directory. On Windows, you can open the sysedge.cf file in the data directory from the SystemEDGE Control Panel.

   Note: The dynamic sysedge.cf file in the 'data' directory is different from the static sysedge.cf file in the 'config' directory.
3. Open the dynamic sysedge.cf file in the port<*number*> directory.

   If SystemEDGE runs in managed mode, the first line specifies a control value (ctrl_value).

   Example:

   ```
   ctrl_value 0x9e30d00e

   # Generated file - DO NOT EDIT
   #
   # Configuration file generated on 2012:03:27 05:37
   #
   # Generated from default.generic.0.prof
   #
   version 5.7
   ```

   If SystemEDGE runs in unmanaged mode, the first line specifies the version.

   Example:

   ```
   version 5.7
   ```

How to Change SystemEDGE from Managed Mode to Unmanaged Mode

The following diagram provides an overview of the required actions to change to unmanaged mode.

How to Change SystemEDGE from Managed Mode to Unmanaged Mode

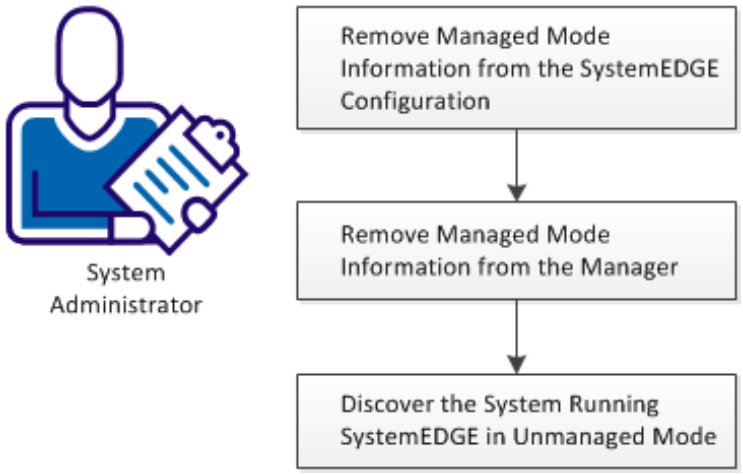

Follow these steps:

Remove Managed Mode Information from the SystemEDGE Configuration (see page 272)
Remove Managed Mode Information from the Manager (see page 273)
Discover the System Running SystemEDGE in Unmanaged Mode (see page 274)

Remove Managed Mode Information from the SystemEDGE Configuration

The following procedure describes how to remove managed mode information from SystemEDGE configuration on a particular server.

Follow these steps:

1. Log in the server on which you want to change the configuration mode for SystemEDGE.
2. Create the following backup directories at a convenient location:

   ```
   data.backup
   config.backup
   ```

3. Stop SystemEDGE, using the normal mechanism.
4. Navigate to the 'data' directory of SystemEDGE and open the port<*number*> directory. The default directory is port161.

   The content of the directory is listed.
5. Move the following files to your data.backup directory so that they no longer appear in the port<*number*> directory:

   ```
   .sysedge.id
   sysedge.cf
   ```
6. Change to the 'config' directory of SystemEDGE.

   The content of the directory is listed.
7. Copy the following file to your config.backup directory:

   ```
   sysedge.cf
   ```
8. Navigate to the 'config' directory, open the sysedge.cf file in a text editor, and scroll down to the bottom of the file.
9. Delete the following line:

   ```
   manager_name <hostname of the manager>
   ```
10. Save the file and start SystemEDGE.

    SystemEDGE creates a sysedge.cf file in the 'data' directory without any managed mode information.

Remove Managed Mode Information from the Manager

The following procedure describes how to remove managed mode information from SystemEDGE configuration on a particular server.

Follow these steps:

1. Log in the CA Virtual Assurance user interface and change to Management.

   The Resources tab opens and shows the Explore pane.
2. Enter the name of the server on which you have modified the SystemEDGE configuration into the Search field, and click (Search).

   The Search window opens and lists the search results.
3. Click one of the search results.

   The resources page for that particular server opens and shows the Quick Start panel.
4. Click Delete from System.

   The server disappears from the Explore pane. All server-related objects are deleted on the manager including managed mode information.

Discover the System Running SystemEDGE in Unmanaged Mode

The following procedure describes how to rediscover the server that runs in unmanaged mode.

Follow these steps:

1. Log in the CA Virtual Assurance user interface and change to Management.

   The Resources tab opens and shows the Explore pane.

2. Right-click Data Center, select Management, Discover, Server.

   The Discover Window opens.

3. Enter the name of the server that you have deleted in the previous procedure and click Finish.

   CA Virtual Assurance discovers the server on which SystemEDGE runs in unmanaged mode.

   After CA Virtual Assurance has completed Discovery, SystemEDGE on the discovered server has not been registered in Policy Configuration. SystemEDGE runs in unmanaged mode.

4. Double-click the server name in the Explore pane.

   The resources page for that server opens.

5. Change to the Summary tab to verify that CA Virtual Assurance has correctly discovered the server. If necessary, you can select a different SNMP community which CA Virtual Assurance uses for monitoring the server.

   You can now configure SystemEDGE on that server through editing the sysedge.cf configuration file.

How to Change SystemEDGE from Unmanaged Mode to Managed Mode

The following diagram provides an overview of the required actions to change to managed mode.

How to Change SystemEDGE from Unmanaged Mode to Managed Mode

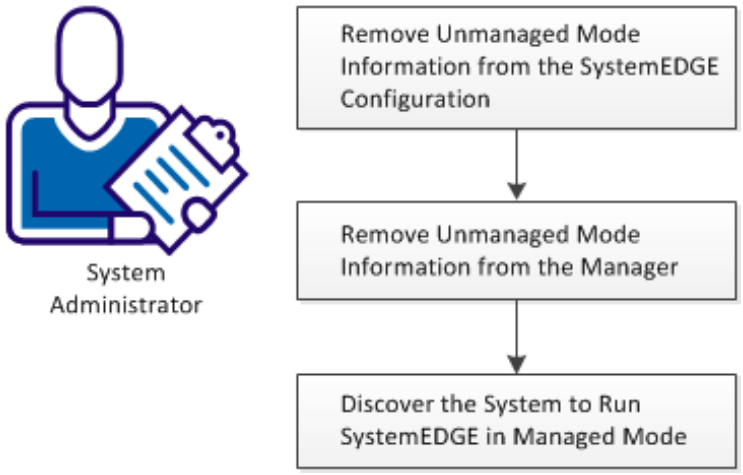

Follow these steps:

Remove Unmanaged Mode Information from the SystemEDGE Configuration (see page 275)
Remove Unmanaged Mode Information from the Manager (see page 276)
Discover the System to Run SystemEDGE in Managed Mode (see page 277)

Remove Unmanaged Mode Information from the SystemEDGE Configuration

The following procedure describes how to remove unmanaged mode information from SystemEDGE configuration on a particular server and to prepare the server to change to managed mode.

Follow these steps:

1. Log in the server on which you want to change the configuration mode for SystemEDGE.
2. Create the following backup directories at a convenient location:

   ```
   data.backup
   config.backup
   ```

3. Stop SystemEDGE, using the normal mechanism.

4. Navigate to the 'data' directory of SystemEDGE and open the port<*number*> directory. The default directory is port161.

   The content of the directory is listed.

5. Move the following file to your data.backup directory so that it no longer appears in the port<*number*> directory:

   `sysedge.cf`

6. Change to the 'config' directory of SystemEDGE.

   The content of the directory is listed.

7. Copy the following file to your config.backup directory:

   `sysedge.cf`

8. Navigate to the 'config' directory, open the sysedge.cf file in a text editor, and scroll down to the bottom of the file.

9. Add the following line:

   `manager_name <hostname of the manager>`

10. Save the file and start SystemEDGE.

    SystemEDGE creates a sysedge.cf file in the 'data' directory.

Remove Unmanaged Mode Information from the Manager

The following procedure describes how to remove unmanaged mode information from SystemEDGE configuration on a particular server.

Follow these steps:

1. Log in the CA Virtual Assurance user interface and change to Management.

   The Resources tab opens and shows the Explore pane.

2. Enter the name of the server on which you have modified the SystemEDGE configuration into the Search field, and click (Search).

   The Search window opens and lists the search results.

3. Click one of the search results.

   The resources page for that particular server opens and shows the Quick Start panel.

4. Click Delete from System.

   The server disappears from the Explore pane. All server-related objects are deleted on the manager.

Discover the System to Run SystemEDGE in Managed Mode

The following procedure describes how to rediscover the server that causes SystemEDGE to run in managed mode.

Follow these steps:

1. Log in the CA Virtual Assurance user interface and change to Management.

   The Resources tab opens and shows the Explore pane.

2. Right-click Data Center, select Management, Discover, Server.

   The Discover Window opens.

3. Enter the name of the server that you have deleted in the previous procedure and click Finish.

   CA Virtual Assurance discovers the server.

   After CA Virtual Assurance has completed Discovery, SystemEDGE on the discovered server has been registered in Policy Configuration. SystemEDGE runs in managed mode.

4. Double-click the server name in the Explore pane.

   The resources page for that server opens.

5. Change to the Summary tab to verify that CA Virtual Assurance has correctly discovered the server. If necessary, you can select a different SNMP community which CA Virtual Assurance uses for monitoring the server.

Verify the SystemEDGE Configuration Mode

Basically, you can repeat the procedure in Verify the Current Configuration Mode (see page 270).

If SystemEDGE runs in unmanaged mode, the first line of the dynamic sysedge.cf file in the 'data' directory specifies the version.

Example

```
release 5.7.1
```

If SystemEDGE runs in managed mode, the first line specifies a control value (ctrl_value).

Example

```
ctrl_value 0x9e30d00e

# Generated file - DO NOT EDIT
#
# Configuration file generated on 2012:03:27 05:37
#
# Generated from default.generic.0.prof
#
release 5.7.1
```

During the discovery process, additional meta information has been added to the end of the dynamic sysedge.cf file.

Example

```
template data_directory <path>
data_directory "C:\Users\Public\CA\SystemEDGE\"
template default_port CA Portal
default_port 161
template manager_name <name>
manager_name manager_server.mycompany.com
template manager_policy_name <policy>
manager_policy_name default.generic
template manager_policy_version <version>
manager_policy_version 1
```

More information:

Verify the Current Configuration Mode of SystemEDGE (see page 270)

Chapter 6: Managing Virtual Environments

This section contains the following topics:

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:

How to Configure the Cisco UCS Management Components

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.

Review Requirements

Review the following requirements before configuring the management components of CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Virtual Assurance and SystemEDGE.
  - You can access a CA Virtual Assurance manager installation that includes:
  - Platform Management Module (PMM)
  - Application Insight Module (AIM)
  - Monitoring Agent (SystemEDGE)
- You can access the CA Virtual Assurance user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote AIM Servers that you want to use.

Cisco UCS Server

Verify the following conditions for Cisco UCS management:

Launch the Cisco Java user interface to verify that the Cisco UCS Manager is running. The link to launch the Cisco Java user interface is http://<UCS_Manager_name> or https://<UCS_Manager_name>.

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 the UCS Manager service profiles information.

Interaction Between Cisco UCS Management Components

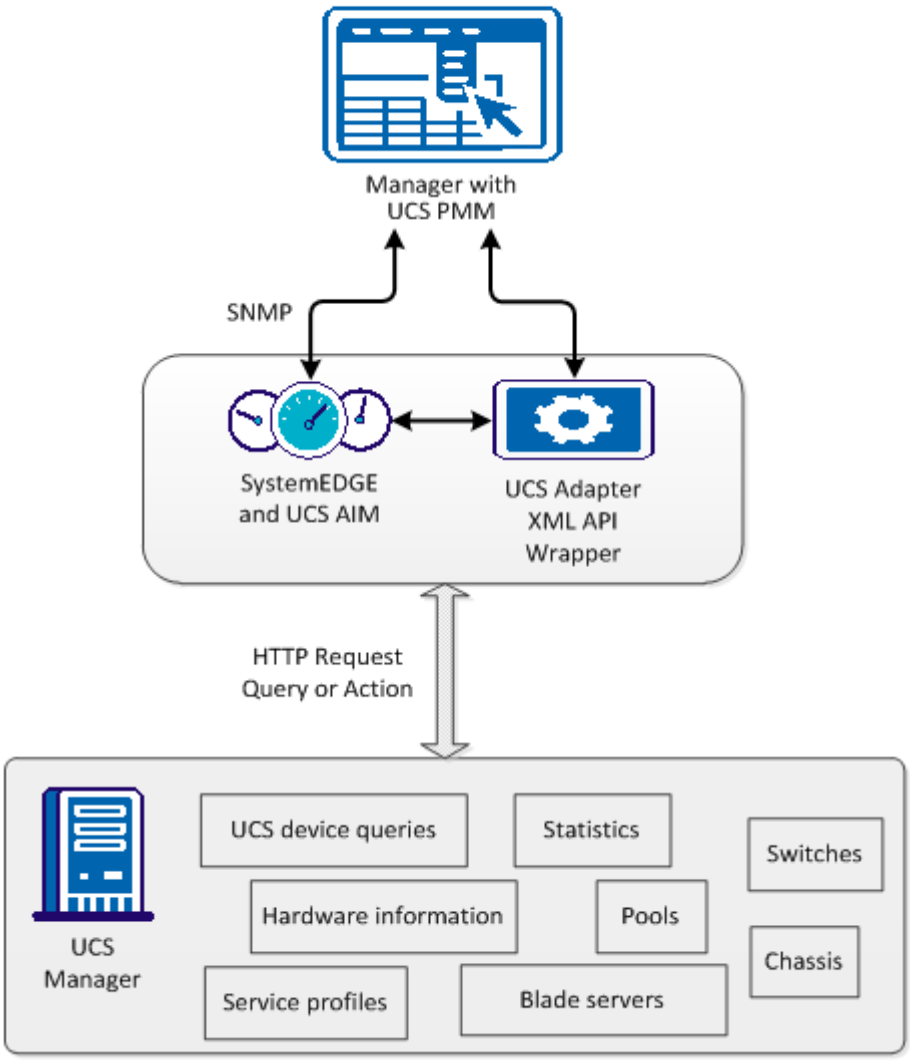

The diagram shows the integration components for the Cisco UCS. The communication protocol between the UCS Adapter and Cisco UCS manger is HTTP or HTTPS.

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.

Add a Cisco UCS to the Manager

You can add a Cisco UCS Manager server using the Administration page of the 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 Cisco UCS from the Provisioning section in the left pane.

3. Click (Add) on the Cisco UCS pane toolbar.

   The Add Cisco Unified Computing System Server dialog appears.

4. Enter the required connection data (server name, user, password, port), specify the preferred AIM, enable Managed Status (checkbox).

5. Enter the required server identification information, and click OK.

   If the network connection is established successfully, the Server is added to the top right pane with a green status icon.

   Note: If the connection fails, the Validation Failed dialog appears. Click Yes, CA Virtual Assurance adds the Server to the list with a red status icon. If you click No, nothing is added.

Manager Connection to the Server Fails

Symptom:

After I have added a server connection under Administration, Configuration, the validation of the connection to the server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used server connection data is still valid. If necessary, update the connection data.
- Verify, if the server system is running and accessible.
- Verify, if all services that are required for the connection are running properly on the server system.

To update the server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 server cannot be established, continue with the next procedure.

To verify if the server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. To find out whether the server has a valid DNS entry and IP address, verify the output of these commands.

   If the server is not in the DNS, add the server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>

   Enter the correct IP address and server name and save the file. For example:

   `192.168.50.50 myServer`

4. Change to the CA Virtual Assurance user interface, Administration tab, Configuration, Server pane, and click (Validate) in the upper-right corner.

   If the server credentials and connection data are correct and you can ping the server, the connection can still fail. In this case, it is possible that the server causes the problem. If the connection to the server cannot be established, continue with the next procedure.

To verify, if all services that are required for the connection are running properly on the server system:

1. To access the server, contact the system administrator.
2. Log in to the server system.
3. Verify, if all services that are required for the connection are running properly.
4. If necessary, start or restart the service.
5. Change to the CA Virtual Assurance user interface, server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the server connection.

   If the connection to the server fails, verify the validity of the data you gathered according to the requirements for this scenario.

   Work with the administrator or support to fix the server connection problem.

Register a UCS AIM Server

After adding a Cisco UCS component to the CA Virtual Assurance manager, add the AIM instance using the Administration page of the user interface to manage the Cisco UCS environment.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.

2. Select Cisco UCS from the Provisioning section in the left pane.

3. Click  (Add) on the UCS AIM Servers pane toolbar.

   The Add Cisco Unified Computing System AIM Server dialog appears.

4. Select the UCS AIM Server from the drop-down list.

   The list of UCS AIM Servers appears.

5. Select the Cisco UCS Server from the drop-down list.

   CA Virtual Assurance populates the Cisco UCS Server drop-down list with the servers listed in the Cisco UCS pane. You can only manage those UCS Servers for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first. After the Discovery, the AIM server appears in the drop-down list.

6. Click OK.

   A new AIM instance for the selected Server is registered.

   Note: If the instance is not in an error or in a stopped state, CA Virtual Assurance starts to discover the associated environment. When the Discovery process is complete, you can start managing the Cisco UCS environment.

Verify the Cisco UCS in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand Cisco UCS group.

   The managed Cisco UCS resources appear.

CA Virtual Assurance is now ready to manage the configured Cisco UCS environment.

Cisco UCS Management

The CA Virtual Assurance integration with Cisco UCS lets you manage your UCS switches, chassis, and blades from a centralized user interface. The UCS Manager running on a switch is the location from which you can view UCS resources and perform management operations such as cloning, snapshots, power operations, and so on.

This section describes the resource management operations that you can perform on Cisco UCS resources from the Resources page. The Resources page displays basic information about the following UCS objects:

- Cisco UCS servers
- UCS Manager servers
- Chassis
- Blade servers
- Fabric interconnects
- Organizations

The Summary page lets you view information associated with that object (for example, a chassis summary can display its blades, or a blade summary can display its storage) and events associated with the resource.

If available, a Details page lets you view other resource information, such as system properties, software, hardware, performance, and so on. After you assign an automation source here, default access and management profiles are automatically created for the system and discovery is run.

Other pages can be available to perform resource management tasks. The right-click menus on Explore pane objects also let you perform UCS management tasks.

More information:

How To Use Centralized Service Profiles

Central service profiles that reside in the CA Virtual Assurance Management DB provide an efficient way to manage configuration information across multiple UCS domains. Use the CA Virtual Assurance user interface to import service profiles into the Management DB from UCS Managers, or create a central service profile in the Management DB.

From the Management DB, you can export central service profiles to any UCS Manager.

Manage Central Service Profiles

You can manage central service profiles from the Resources page. For access, select Cisco UCS Server in the Explore pane, and click Central Service Profiles in the right-hand pane.

To import service profiles from UCS Managers

1. Click the white triangle (Import) icon.
2. Use the Import Service Profiles dialog to select UCS Managers. Click Refresh to populate the Service Profiles list and select Import All Service Profiles, or select one or more from the list. To remove imported profiles from their UCS managers after import, select Delete Source Service Profiles.
3. Click OK.

   The selected service profiles are imported into the Management DB.

To create or update a central service profile in the Management DB

1. Click the + (Create) icon, or select a central service profile and click the tool (Edit) icon.
2. Use the wizard pages to create or update the central service profile.

   Note: You cannot specify pools and policies when you create a service profile in the Management DB; this information is for reference only. You can specify this information after you export the central service profile to the UCS Manager.
3. When the service profile is created or updated, click Finish.

To export service profiles to a UCS Manager

1. Select one or more central service profiles
2. Click the blue triangle (Export) icon.

   The Export Service Profiles dialog appears.
3. On the Available UCS Managers list, select one UCS Manager and click a right arrow to transfer to the Selected UCS Managers list. Click the double right arrows to transfer all.
4. Click OK.

   Note: Pools and policies are not exported; they must already reside on the target UCS Manager.

To delete service profiles from the Management DB

1. Select the service profiles to delete.
2. Click the - (Delete) icon.

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.

More information:

Configure Service Profile: Cisco UCS (see page 638)

How to Manage Port Profiles

Use the following process to manage port profiles using CA Virtual Assurance.

1. Export plug-ins.

   To establish the communication between Cisco UCS Manager and vCenter, generate and install one or more extension XML files in the target vCenter as follows:

   - For vCenter 4.0, multiple extension files are required.
   - For vCenter 4.0 update 1 version and above, export a single extension file from Cisco UCS Manager.

   After the required files are exported, import them into vCenter as new plug-ins using vSphere Client. This is a one-time requirement for each UCS Manager and vCenter combination; a Cisco UCS Manager cannot use files exported from a different UCS Manager.

2. Export .vib file to ESX server.

   Based on the ESX version, configure the ESX server by installing the appropriate .vib file component from the Cisco Nexus 1000V Virtual Ethernet Module Software. This package (jointly designed by Cisco and VMware) enables a distributed virtual switch solution that is fully integrated with the VMware Virtual Infrastructure.

3. Create Port Profile Network Topology (see page 293)
4. Create Port Profiles and Port Profile Clients (see page 293)

Create Port Profile Network Topology

To push a port profile to VMware, the Cisco UCS Manager must have defined vCenter, datacenter, DVS folder, DVS, and profile client objects. The topology of these objects must match the topology in VMware. CA Virtual Assurance lets you create the required topology.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Right-click a UCS Manager in the Explore tree, and click VMware to launch the vCenter Layout dialog
3. Expand vCenter and highlight a vCenter, datacenter, DVS folder, DVS, or profile client object.
4. Select Create New in the Actions drop-down menu.
5. Enter the required information, and click Finish. Selecting Enable on the DVS panel automatically pushes the associated port profiles to vCenter.

   The vCenter, datacenter, DVS folder, DVS, or profile client object is added.

Note: To use this dialog to delete topology objects, select Delete on the Actions drop-down menu.

Create Port Profiles and Port Profile Clients

You can use the CA Virtual Assurance vCenter Layout dialog to create port profiles and port profile clients.

Follow these steps:

1. Right-click a UCS Manager in the Explore tree and click VMware to launch the vCenter Layout dialog.
2. Highlight Port Profiles to create a port profile or an existing port profile to create a port profile client.
3. Select Create New in the Actions drop-down menu.
4. Enter the required information, and click OK.

   The port profile or port profile client is created.

Note: You can also use this procedure to edit or delete port profiles and port profile clients. Highlight an existing port profile or port profile client, and select Edit or Delete on the Actions drop-down menu.

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.

View Cisco UCS Resources

The Resources page lets you view UCS resources at any level on the UCS object tree. For example, you can inspect the following objects to determine:

- Cisco UCS Servers - UCS resources by category, blade allocation, and imported service profiles
- Centralized service profiles - UCS Manager assignments, import, and export
- UCS Managers - Fabric interconnects, chassis, and organizations
- Chassis - Number of blades mounted, number of fans (and their status), and input/output modules
- Blade servers - Number of blades, whether powered on or off, whether associated with service profiles, and the OS host name

  Note: Expand the blade in the Explore tree to see the OS host. The OS host name for a blade will be available after provisioning and discovery are complete.
- Individual blade - High-level inventory, including motherboard, CPU, memory, and storage
- Fabric interconnects - High-level hardware and fans
- Organizations - Pools, service profiles, and service profile templates

To view a Cisco UCS resource

1. Click Resources.

   The Resources page appears.
2. Open the Explore pane.

   Available groups, services, and systems appear.
3. Find and select a Cisco UCS resource.

   The resource page appears in the right pane.

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.

More information:

Configure Service Profile: Cisco UCS (see page 638)

Back Up a UCS Manager Configuration

CA Virtual Assurance supports backup of the following types of UCS Manager configuration information:

- Full state
- All configurations
- System configuration
- Logical configuration

The export/import capability emulates the Cisco UCS capability and allows you to create and rerun backup jobs.

To import/ export a UCS Manager configuration

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Select a Cisco UCS Manager.

   The UCS Manager page appears in the right pane.

4. Click Export/Import.

   The Export/Import page appears.

5. Click + (Create) in the Import job or Export jobs section as required.

   The Create Backup Operation dialog appears.

6. Enter import information or export information as required and click OK.

   The job starts and displays in the respective list.

vNIC Templates

CA Virtual Assurance supports the creation and management of vNIC templates. You can specify the template target as either the service profile or a VM.

To create a vNIC template, select Use vNIC Template in the Service Profile wizard, and launch the Create vNIC Template dialog. You also can right-click on a UCS organization in the Explore pane.

UCS Organizations

You can use organizations to group related UCS resources to create a nested hierarchy of pools, service profiles, and service profile templates for UCS resource management. Organizations and sub-organizations can be created and deleted.

Create a Sub Organization

You can create organizations or sub-organizations on the UCS root tree.

To add an organization

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

   Navigate to root, or open root and click a sub-organization

3. Right-click root or a sub-organization to select Management, Create Sub Organization.

4. Use the Create Sub Organization dialog to create the new sub-organization.

   The sub-organization is created.

UCS Pools

CA Virtual Assurance lets you create pools to manage UCS resources more efficiently.

Note: When pool range conflicts occur, a warning displays.

The following types of pools are available:

- UUID Pools
- MAC Pools
- World Wide Node Name (WWNN) Pools
- World Wide Port Name (WWPN) Pools
- Server Pools

WWNN and WWPN pools can be used for configuring blade to use remote storage (SAN).

More information:

View a UCS Pool (see page 300)
Create a UCS Pool (see page 300)
Rename a UCS Pool (see page 301)
Delete a UCS Pool (see page 302)

View a UCS Pool

The Resources page lets you view a UCS pool.

To view a UCS pool

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Click root at the top of the UCS organization tree, and navigate to the desired organization.

4. Click Summary.

   The Summary page appears.

5. In the Components section, click the desired Pools type in the drop-down menu.

6. Select a pool to view, and click the tool (View) icon.

7. Click Cancel to return to the Pools list.

Create a UCS Pool

The Resources page lets you create UCS pools to manage UCS resources more efficiently.

To create a resource pool

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Click root at the top of the UCS organization tree, and navigate to the desired organization.

4. Click Summary.

   The Summary page appears.

5. In the Components section, click the desired Pools type in the drop-down menu.

6. Click + (Add new pool).

   The Create Pool dialog appears.

7. Use the dialog to complete the definition.

   The pool is added to the pool list.

Note: You can customize the Quick Start menu to provide the function.

Rename a UCS Pool

The Resources page lets you rename a UCS pool.

To rename a UCS pool

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Click root at the top of the UCS organization tree, and navigate to the desired organization.

4. Click Summary.

   The Summary page appears.

5. In the Components section, click the desired Pools type in the drop-down menu.

6. Select the pool to rename.

7. Click the double-arrow icon (>>).

   The Rename Pool dialog appears.

8. Use the dialog to rename the pool.

   The pool is renamed in the list.

Delete a UCS Pool

The Resources page lets you delete a UCS pool.

To delete a UCS pool

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Click root at the top of the UCS organization tree, and navigate to the desired organization.

4. Click Summary.

   The Summary page appears.

5. In the Components section, click the desired Pools type in the drop-down menu.

6. Select the pool to delete.

7. Click the - (Delete) icon.

   The Delete confirmation appears.

8. Click Yes.

   The pool is deleted.

UCS Action Types

Cisco UCS resources can use CA Virtual Assurance action types to create new actions that automate UCS power, resource allocation, and other operations when the assigned rule criteria are met. You can also schedule these actions to occur at specific times.

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.

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.

More information:

Configure Power: Cisco UCS (see page 629)

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.

   The Cisco UCS 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.

Remove a UCS AIM

You can remove a UCS AIM 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.

   The Cisco UCS page appears.

3. Select the UCS AIM server that you want to remove.

4. Click - (Delete) on the server toolbar.

   A confirmation prompt appears.

5. Click OK.

   The UCS AIM server is removed.

Citrix XenServer

Citrix XenServer is a virtualization platform that offers near bare metal virtualization performance for virtualized server and client operating systems. XenServer uses the Xen hypervisor to virtualize each server on which it is installed, enabling each server to host multiple virtual machines (VMs) simultaneously with guaranteed performance. XenServer provides its own operating system to administer the physical and virtual resources of a XenServer host, and therefore, it does not require a specific operating system. XenServer supports Linux and Windows guest operating systems.

XenServer resources can be managed at three levels:

Host Management

A *XenServer host* object represents a physical host on which XenServer and its VMs run. A XenServer host can be a stand-alone host or associated with a XenServer pool. You can monitor virtual and physical resources available on a XenServer host, manage storage repositories containing virtual disk images, manage tasks, or run the XenServer host in maintenance mode.

Resource Pool Management

A *resource pool* is a connected group of up to 16 XenServer hosts. Combined with shared storage and dynamically controlled memory, CPU, and networking resources, the XenServer hosts in a resource pool provide an operating environment on which VMs run. You can manage the membership or role of the XenServer hosts in a pool, and can let XenServer monitor the health of the pool members for high availability. If necessary, VMs are live migrated between pool hosts to avoid downtime.

Virtual Machine Management

On the VM management level, you can perform the following tasks:

- Control VMs (Discover, Start, Suspend, Shutdown, Delete From Disk)
- Manage VMs (clone)

How to Configure XenServer Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the Management Components

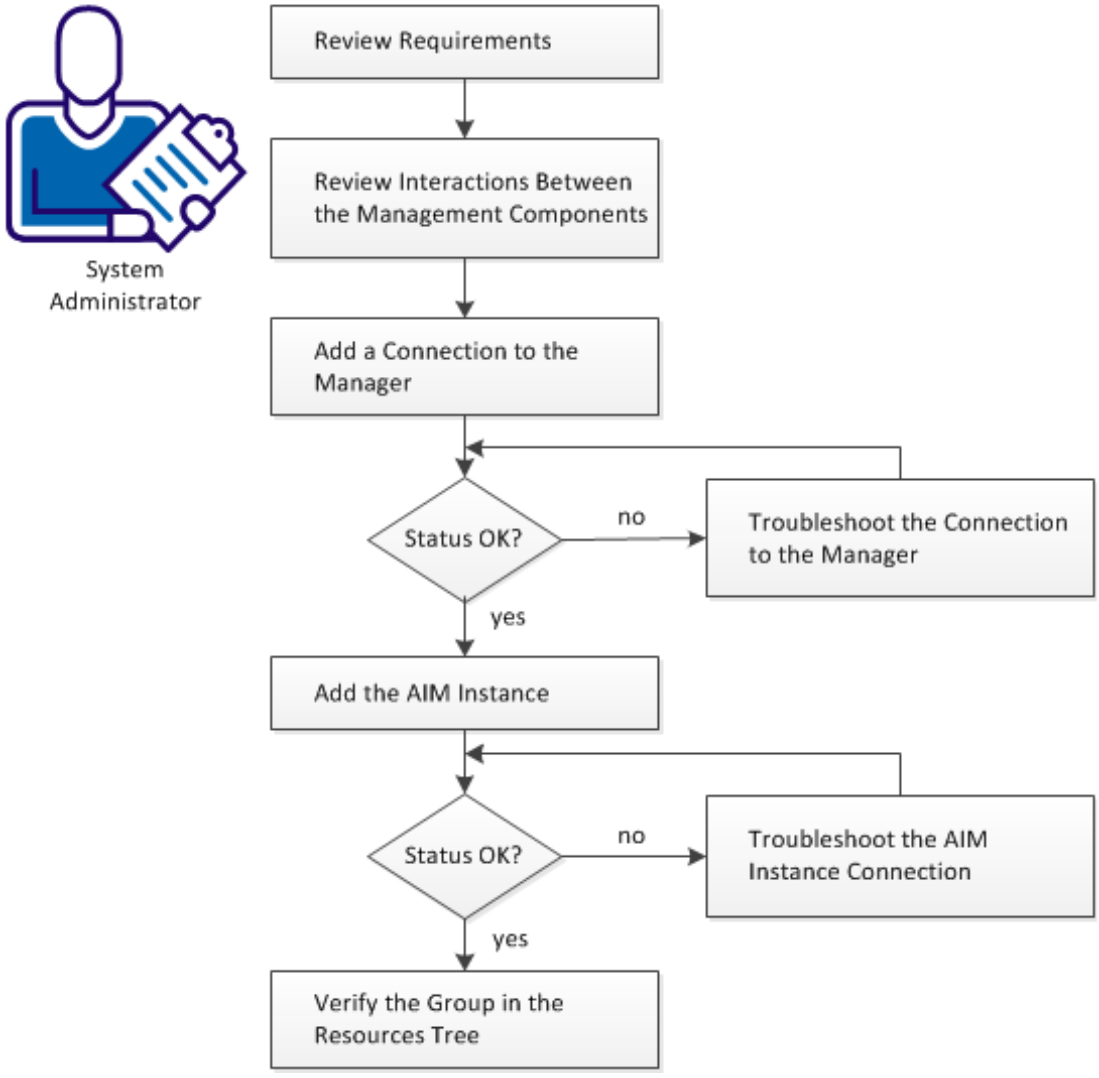

Follow these steps:

Review Requirements

Review the following requirements before configuring the management components of CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Virtual Assurance and SystemEDGE.
  - You can access a CA Virtual Assurance manager installation that includes:
  - Platform Management Module (PMM)
  - Application Insight Module (AIM)
  - Monitoring Agent (SystemEDGE)
- You can access the CA Virtual Assurance user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote AIM Servers that you want to use.

Interactions Between Citrix XenServer Management Components

The Citrix XenServer AIM is implemented as a multi-instance, remote AIM. CA Citrix XenServer AIM is able to monitor multiple standalone Citrix XenServers and Citrix XenServer resource pools remotely. The Citrix XenServer AIM is implemented as x86 and x64 module.

The management API for Citrix XenServer is based on XML RPC. For a Citrix XenServer resource pool, all XML RPC communication is taking place between the AIM, PMM, and the pool master only.

Interaction Between XenServer Management Components

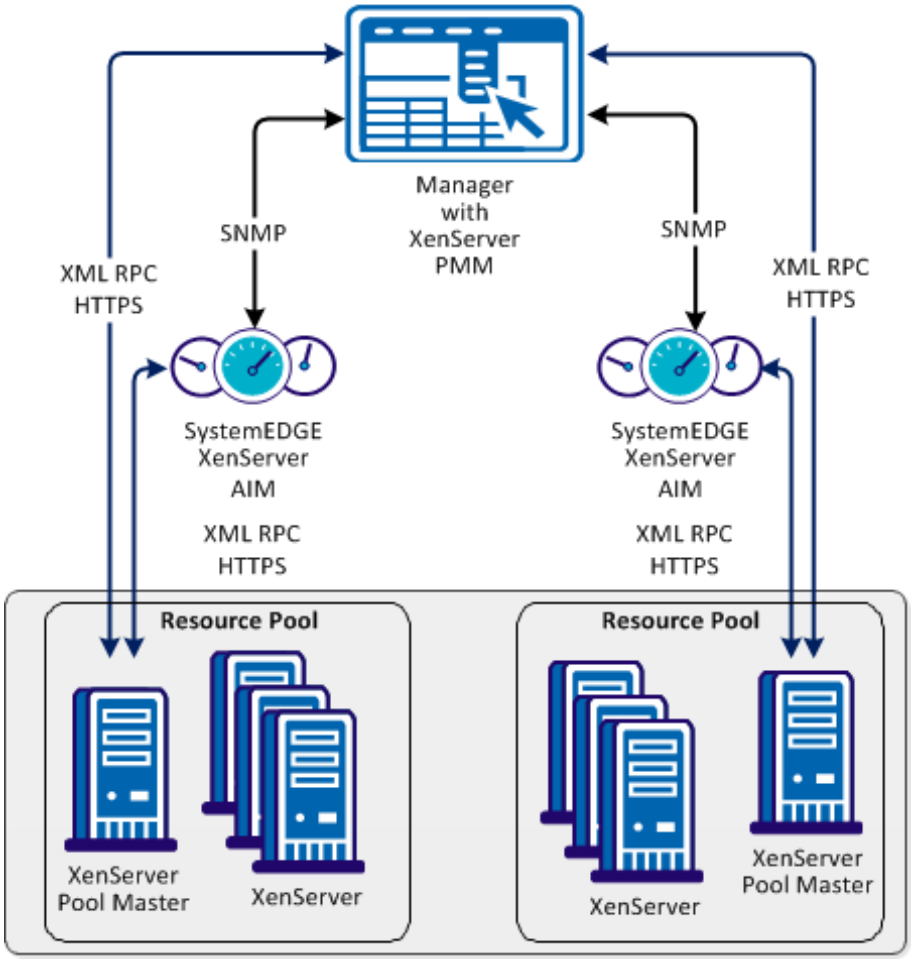

Add a Citrix XenServer Connection to the Manager

You can add a Citrix XenServer 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 Citrix XenServer from the Provisioning section in the left pane.

3. Click (Add) on the Registered Citrix XenServers pane toolbar.

   The Add Citrix XenServer dialog appears.

4. Enter the required connection data (server name, username, password, resource pool UUID), specify the preferred AIM, and enable Managed Status (checkbox).

   Important! Verify that you add the pool master to the Registered Citrix XenServers.

5. Click OK.

   If the network connection has been established successfully, the Server is added to the top right pane with a green status icon. CA Virtual Assurance discovers the Citrix XenServer system automatically.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Server Connection to the Manager Failed (Citrix XenServer)

Symptom:

After I have added a server connection under Administration, Configuration, the validation of the connection to the server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used server connection data is still valid. If necessary, update the connection data.
- Verify, if the server system is running and accessible.
- Verify, if the Management Service on the server system is running properly.

To update the Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 server cannot be established, continue with the next procedure.

To verify if the Server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. Verify the output of the commands to find out whether the server has a valid DNS entry and IP address.

   If the server is not in the DNS, add the server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>`

   Enter the correct IP address and Server name. For example:

   `192.168.50.50 myServer`
4. Click  (Validate) in the upper-right corner.

   If the server credentials and connection data are correct and you can ping the server, the connection can still fail. In this case, it is possible that the server causes the problem. If the connection to the server cannot be established, continue with the next procedure.

To verify, if the Management Service on the Server system is running properly:

1. Contact the Administrator to access the server system.
2. Log in to the server system and execute xsconsole command.

   The Service control console launches.
3. Verify the status of the service and resolve any reported issues.
4. Change to the CA Virtual Assurance user interface, server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the server connection.

   If the connection to the server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

   Work with the administrator or support to fix the server connection problem.

Add the Discovered Citrix XenServer AIM Instance

After adding a Citrix XenServer connection to the CA Virtual Assurance manager, add an AIM instance to manage the Citrix XenServer.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.
2. Select Citrix XenServer from the Provisioning section in the left pane.

3. Click  (Add) on the Discovered Citrix XenServer AIM Instances pane toolbar.

   The Add Citrix XenServer AIM dialog appears.

4. Select the Citrix XenServer AIM Server from the drop-down list.

   The list of discovered XenServer AIM Servers appears. If you have installed the XenServer AIM on the local system, the name of the local system appears in the list too.

5. Select the Citrix XenServer from the drop-down list.

   CA Virtual Assurance populates the XenServer drop-down list with the XenServers listed in the Registered Citrix XenServers pane. You can only manage those XenServers for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

   A new AIM instance for the selected Server is added. If the instance is not in an error or in a stopped state, CA Virtual Assurance starts to discover the associated environment. When the discovery process is complete, you can start managing your Citrix XenServer environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the number of managed objects that are related to virtual and physical resources in your environment. You can move the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Virtual Assurance adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance 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 AIM:

- Verify that the AIM Server is accessible.
- Verify that SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping servername`

2. Verify that the output of the commands has a valid DNS entry and IP address for the AIM server.

   If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 servername`

   Enter the correct IP address and AIM server name. For example:

   `192.168.50.51 myAIM`

4. Click (Validate) in the upper-right corner of the 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 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, AIM Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the AIM Server connection.

   If the error status remains unchanged, verify that the data you gathered is according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Virtual Assurance discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Virtual Assurance manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Citrix XenServer Group in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand Citrix XenServer group.

   The managed Citrix Resource Pools appear.
3. Expand the Resource Pool entry.

   The managed Citrix XenServers appear.

CA Virtual Assurance is now ready to manage the Citrix XenServer environment that was configured.

How to Prepare Linux template for XenServer Provisioning

CA Virtual Assurance supports customized provisioning of new virtual machines (VM) running the following operating systems:

- Red Hat Enterprise Server 6.0
- SUSE Linux Enterprise Server 11

Customization options include hostname, password, domain, or network configuration.

The following diagram illustrates how a system administrator prepares Linux template for VM provisioning.

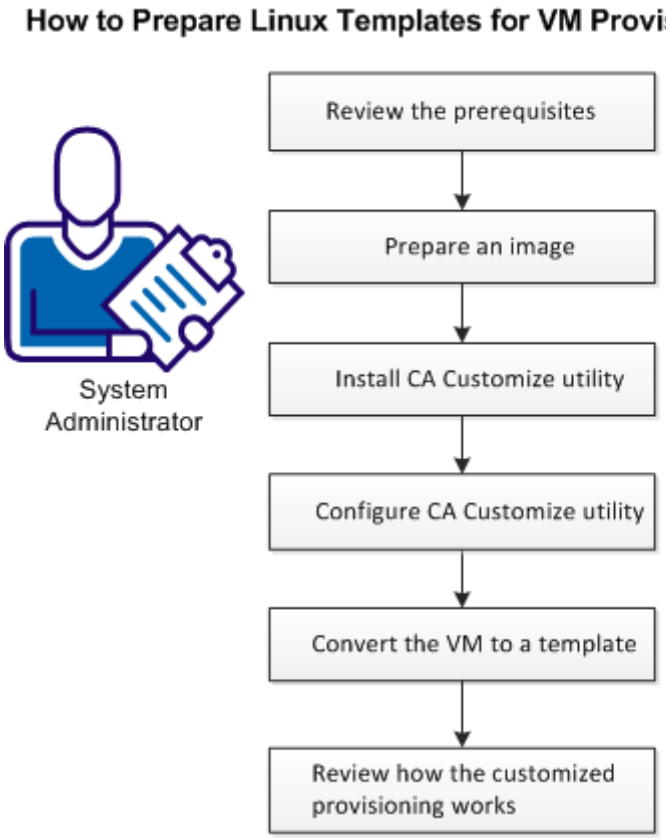

Follow these steps:

Prerequisites for Customized VM Provisioning (see page 316)
Prepare a Linux Image (XenServer) (see page 317)
Install CA Customize Utility (see page 317)
Configure CA Customize Utility (see page 318)
Convert the VM to a Template (see page 318)
How the Customized Provisioning Works (see page 319)
Customization Log (see page 319)

Prerequisites for Customized VM Provisioning

To customize the Linux guest, one needs direct access to the file system or console.

Ensure that the following prerequisites are met for the XenServer environment:

- Each XenServer in the resource pool must have SSH or SFTP access enabled.

Prepare a Linux Image (XenServer)

Before you create a template containing the Linux operating system, prepare the image by following this procedure. The specific steps may differ based on the Linux Distribution.

Follow these steps:

1. Install the Linux operating system on a new virtual machine from scratch.
2. Install the XenTools for Citrix XenServer on the virtual machine.
3. Apply any customizations like user accounts, policy, applications, hotfixes, that you would like to apply on the new virtual machines.

This Linux image is ready for further customization using the CA Customize utility.

Install CA Customize Utility

CA Customize utility enables CA Virtual Assurance to change the virtual machine settings externally. The guest utility monitors CD drive on the OS start. If a special ISO is connected, the following actions are executed:

1. A set of commands customizes the guest.
2. The guest system is marked as customized.
   The system cannot be modified again until someone resets this state.
3. The system is halted to indicate that the customization succeeded.

To install correct ca-customize guest utility:

1. Find this utility at:
   - Valid for Red Hat Enterprise Server 6.0

     `<InstallationRoot>\Utilities\linuxCustomization\rh6`
   - Valid for SUSE Linux Enterprise Server 11

     `<InstallationRoot>\Utilities\linuxCustomization\sles11`
2. Transfer this executable file to the following location on a hard drive of the VM being prepared:

   `/usr/bin/ca-customize`
3. (Optional) Provide your own version of ca-customize script to support other guest systems that we do not support.
4. Enable executable bit of the ca-customize utility:

   `chmod 755 /usr/bin/ca-customize`

Configure CA Customize Utility

You can set up the template for Linux provisioning. To customize the guest, use the available scripts. You can also use your own scripts to allow further setup.

Follow these steps:

1. Disable the network interfaces so that the network does not affect the customization process.

   Note: The network is enabled automatically during the customization.

2. Override the default CDROM device name if needed using the */etc/ca-customize.conf* file.

   CD_DEVICE=/dev/cdrom

   Defines the device name that is used for CD drive.

   Default: /dev/cdrom

3. Set up the automatic start at the end of the boot process.

   - (Valid for SUSE Linux) Create or modify the */etc/init.d/after.local* file:

     ```
     #!/bin/bash
     [ -e /etc/ca-customized ] || /usr/bin/ca-customize
     ```

   - (Valid for Red Hat Linux) Add the following line to the */etc/rc.local* file:

     ```
     [ -e /etc/ca-customized ] || /usr/bin/ca-customize
     ```

4. Shut down the system.

Convert the VM to a Template

The template allows you to create any number of customized virtual machines.

Follow these steps:

1. Shut down the VM.
2. To convert the prepared image to a XenServer template, use XenCenter.

The template appears in CA Virtual Assurance and can be used for customized provisioning.

Once these steps have been performed, the new template can be used to create any number of new customized virtual machines.

How the Customized Provisioning Works

The following steps represent the Customized VM Provisioning Workflow.

1. The platform management service provisions new Linux VM.
2. The platform management service prepares new ISO using customization parameters and attach it to new VM.
3. The platform management service starts the VM.
4. The VM detects that customization ISO is attached. The VM applies the customization changes.
5. If the customization is successful, the VM shuts down. The PMM detects that the VM is stopped. The platform management service starts VM again and finishes provisioning.
6. If the customization failed, the VM is not halted. The platform management service takes the following actions:
   a. Returns a provisioning failure
   b. Sets the provisioning job in exception state

Customization Log

A successful customization is stored in the */etc/ca-customized* file. This file contains a list of the customization changes.

If the customization fails, the logs are stored in the */etc/ca-customized.tmp* file.

How to Prepare Windows Templates for XenServer Provisioning

CA Virtual Assurance supports customized provisioning of new virtual machines (VM) running Windows 2003 R2 Server (32 bit and 64 bit), Windows 2008 (32 bit and 64 bit) or Windows 2008 R2 Server (64 bit). Customization options include a number of settings. For example, changing the built-in Administrator account password, computer name, and the network configuration.

The following diagram illustrates how a system administrator prepares Windows templates for XenServer provisioning.

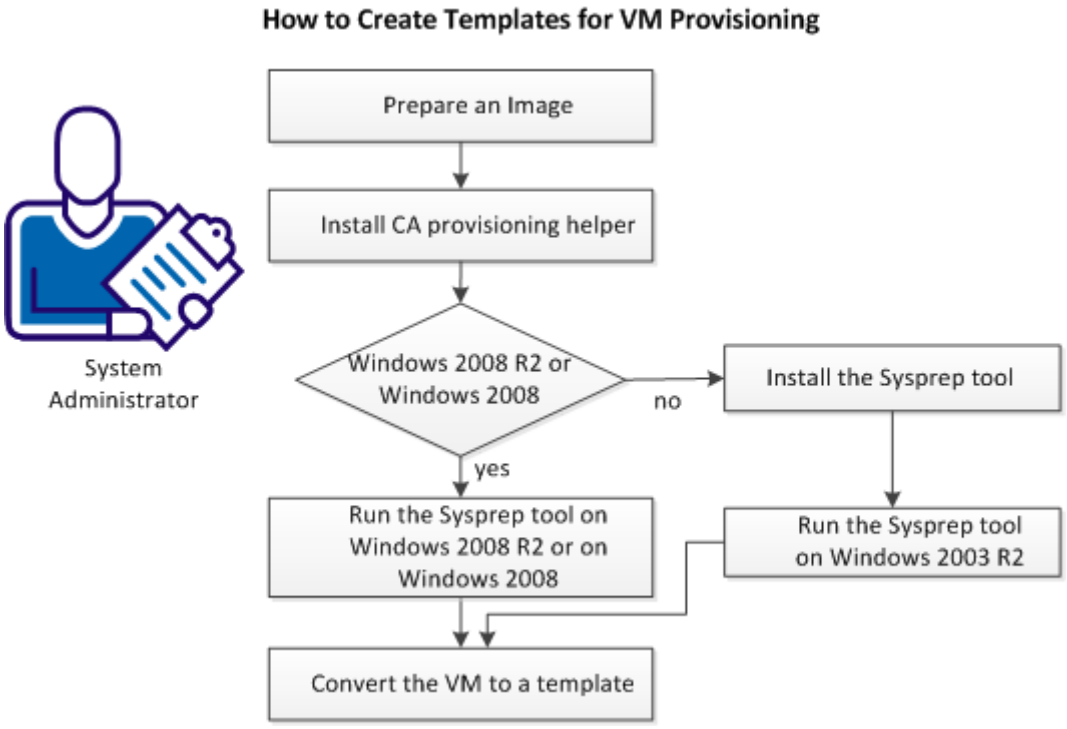

Follow these steps:

1. Prepare a Windows Image (see page 321).
2. Install CA provisioning helper. (see page 321)
3. (Valid on Windows 2003 R2) Install the Sysprep tool. (see page 322)
4. Depending on your operating system select *one* of the following actions:
   - Run Sysprep tool on Windows 2003 R2. (see page 322)
   - Run Sysprep tool on Windows 2008 or on Windows 2008 R2. (see page 322)
5. Convert the VM to a Template in XenCenter (see page 322).

Prepare a Windows Image

When creating a template containing the Windows operating system, prepare the image by following this procedure. Follow the steps to enable CA Virtual Assurance provisioning operations to customize the template. The specific steps differ based on the Windows version.

Follow these steps:

1. Install the Windows operating system on a new virtual machine from scratch.
2. Install the XenTools for Citrix XenServer on the virtual machine.
3. Apply any customizations like user accounts, policy, applications, hotfixes, and so on, that you would like to apply on the new virtual machines.
4. (Valid on Windows 2003) Blank out the built-in Administrator account password.

   Note: If the Administrator password is not empty, SysPrep is unable to set a new password during provisioning and the existing password remains.

Prerequisites for XenServer Environments

Ensure that the following prerequisites are met for the XenServer environment:

- Each XenServer in the resource pool must have SSH or SFTP access enabled.

Install CA provisioning helper

CA provisioning helper enables CA Virtual Assurance to change the virtual machine settings externally.

Follow these steps:

1. Find this utility at <InstallationRoot>\Utilities\Sysprep\CAProvisioningHelper.exe
2. Transfer this executable file to any location on a hard drive of the VM being prepared.
3. Execute CA provisioning helper once from the command line.

The Sysprep Tool

The Microsoft provided Sysprep tools to generalize, freeze and shut down the readily configured Windows installation. The following sections describe how to use the Sysprep tool for Windows 2003 R2 and Windows 2008 R2 in detail.

Install and Run the Sysprep Tool on Windows 2003 R2

On Windows 2003 the Sysprep tools are not installed by default, but can be found on the Windows installation CD-ROM.

Install the Sysprep Tool

Install the Sysprep tool from the Windows installation CD-ROM.

Run the Sysprep Tool on Windows 2003 R2

After you configure the Sysprep tool installation, run the Sysprep tool.

Follow these steps:

1. Locate and open the following CAB file:

   `\SUPPORT\TOOLS\DEPLOY.CAB`

2. Select all files contained in the CAB file and copy them to the following location: %SystemDrive%\Sysprep (normally C:\Sysprep).

   Note: Do not change the directory name.

3. Change to the Sysprep directory and run:

   ```
   sysprep -quiet -reseal -mini -forceshutdown
   ```

Run the Sysprep Tool on Windows 2008 R2

The regular Windows installation process installs all files to perform the SysPrep process. After you configure the Windows installation, perform the following steps:

1. Generate a valid XML response file using the Windows Automated Installation Kit (WAIK) for Windows Server 2008 R2. WAIK is available from the Microsoft Web site.

   Note: The way provisioning requires a dummy unattended response file, or it cannot shut down. The content of the response file is irrelevant, since the provisioning process replaces it, but the file must follow the SysPrep-specific XML schema.

2. Name the generated XML file “sysprep.xml” and place it into the Sysprep directory:

   %SystemRoot%\system32\sysprep

3. Run the following command:

   ```
   sysprep /generalize /oobe /shutdown /unattend:sysprep.xml
   ```

Convert the VM to a Template in XenCenter

The template allows you to create any number of customized virtual machines.

Follow these steps:

1. Shut down the virtual machine.
2. To convert the prepared image to a XenServer template, use XenCenter.

   The template appears in CA Virtual Assurance and can be used for customized provisioning.

Manage VM Status (XenServer)

You can control the status of virtual machines by performing one of the following operations:

- Discover
  - Server
  - Network
- Start
- Suspend
- Shutdown
- Delete From Disk

To control VM status:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Right-click a VM, select Management and one of the following options:

   Discover

   Discovers a server or network.

   Start

   Starts a VM on the specified XenServer host.

   Suspend

   Suspends a running VM on the specified XenServer host and saves its current state. All activity is suspended until you resume the VM.

   Shutdown

   Shuts down a running VM on the specified XenServer host.

   Delete From Disk

   Deletes a VM from the Disk.

   A corresponding wizard appears.

3. Fill in the required information and proceed to the next step.

4. Submit.

   The status operation occurs, and a confirmation message appears. Refresh the interface to view the new VM status. An event appears confirming the result of the operation.

Provision a Citrix XenServer Virtual Machine

You can provision virtual machines by performing the following procedure. Ensure that you prepare a Windows template for VM provisioning.

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Right-click the Citrix XenServer group, select Provisioning, Provision Citrix XenServer Virtual Machine.

   A provisioning wizard appears.

3. Fill in the required information:

   VM Name

   Defines the new VM name.

   Template

   Specifies the Windows provisioning template.

   Administrator Password

   Defines the administrator password for the new VM.

   Product Activation Key

   Defines the Windows 2003 Activation Key.

   Full Name

   Defines the full VM name.

4. (Optional) Fill in the additional information (Workgroup, Memory, CPUs, VM Host, Organization). If you want to use a static IP address, disable the DHCP and provide the IP address, mask, and default gateway.

   Note: The Memory and CPUs settings depend on the Windows provisioning template used.

5. Submit.

   The confirmation message appears.

6. Refresh the Jobs panel to view the progress.

   An event appears confirming the result of the operation.

Huawei GalaX

Huawei GalaX contains the following platforms:

Virtualization Infrastructure Platform

Virtualizes physical resources, such as computing, storage, and network, into virtual resources that are centrally managed, flexibly scheduled, and dynamically allocated. Virtualization Infrastructure is a key platform that is used for building cloud-computing-based data centers.

Cloud Computing Infrastructure Platform

Encapsulates and manages virtual resources that are provided by the virtualization infrastructure platform. Helps carriers and enterprises to build their data center with OMM capability. The management function includes resource management, image management, billing management, scheduling management, and user management.

Operation and Maintenance Management (OMM) Platform

Provides a unified OMM interface for OMM users. OMM users can remotely access the SingleCLOUD OMM System through a web interface. The users can perform operations such as resource management, resource monitoring, and resource statistics reporting.

More information:

How to Configure Huawei GalaX Management Components (see page 326)
How to Create Virtual Private Cloud VLAN (see page 338)
How to Manage Huawei SingleCLOUD Environments (see page 347)
How to Prepare Windows Templates for GalaX Provisioning (see page 356)

How to Configure Huawei GalaX Management Components

As a System Administrator you can configure CA Virtual Assurance to connect to your Huawei GalaX environment and monitor its performance.

How to Configure the GalaX Management Components

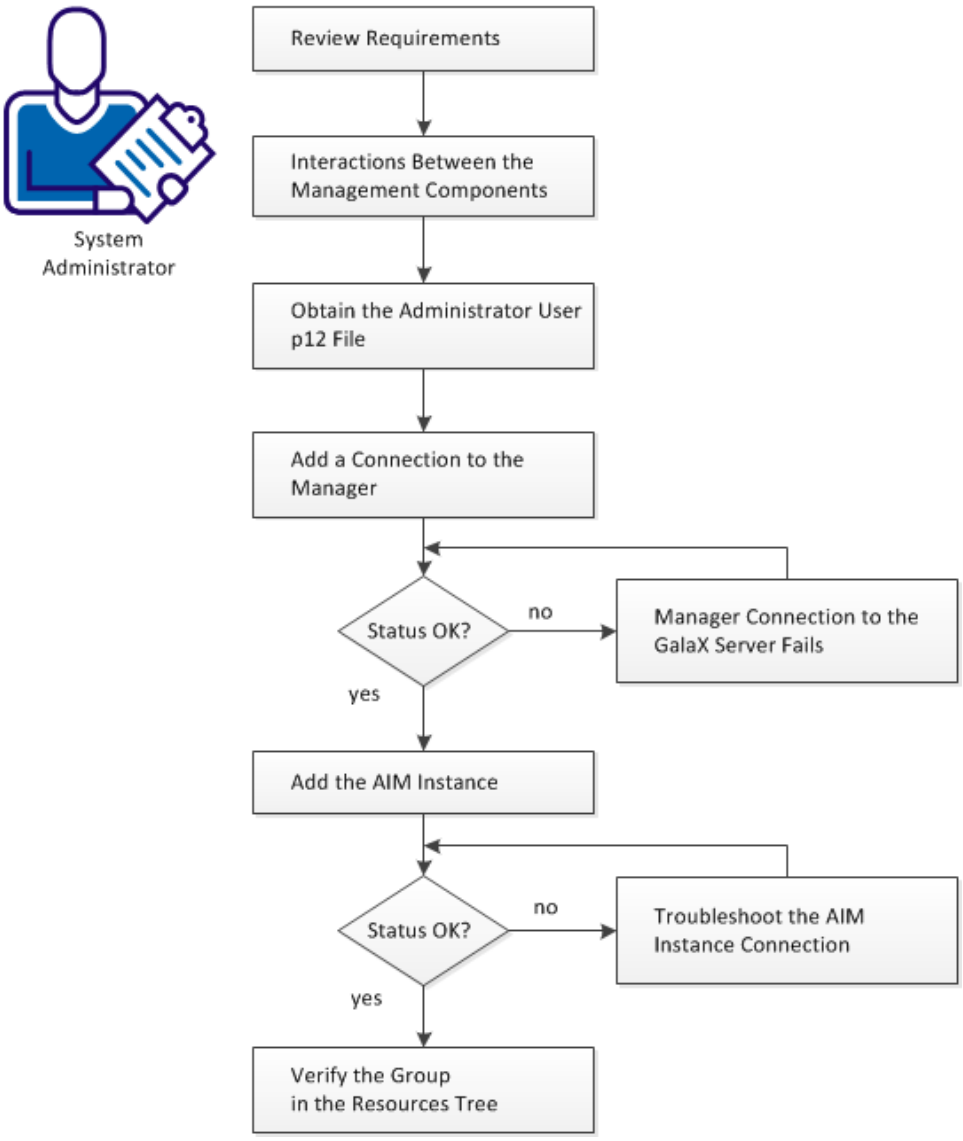

Follow these steps:

Review Requirements

Review the following requirements before configuring the management components of CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Virtual Assurance and SystemEDGE.
- You can access a CA Virtual Assurance manager installation that includes:
  - Platform Management Module (PMM)
  - Application Insight Module (AIM)
  - Monitoring Agent (SystemEDGE)
- You can access the CA Virtual Assurance user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which port to use to access the server in your environment through web services.
  Default HTTP Port: 8773.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote AIM Servers that you want to use.

Review Interactions Between Huawei GalaX Management Components

As a System Administrator, you want to manage a new Huawei GalaX environment with CA Virtual Assurance. CA Virtual Assurance allows you to manage the physical and virtual resources of one or more GalaX environments dynamically. Huawei GalaX consists of Elastic Service Controller (ESC) that communicates with one or more Computing Resource Managers (CRMs). The CRMs communicate with multiple Computing Node Agents (CNAs).

To manage GalaX, CA Virtual Assurance requires network connections between its GalaX Platform Management Module (PMM), GalaX Application Insight Module (AIM), and the Elastic Service Controller (ESC). To establish these network connections, configure the CA Virtual Assurance GalaX management components, that is, GalaX PMM and GalaX AIM.

The GalaX AIM is a SystemEDGE agent plug-in which extends the functional scope of SystemEDGE. The GalaX AIM enables SystemEDGE to monitor the performance of multiple GalaX environments and to evaluate the states of monitored GalaX resources. Based on thresholds, SystemEDGE and the AIM determine the status of a monitored resource and propagate this information to the CA Virtual Assurance manager using SNMP.

The GalaX PMM is a component of the CA Virtual Assurance manager. The PMM is responsible for providing connection and support for all Huawei GalaX operations using SOAP. The PMM manages connections with Computing Resource Manager, performs GalaX-related operations, retrieves data from the AIM, and populates the CA Virtual Assurance Management Database.

The following diagram shows the interaction of the affected components in an example environment with two GalaX ESCs. In general, the GalaX PMM and each GalaX AIM with its multi-instance support can connect to multiple Elastic Service Controllers. The connections shown in the diagram do not specify any limitations. The required network connections are based on TCP/IP, SNMP, and SOAP.

Interactions Between Huawei GalaX Management Components

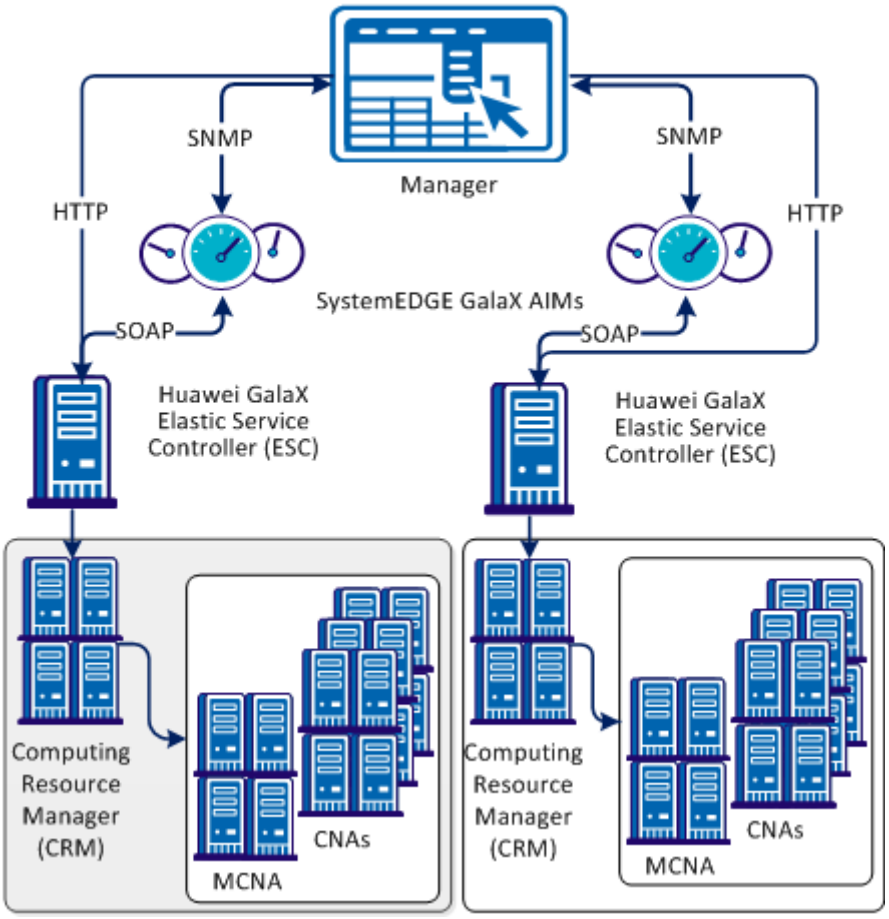

Obtain the Administrator User p12 File

To perform operations in the CA Virtual Assurance UI, obtain the administrator user p12 file from the GalaX environment. The p12 file provides you administrator privileges to configure, monitor, and manage the GalaX environment.

The p12 certification file is generated during the GalaX installation. The certification file is globally unique and only valid for a specific Elastic Service Controller (ESC) API. You cannot use the file to access other GalaX ESC servers.

Before you perform the following procedure, verify the IP address of your GalaX ESC server and the password for the user root.

Follow these steps:

1. Specify a password that you want to use for generating the p12 file.

   You also require this password when you configure the connection between the CA Virtual Assurance manager and the GalaX ESC server.

2. Log in the GalaX ESC server using root.

3. Open a terminal window and run the following command:

   ```
   cd /opt/eucalyptus/.euca
   ```

   This directory contains certification files.

4. To get the names of the digit-signed certification file and private key certification file, run the ls command.

   The file names have the following format:

   - The digit-signed certification file: euca2-admin-*-cert.pem
   - The private key certification file: euca2-admin-*-pk.pem

5. Run the following command:

   ```
   Openssl pcks12 –export –in <digit-signed certification file> -out admin.p12 –inkey <private key certification file>
   ```

   Example:

   ```
   openssl pkcs12 -export -in euca2-admin-109f9d47-cert.pem -out admin.p12 -inkey euca2-admin-109f9d47-pk.pem
   ```

6. The system prompts you: "Enter Export Password"

7. Enter the password that you have specified in Step 1.

   The system generates the required admin.p12 certification file in the /opt/eucalyptus/.euca directory.

8. Copy the admin.p12 file to the server on which the CA Virtual Assurance manager resides. The directory on the server can be arbitrary. You can use a tool like WinSCP to copy the admin.p12 file to that Windows system.

9. The admin.p12 file and your password can now be used to establish a connection between the CA Virtual Assurance manager and the GalaX ESC server.

Add a New GalaX Connection to the Manager

You can add a GalaX 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 Huawei SingleCLOUD from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed GalaX Servers and the associated GalaX AIM Servers.

3. Click  (Add) on the GalaX Servers pane toolbar.

   The New GalaX Server dialog appears.

4. Enter the required connection data (user name, Server, port, P12 file path, and password) and click OK.

   If the network connection is established successfully, the GalaX Server is added to the top right GalaX Servers pane with a green status icon. CA Virtual Assurance discovers the GalaX Server automatically.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the GalaX Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Manager Connection to the GalaX Server Fails

Symptom:

After I have added a server connection under Administration, Configuration, the validation of the connection to the server fails.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify that the used server connection data is still valid. If necessary, update the connection data.
- Verify that the server system is running and accessible.
- Verify that the time difference between the CA Virtual Assurance server and the GalaX server is less than 5 minutes.
- Verify that the services required for the connection are running properly on the server.

To update the server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 server cannot be established, continue with the next procedure.

To verify, if the server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. To find out whether the server has a valid DNS entry and IP address, verify the output of these commands.

   If the server is not in the DNS, add the server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>`

   Enter the correct IP address and server name and save the file. For example:

   `192.168.50.50 myServer`
4. Change to the CA Virtual Assurance user interface, Administration tab, Configuration, Server pane, and click  (Validate) in the upper-right corner.

   If the server credentials and connection data are correct and you can ping the server, the connection can still fail. In this case, it is possible that the server causes the problem. If the connection to the server cannot be established, continue with the next procedure.

To verify, if the time difference between the CA Virtual Assurance server and the GalaX server is less than 5 minutes:

1. To access the GalaX server, contact the system administrator.
2. Check the system time on the GalaX server.
3. Check the system time on the CA Virtual Assurance manager system.
4. If the system time difference is greater than 5 minutes, update the time settings accordingly.

To verify, if all services that are required for the connection are running properly on the server system:

1. Log in to the GalaX server.
2. Verify that the services required for the connection are running properly.
3. If necessary, start or restart a service.
4. Change to the CA Virtual Assurance user interface, server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the server connection.

   If the connection to the server fails, verify the validity of the data you gathered according to the requirements for this scenario.

   Work with the administrator or support to fix the server connection problem.

Add the AIM Instance for GalaX Server

After adding a new GalaX connection to the CA Virtual Assurance manager, add a GalaX AIM instance to manage the new GalaX Server. CA Virtual Assurance then discovers the entire Huawei GalaX environment with all its physical and virtual components.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.

2. Select Huawei SingleCLOUD from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed GalaX Servers and the associated GalaX AIM Servers.

3. Click (Add) on the GalaX AIM Servers pane toolbar.

   The New GalaX AIM Server dialog appears.

4. Open the GalaX AIM Server drop-down list.

   The list of discovered GalaX AIM Servers appears.

5. Select a GalaX AIM Server from the drop-down list.

   CA Virtual Assurance populates the GalaX Server drop-down list with the GalaX Servers listed in the GalaX Servers pane. You can only manage those GalaX Servers for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first so that the AIM server appears in the drop-down list.

6. Select the GalaX Server that you want to manage and click OK.

   A new AIM instance for the selected GalaX Server is added. If the instance is not in an error or stopped state, CA Virtual Assurance starts to discover the associated Huawei GalaX environment. When the discovery process is complete, you can start managing the virtual and physical resources of Huawei GalaX.

Verify the Huawei GalaX in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand Huawei SingleCLOUD group.

   The Huawei GalaX resources appear.

CA Virtual Assurance is now ready to manage the configured Huawei GalaX environment. You can monitor the status and the properties of your resources.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the number of managed objects that are related to virtual and physical resources in your environment. You can move the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Virtual Assurance adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance 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 AIM:

- Verify that the AIM Server is accessible.
- Verify that SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping servername`

2. Verify that the output of the commands has a valid DNS entry and IP address for the AIM server.

   If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 servername

   Enter the correct IP address and AIM server name. For example:

   192.168.50.51 myAIM

4. Click (Validate) in the upper-right corner of the 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 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, AIM Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the AIM Server connection.

   If the error status remains unchanged, verify that the data you gathered is according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered AIM instances in the network, the status icons of several instances show (Disabled). This AIM instance is not managed.

This status appears, if CA Virtual Assurance discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Virtual Assurance manager.
- Edit the existing Server connection and change its managed status to enabled.

How to Create Virtual Private Cloud VLAN

As a System Administrator you want create Virtual Private Clouds with associated virtual machines and virtual disks in the GalaX environment. *A Virtual Private Cloud (VPC)* is a private local network for a Huawei SingleCLOUD user with several virtual machines and associated virtual disks. Because CA Virtual Assurance has already discovered the GalaX environment (see Review Requirements (see page 339)), the CA Virtual Assurance user interface provides the infrastructure for creating the required VPC VLAN resources.

The following diagram illustrates the required steps on how to create VPC VLAN.

How to Create a VPC VLAN

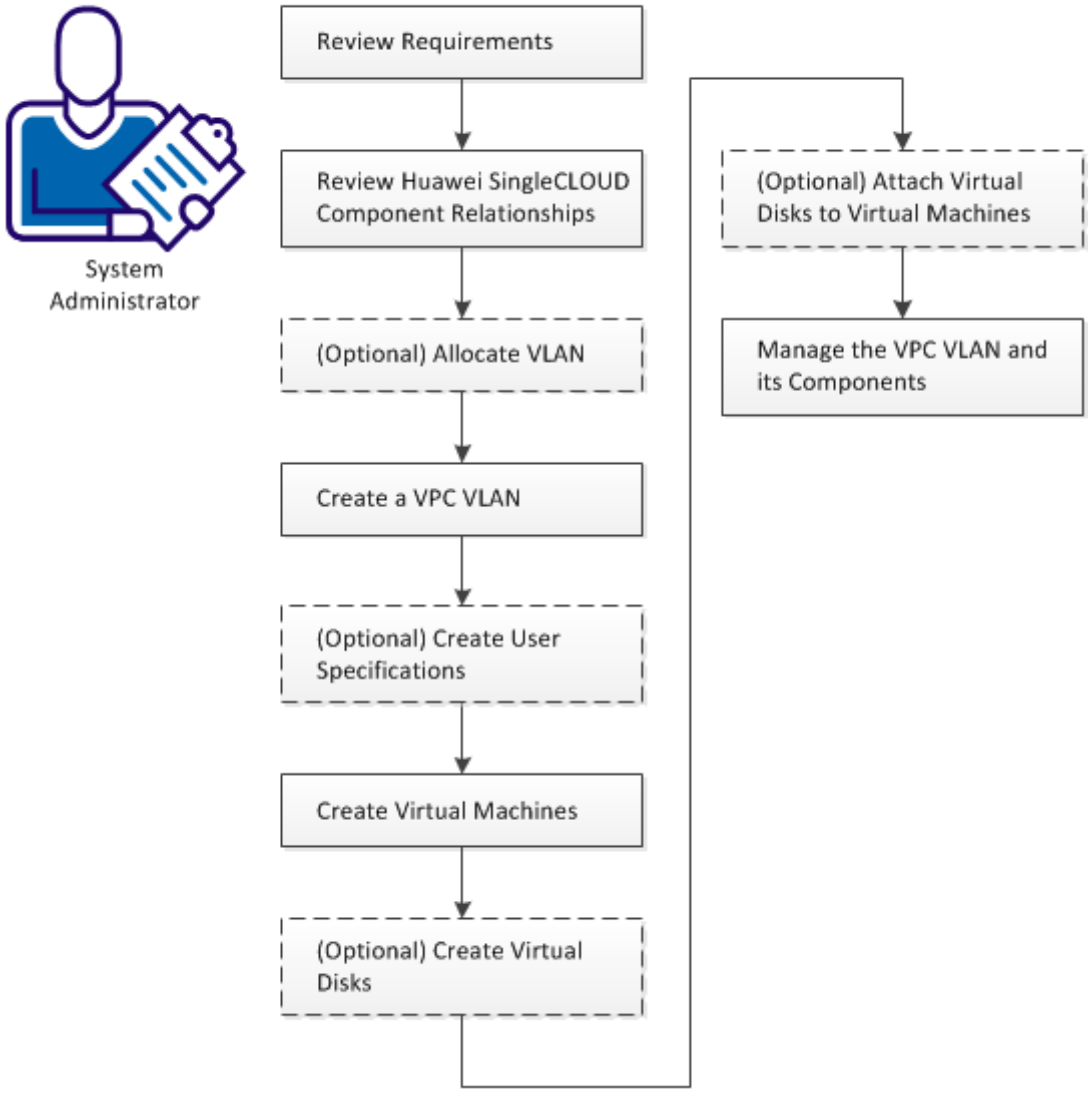

Follow these steps:

Review Requirements

Review the following prerequisites before you set up a Huawei SingleCLOUD instance in CA Virtual Assurance:

- You are familiar with the Huawei GalaX environment.
- You are familiar with the CA Virtual Assurance user interface, and how to provision resources.
- You are familiar with deploying and configuring monitoring software (SystemEDGE).
- CA Virtual Assurance is installed, and you can access the CA Virtual Assurance user interface.
- The Huawei GalaX environment is available and running.
- The servers for Computing Clusters (for virtual machines) and Storage Clusters (for virtual disks) are available in the Huawei GalaX environment.
- Images with operating systems to apply to virtual machines are available in the Huawei GalaX environment.
- A connection between CA Virtual Assurance and a Huawei GalaX server is established.
- The GalaX AIM is configured to monitor the Huawei GalaX server.
- The servers for user VLAN pool and VPC VLAN pool are available.
- CA Virtual Assurance has discovered the Huawei GalaX server and their associated resources such as clusters, storage clusters, and virtual machines.
- Shared disks require Microsoft Cluster Service (MSCS).

Review Huawei SingleCLOUD Component Relationships

The Huawei GalaX environment is part of the Huawei SingleCLOUD solution and designed for cloud computing data centers of cloud service providers or enterprise customers.

The Huawei SingleCLOUD solution consists of a layered architecture. The devices on the physical layer and network layer are integrated into the solution. Based on cluster, distributed storage, NAS storage, and virtualization technologies, these integrated devices provide the storage, computing, and network services to upper-layer services. A Huawei SingleCLOUD instance in CA Virtual Assurance includes the required infrastructure to manage and monitor your Huawei GalaX environment. The Huawei GalaX environment consists of clusters and their associated resources.

The following diagram illustrates the GalaX components of a SingleCLOUD solution that you can manage through CA Virtual Assurance and the dependencies between these components:

Huawei SingleCLOUD GalaX Components and Their Relationships

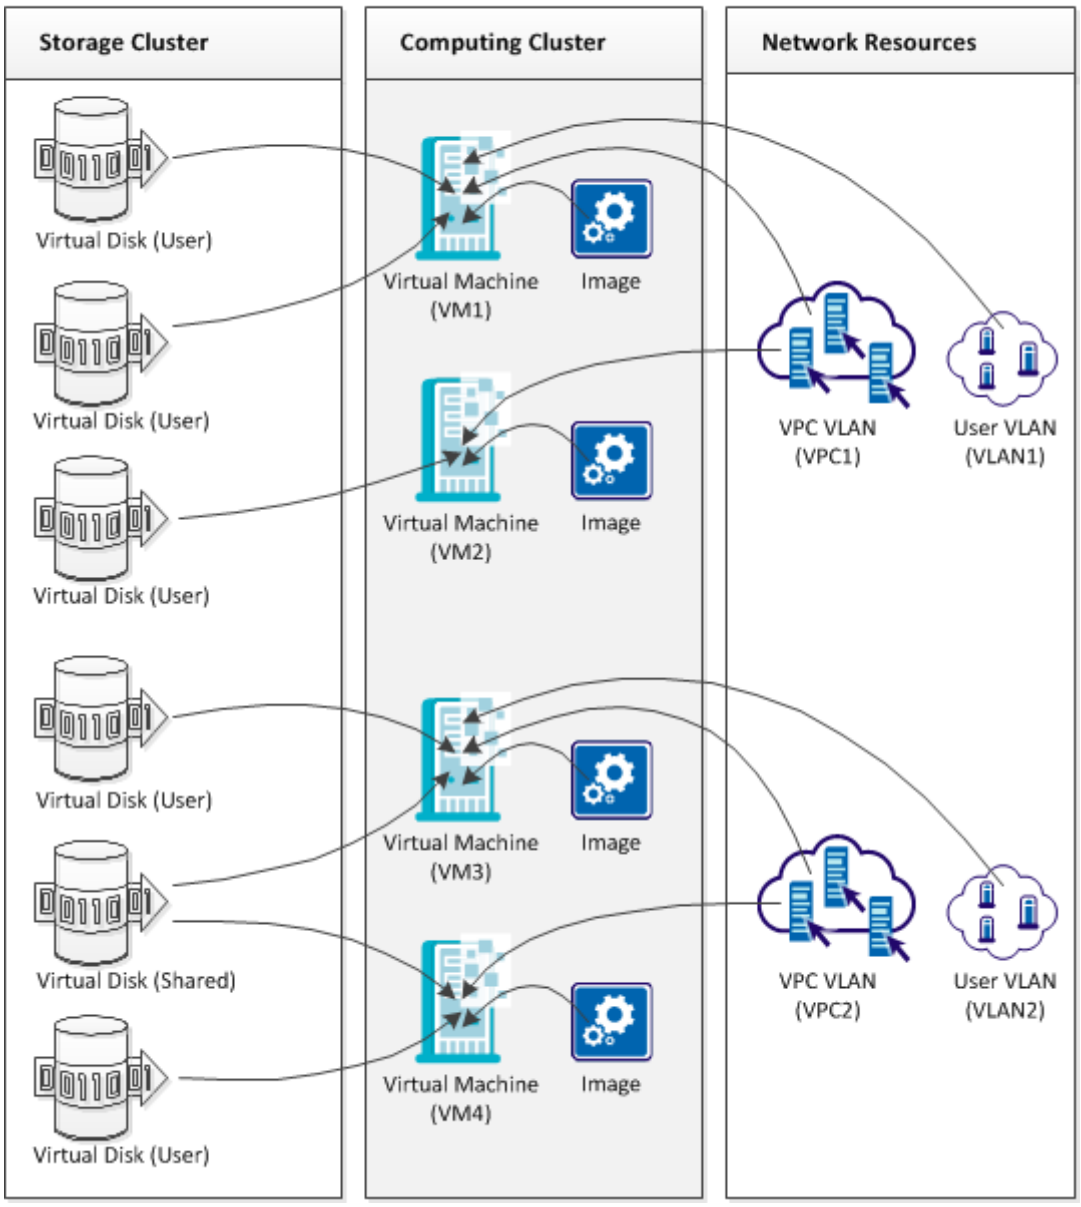

Initially, create a VPC VLAN that provides VLAN access to the virtual machines in the cloud and their users. Optionally, you can add a User VLAN to a virtual machine. A virtual machine in the Computing Cluster requires an appropriate image and the VPC VLAN to which the virtual machine belongs. The image contains the operating system and applications for this virtual machine.

You can then create virtual disks in the Storage Cluster and attach these disks to the appropriate virtual machines to store the user-specific data. Two types of virtual disks are supported: User disks and shared disks. User disks have a one-to-one relationship to virtual machines and shared disks can have a one-to-many relationship. Shared disks require Microsoft Cluster Service (MSCS) support.

(Optional) Allocate VLAN

Because a Virtual Private Cloud object requires VLAN, allocate VLAN initially.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder and select the appropriate SingleCLOUD server.

   The right pane refreshes and displays the Resource Management and Network Management tabs.

3. Click Network Management, VLAN.

   A list of existing VLAN objects appears.

4. Click  (Add) on the VLAN pane toolbar.

   The Allocate VLAN dialog appears.

5. Specify a VLAN name, select a cluster from the drop-down menu, specify the Method (Automatically or Manually Fill), and click OK.

   The VLAN is allocated.

Create a VPC VLAN

A VPC serves as a private local network for a cloud user with several virtual machines and associated virtual disks.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder and select the appropriate SingleCLOUD server.

   The right pane refreshes and displays the Resource Management and Network Management tabs.

3. Click Network Management, VPC.

   A list of existing VPC instances appears.

4. Click (Add) on the VPC pane toolbar.

   The Create VPC dialog appears.

5. Specify a VPC name, select a cluster from the drop-down menu, assign a VLAN (automatically or manually from the list), and click OK.

   The VPC instance is created.

(Optional) Create User Specifications

A User Specification is a set of configuration values for CPU, memory, and system volume size that you can use for creating virtual machines.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder and select the appropriate SingleCLOUD server.

   The right pane refreshes and displays the Resource Management and Network Management tabs.

3. Click Resource Management, User Specification.

   A list of existing User Specifications appears.

4. Click (Add) on the User Specification pane toolbar.

   The Create User Specification dialog appears.

5. Specify a User Specification name and values for CPU, memory and system volume size. Click OK.

   The User Specification is created.

Create Virtual Machines

A virtual machine requires an image for the system volume, CPU, memory and disk space settings, VPC and NIC specifications.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder and right-click the appropriate Computing Cluster.

   A pop-up menu opens.

3. Select Management, Create VM from a template.

   The Create VM dialog opens.

4. Specify the following parameters and click OK.

   - Number of VMs
   - VM Name
   - Image ID
   - User Specification or CPU, memory, disk space
   - VPC VLAN
   - (Optional) Additional Network Interface Controllers (NIC)
   - Quality of Service (QoS) settings
     - Memory Reserved
     - CPU Reserved
     - CPU Limit
     - High Availability
     - NIC Speed Limit

   CA Virtual Assurance creates the specified virtual machines. The virtual machines belong to the assigned VPC VLAN. To get a list of virtual machines, open the Details tab in the Computing Cluster panel.

The following parameters require further explanation:

Memory Reserved

Specifies the minimum proportion of physical memory that is allocated to a virtual machine. The reservation is defined as a percentage (%) and you can assign values from 0 to 100 percent.

Example: If you set memory to 2 GB and reservation to 25 percent, the system ensures at least 512 MB for the virtual machine.

CPU Reserved

Specifies the minimum proportion of the physical CPU performance that is reserved for this virtual machine. The reservation is defined in the percent (%) and you can assign a value of 0, 50, or 100 percent.

Example: If you set reservation to 50 percent, the system ensures at least 50 percent CPU time for each CPU core.

CPU Limit

Specifies the maximum percentage of CPU performance that this virtual machine can allocate.

Note: The value of the limit must be greater than or equal to the value that you have specified for the reservation.

(Optional) Create Virtual Disks

Virtual disks are intended to store user-specific data and are attached to virtual machines.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder and right-click the appropriate Storage Cluster.

   A pop-up menu opens.

3. Select Management, Create Disk.

   The Create Disk dialog opens.

4. Specify the following parameters and click OK.
   - Disk Name
   - Disk Type (User Disk or Shared Disk). A User Disk can be attached to one virtual machine. A Shared Disk can be attached to multiple virtual machines.
   - Dynamic Allocation (Ordinary or Thin Provisioning)

     Thin Provisioning requires IP SAN device support.
   - Disk Size (GB)
   - Description of the virtual disk

   CA Virtual Assurance creates the virtual disk. To get a list of virtual disks, open the Details tab in the Storage Cluster panel.

(Optional) Attach Virtual Disks to Virtual Machines

According to the specified virtual disk type, you can attach User Disks to one virtual machine and Shared Disks to multiple virtual machines.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.
2. Expand the Huawei SingleCLOUD folder and select the appropriate Storage Cluster.

   The Storage Cluster panel opens and lists the specified virtual disks.
3. Select the virtual disk that you want to attach and click the attach icon.

   A list of available virtual disks appears.
4. Select the appropriate virtual machines and click OK.

   The virtual disk is attached.

Manage the VPC VLAN and its Components

You have specified virtual machines with attached virtual disks which use VLAN to communicate. These resources belong to a Virtual Private Cloud that you can manage through CA Virtual Assurance.

How to Manage Huawei SingleCLOUD Environments

Because the most part of the user interface is self-explanatory, this scenario is a just a guideline to walk through the object hierarchy of Huawei SingleCLOUD environments and to explore its associated management capabilities.

How to Manage Huawei SingleCLOUD Environments

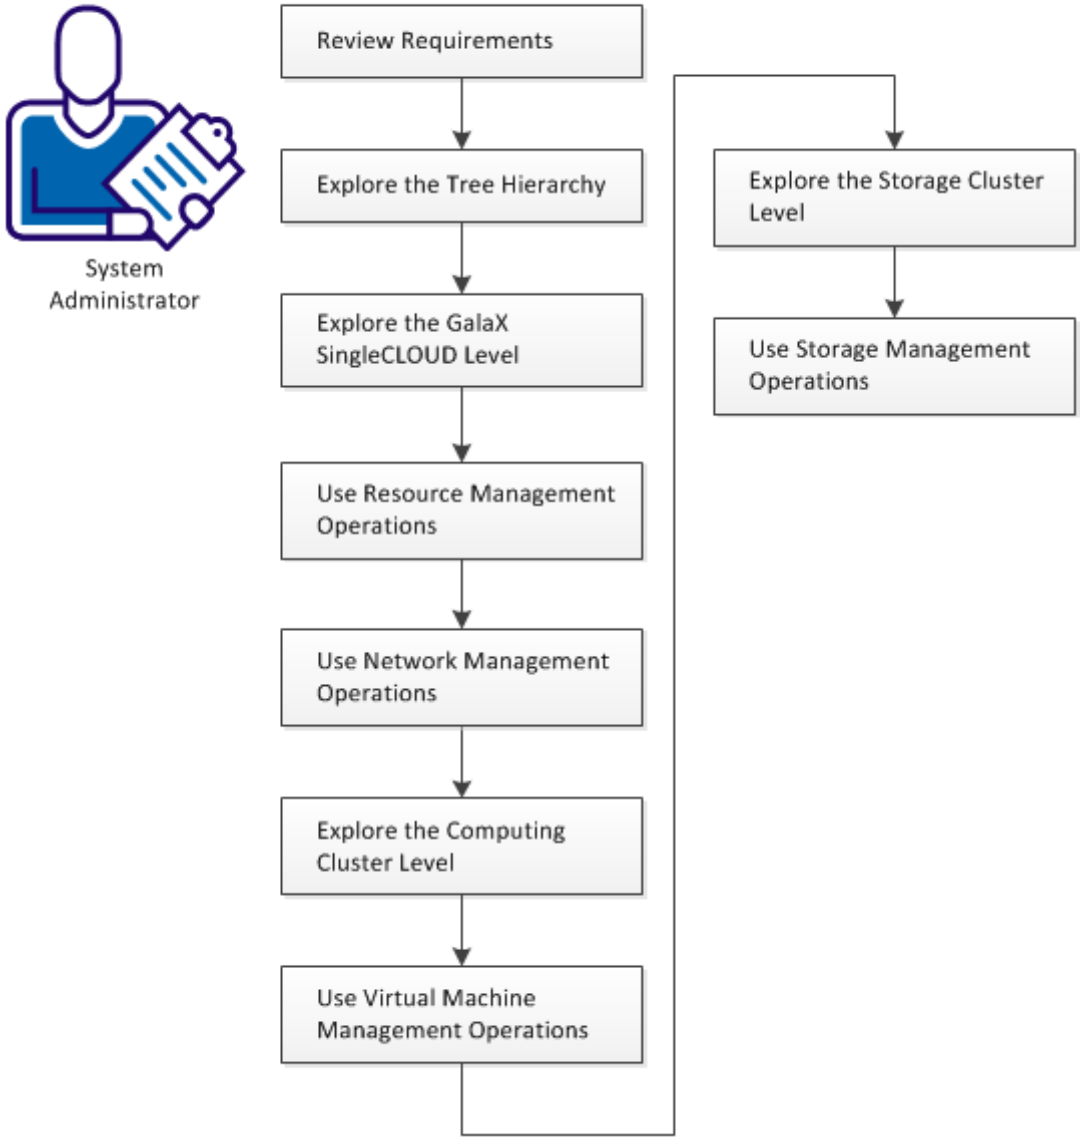

Follow these steps:

Review Requirements

Review the following prerequisites before you manage Huawei SingleCLOUD instance in CA Virtual Assurance:

- You are familiar with the Huawei GalaX environment.
- You are familiar with the CA Virtual Assurance user interface, and how to provision resources.
- You are familiar with deploying and configuring monitoring software (SystemEDGE).
- CA Virtual Assurance is installed, and you can access the CA Virtual Assurance user interface.
- The Huawei GalaX environment is available and running.
- The servers for Computing Clusters (for virtual machines) and Storage Clusters (for virtual disks) are available in the Huawei GalaX environment.
- Images with operating systems to apply to virtual machines are available in the Huawei GalaX environment.
- A connection between CA Virtual Assurance and a Huawei GalaX server is established.
- The GalaX AIM is configured to monitor the Huawei GalaX server.
- The servers for user VLAN pool and VPC VLAN pool are available.
- CA Virtual Assurance has discovered the Huawei GalaX server and their associated resources such as clusters, storage clusters, and virtual machines.
- Virtual Private Clouds with Virtual Machines are available.

Explore the Tree Hierarchy

The Huawei SingleCLOUD folder represents the service level at the top. The SingleCLOUD service consists of one or more Elastic Service Controllers (ESC). Each ESC can control multiple Computing Clusters and Storage Clusters with virtual machines and virtual disks.

The following diagram shows the object hierarchy of the Huawei SingleCLOUD folder:

```
..
└──Huawei SingleCLOUD
    └──ESC Node Name and IP Address
        └──Computing Clusters
            └──Virtual Machines
        └──Storage Clusters
            └──Virtual Disks
```

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder.

   - To open the list of Huawei SingleCLOUD events, select the Huawei SingleCLOUD object.
   - To access Resource Management and Network Management, select the ESC node.
   - To get a list of available virtual machines or to create a virtual machine, select a Computing Cluster.
   - To get a list of available virtual disks or to create a virtual disk, select a Storage Cluster.

Explore the GalaX SingleCLOUD Server Level

The GalaX SingleCLOUD resides at the second level in the tree hierarchy.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder.

   The folder hierarchy appears.

3. Select an ESC node.

   The Resource Management and Network Management tabs appear.

   - Use Resource Management for snapshots, images, and user specifications.
   - Use Network Management for VPC VLAN and User VLAN.

Use Resource Management Operations

The user interface provides the following operations under the Resource Management tab:

- View snapshots and their properties
- Restore a snapshot to a virtual machine
- Delete a snapshot
- View images and their properties
- View User Specifications and their properties
- Create a User Specification
- Edit a User Specification
- Delete a User Specification

The usage and the dialogs are self-explanatory. If necessary, you can move the cursor over an icon to get a tooltip.

Use Network Management Operations

The user interface provides the following operations under the Resource Management tab:

- Create a VPC VLAN
- View VPC VLANs and their properties
- Delete a VPC VLAN
- Allocate User VLAN
- Delete User VLAN

The usage and the dialogs are self-explanatory. If necessary, you can move the cursor over an icon to get a tooltip.

Explore the Computing Cluster Level

The Computing Cluster resides at the third level in the tree hierarchy.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder.

   The folder hierarchy appears.

3. Select or expand a Computing Cluster.

   The list of available virtual machines appears.

4. Right-click a Computing Cluster to create a virtual machine.

   You require an image that contains a disk and an operating system for the system volume (Resource Management), a VPC VLAN, a User Specification (optional), a User VLAN (optional).

5. Right-click a virtual machine to perform a virtual machine management operation.

   Operations which are not applicable are unavailable.

Use Virtual Machine Management Operations

The user interface provides the management operations for virtual machines when you right-click a virtual machine. The usage and the dialogs are self-explanatory.

- Delete VM
- Restart VM
- Power On VM
- Power Off VM
- Safe Restart VM (Shuts down the operating system)
- Safe Power Off VM (Shuts down the operating system)
- Hibernate VM
- Wake Up VM
- Modify VM Name
- View Initial Password
- Set Boot Order
- Rollback Snapshot
- Create VM Snapshot

The following management operations require more explanation:

- Modify CPU Configuration and QoS
- Modify Memory Configuration and QoS
- VNC Login
- Mount/Unmount Tools

Modify CPU Configuration and QoS

Specify the following values:

Number of CPUs

Specifies the number of CPU cores that are allocated to a virtual machine. The maximum number of CPU cores you can assign to a virtual machine is eight.

Example: If you set the number to five, then five CPU cores are available for the virtual machine.

Reservation

Specifies the minimum proportion of the physical CPU performance that is reserved for this virtual machine. The reservation is defined in the percent (%) and you can assign a value of 0, 50, or 100 percent.

Example: If you set reservation to 50 percent, the system ensures at least 50 percent CPU time for each CPU core.

Limit

Specifies the maximum percentage of CPU performance that this virtual machine can allocate.

Note: The value of the limit must be greater than or equal to the value that you have specified for the reservation.

Modify Memory Configuration and QoS

Specify the following values:

Memory

Specifies the amount of memory that you assign to a virtual machine. The memory is defined in megabytes (MB) and ranges between 512 MB and 256 GB.

Example: If you set the memory to 512 MB, 512 MB is the maximum amount of memory that the virtual machine can allocate.

Reservation

Specifies the minimum proportion of physical memory that is allocated to a virtual machine. The reservation is defined as a percentage (%) and you can assign values from 0 to 100 percent.

Example: If you set memory to 2 GB and reservation to 25 percent, the system ensures at least 512 MB for the virtual machine.

VNC Login

Before you can use VNC to access your VMs, VNC Login requires an initial setup: Download VncViewer.jar and install it on your CA Virtual Assurance manager system.

Follow these steps:

1. Log in the CA Virtual Assurance manager server, open the user interface, expand the Explore Tree, right-click a Huawei SingleCloud VM, and select Management, VNC Login.

   A message appears and gives you instructions how to proceed.

2. From the CA Virtual Assurance manager server, connect to your ESC or OMM server and download VncViewer.jar from the following directory:

   `/opt/omm/oms/webapps/oms/business/resourcemanage/virtualresources`

3. Click VNC Login again.

   The message dialog opens.

4. Click the message in the dialog.

   The Upload File dialog opens.

5. Click Browse ..., navigate to the downloaded VncViewer.jar file, and click open.

   The File Path appears in the dialog.

6. Click OK.

   CA Virtual Assurance uploads VncViewer.jar to the *Install_Path**product*\tomcat\webapps\UI directory.

   The VNC Viewer automatically opens and connects to the VM.

When you have completed this one-time procedure, VNC Login is available and you can remotely access any Huawei SingleCloud VM in your environment.

Mount/Unmount Tools

To provide the maximum of functionality, install the SingleCloud Tools on your VMs.

Follow these steps:

1. Log in the CA Virtual Assurance manager server, open the user interface, expand the Explore Tree, right-click the VM, and select Management, Mount/Unmount Tools.

   CA Virtual Assurance displays the current VM status and SingleCloud Tools status in a dialog.

2. To change the SingleCloud Tools status to mount/unmount, click OK.

3. After successfully mounting the SingleCLOUD tools on the VM, install the PV driver. If the VM runs on the Linux OS, restart the VM and install the PV driver.

Resource Allocation Best Practices

Specify resource allocation settings (reservation and limit) that are appropriate for the virtual machines in your Huawei SingleCLOUD environment.

The following guidelines help you to achieve better performance for your virtual infrastructure:

- Use reservations to specify the *minimum* acceptable amount of CPU or memory, not the amount that you want to have available. The host assigns more resources as available based on the estimated demand and the limit for your virtual machine. The amount of CPU or memory that you specified through reservations remains unchanged when you modify the environment, such as adding or removing virtual machines.
- When specifying reservations for virtual machines, do not commit all resources. Plan to leave an appropriate portion unreserved, because when you move closer to reserving the full system capacity, changing reservations becomes increasingly difficult.

Explore the Storage Cluster Level

The Storage Cluster resides at the third level in the tree hierarchy.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Management, Resources.

   The Explore tree opens.

2. Expand the Huawei SingleCLOUD folder.

   The folder hierarchy appears.

3. Expand a Storage Cluster.

   The list of available virtual disks appears.

4. Right-click a Storage Cluster to create a virtual disk.

   The following parameters require more explanation:

   Disk type: User disk

   Can be attached to one virtual machine.

   Disk type: Shared disk

   Can be attached to multiple virtual machines.

   Dynamic allocation: Thin provisioning

   Reserves the specified disk space, but does not dedicate the entire reserved space to the disk until the space is required to store data. The size of a thin provisioned virtual disk grows according to the amount of data that is stored.

   Thin provisioning allows you to overcommit the datastores and to increase the storage utilization by minimizing the disk space that is reserved but not used.

5. Right-click a virtual disk in the Explore tree to perform a virtual disk management operation.

   You can view the details of the virtual disk or can delete the virtual disk.

Use Storage Management Operations

The user interface provides the management operations for virtual machines when you right-click a virtual machine. The usage and the dialogs are self-explanatory.

- Delete Virtual Disk
- View Details of Virtual Disks
- Select a virtual disk to view events of the virtual disk

Select a Storage Cluster to open a list of available virtual disks. The following operations are available:

- Attach Virtual Disk
- Detach Virtual Disk
- Delete Virtual Disk

The usage is self-explanatory. If necessary, you can move the cursor over an icon to get a tooltip.

How to Prepare Windows Templates for GalaX Provisioning

CA Virtual Assurance supports customized provisioning of new virtual machines (VM) running Windows 2003 R2 Server (32 bit and 64 bit), Windows 2008 (32 bit and 64 bit) or Windows 2008 R2 Server (64 bit). Customization options include a number of settings. For example, changing the built-in Administrator account password, computer name, and the network configuration.

The following diagram illustrates how a system administrator prepares Windows templates for GalaX provisioning.

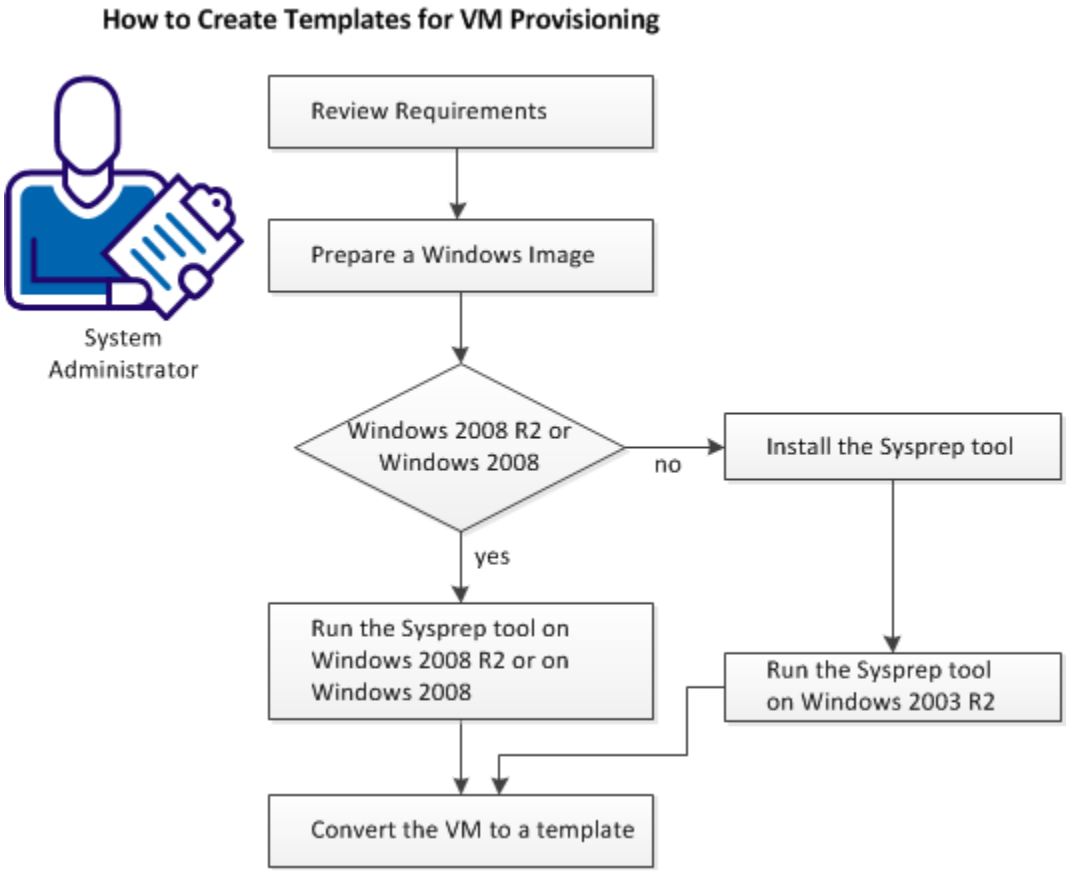

The Microsoft sysprep tool lets you generalize, freeze and shut down the readily configured Windows installation. The following sections describe how to use the Sysprep tool for Windows 2003 R2 and Windows 2008 R2 in detail.

On Windows 2003 the Sysprep tools are not installed by default, but can be found on the Windows installation CD-ROM.

Follow these steps:

Review Requirements (see page 357)
Prepare a Windows Image (see page 358)
Run the Sysprep Tool on Windows 2003 R2 (see page 358)
Run the Sysprep Tool on Windows 2008 R2 (see page 358)
Convert the VM to a Template in GalaX (see page 359)
Using Provisioned Virtual Machines (see page 359)

Review Requirements

Review the following prerequisites before you create templates for virtual machine provisioning in CA Virtual Assurance:

- You are familiar with the Huawei GalaX environment.
- You are familiar with the CA Virtual Assurance user interface, and how to provision resources.
- CA Virtual Assurance is installed, and you can access the CA Virtual Assurance user interface.
- The Huawei GalaX environment is available and running.
- The servers for Computing Clusters (for virtual machines) and Storage Clusters (for virtual disks) are available in the Huawei GalaX environment.
- The servers for user VLAN pool and VPC VLAN pool are available.
- CA Virtual Assurance has discovered the Huawei GalaX server and their associated resources such as clusters, storage clusters, and virtual machines.

Prepare a Windows Image

When creating a template containing the Windows operating system, prepare the image by following this procedure. Follow the steps to enable CA Virtual Assurance provisioning operations to customize the template. The specific steps differ based on the Windows version.

Follow these steps:

1. Install the Windows operating system on a new virtual machine from scratch.
2. Install the SingleCloud Tools on the virtual machine.
3. Apply any customizations like user accounts, policy, applications, hotfixes, and so on, that you would like to apply on the new virtual machines.

Run the Sysprep Tool on Windows 2003 R2

After you configure the Sysprep tool installation, run the Sysprep tool.

Follow these steps:

1. Locate and open the following CAB file:

   ```
   \SUPPORT\TOOLS\DEPLOY.CAB
   ```

2. Select all files contained in the CAB file and copy them to the following location: %SystemDrive%\Sysprep (normally C:\Sysprep).

   Note: Do not change the directory name.

3. Change to the Sysprep directory and run:

   ```
   sysprep -quiet -reseal -mini -forceshutdown
   ```

Run the Sysprep Tool on Windows 2008 R2

The regular Windows setup installs all files to perform the Sysprep process. After you configure the Windows installation, perform the following steps:

1. Change to the following directory:

   C:\Windows\system32\sysprep

2. Run the following command:

   ```
   sysprep /generalize /shutdown
   ```

The sysprep command prepares the image for the installation and shuts down the virtual machine. The generalize parameter removes all unique system information such as computer name, log files, restore points, and hardware-specific information.

Convert the VM to a Template in GalaX

After the sysprep command has shut down the virtual machine, change to the SingleCloud user interface to create the template.

Follow these steps:

1. Log in the SingleCloud user interface.
2. Click the VM tab and select the virtual machine that you have prepared with sysprep.
3. Right-click the virtual machine and select Export Image.

   The Export Image dialog opens.
4. Specify a file name, set the Image Type to Ghost, and click OK.

   The virtual machine is saved as a Ghost image.
5. Register the Ghost image in the SingleCloud user interface.

   You can now use the Ghost image as a template for provisioning.

Using Provisioned Virtual Machines

A template that has been created according to the previous scenario (see page 356), causes the following behavior for provisioned virtual machines:

When you initially start a provisioned virtual machine, the start process waits for your input, such as locale setting, product key, EULA, and lets you specify the hostname for this particular machine.

To access the virtual machine, verify that VNC is available.

IBM PowerVM (LPAR)

IBM PowerVM systems provide 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 systems. IVM is an extension to the Virtual I/O Server and can only manage the local PowerVM system.

The LPAR AIM enables SystemEDGE to monitor PowerVM resources.

The LPAR Platform 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

IBM PowerVM Server Administration Overview

The CA IBM PowerVM component of CA Virtual Assurance lets you monitor and manage IBM PowerVM resources. The monitored and managed resources consist of the following types:

- Hardware Management Console (HMC)
- Integrated Virtualization Manager (IVM)
- Virtual IO Server (VIOS)
- Managed System (POWER Server)
- Logical Partition (LPAR)

The *Hardware Management Console (HMC)* is an external appliance that is used to perform management tasks on IBM PowerVM Systems. 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 Systems, providing a single point of control in large PowerVM environments.
The *Integrated Virtualization Manager (IVM)* is an enhancement of the Virtual I/O Server (VIOS) and allows you to manage a single POWER System. IVM lets you create and manage LPARs. IVM enables management of VIOS functions and provides a web-based user interface.

A *Virtual I/O Server (VIOS)* is a special logical partition that is 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.

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 communicate in the network as separate systems.

To manage IBM PowerVM resources, provide SSH access credentials to HMC/IVM Servers and Virtual I/O Servers.

You can configure CA Virtual Assurance to manage PowerVM resources by using CA Virtual Assurance Administration, Configuration, Provisioning, the IBM PowerVM Group.

The following panels are available:

- HMC/IVM Servers
- Virtual I/O Servers
- LPAR AIM Servers

The LPAR AIM Server is the system on which SystemEDGE and the LPAR AIM run. The LPAR AIM can run on the local CA Virtual Assurance manager system or on a remote Windows server. The LPAR AIM is a multi-instance AIM that can connect to multiple HMCs or IVMs. Once the AIM starts managing an HMC or IVM server, the AIM discovers and manages all P-Servers that are connected to this HMC or IVM server.

How to Configure the PowerVM Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

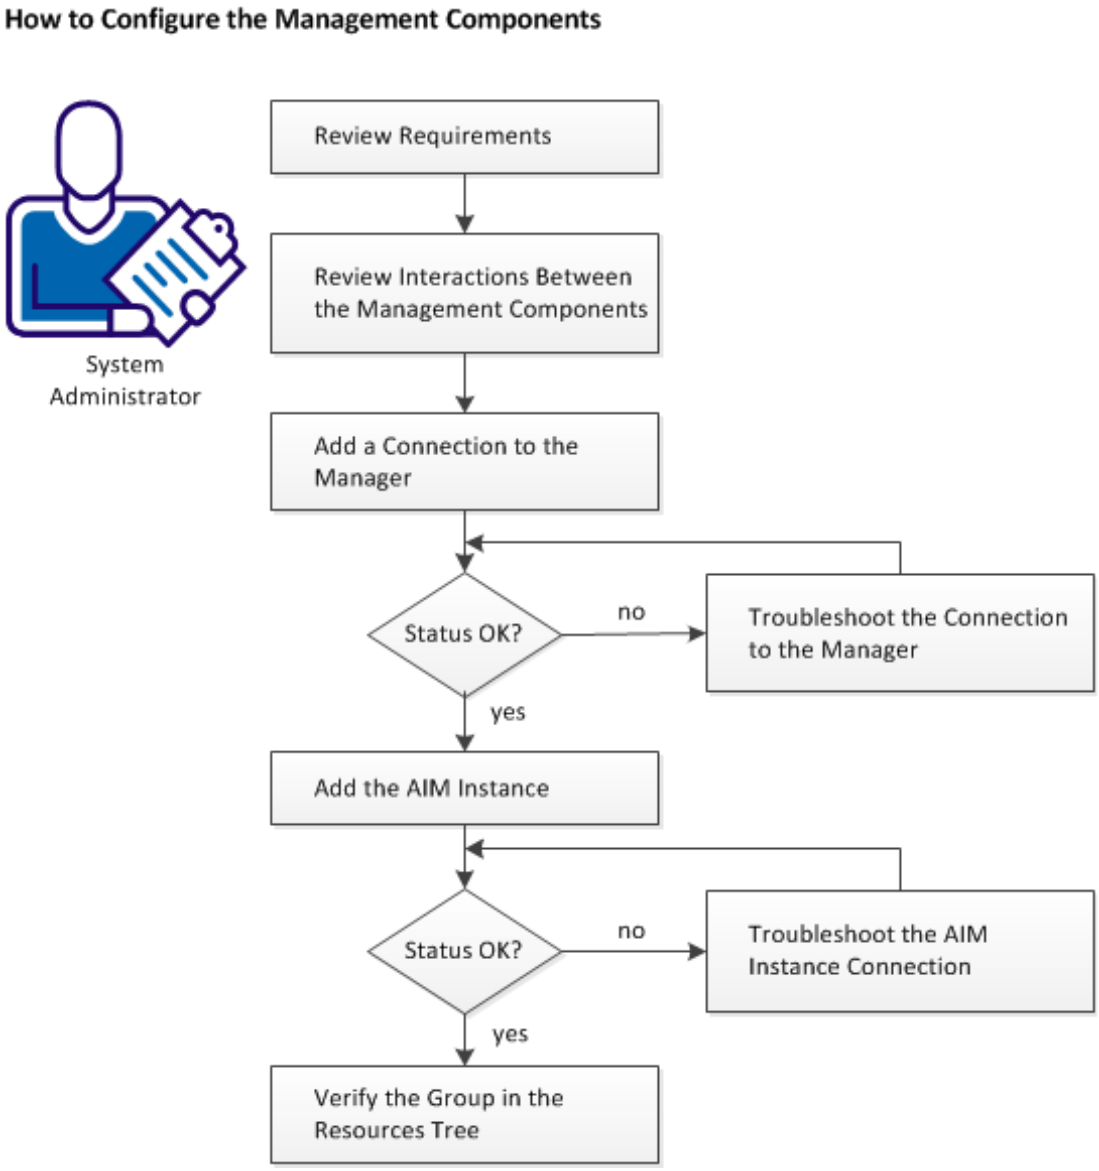

Follow these steps:

Review Requirements

Review the following requirements before configuring the management components of CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Virtual Assurance and SystemEDGE.
  - You can access a CA Virtual Assurance manager installation that includes:
  - Platform Management Module (PMM)
  - Application Insight Module (AIM)
  - Monitoring Agent (SystemEDGE)
- You can access the CA Virtual Assurance user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote AIM Servers that you want to use.

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

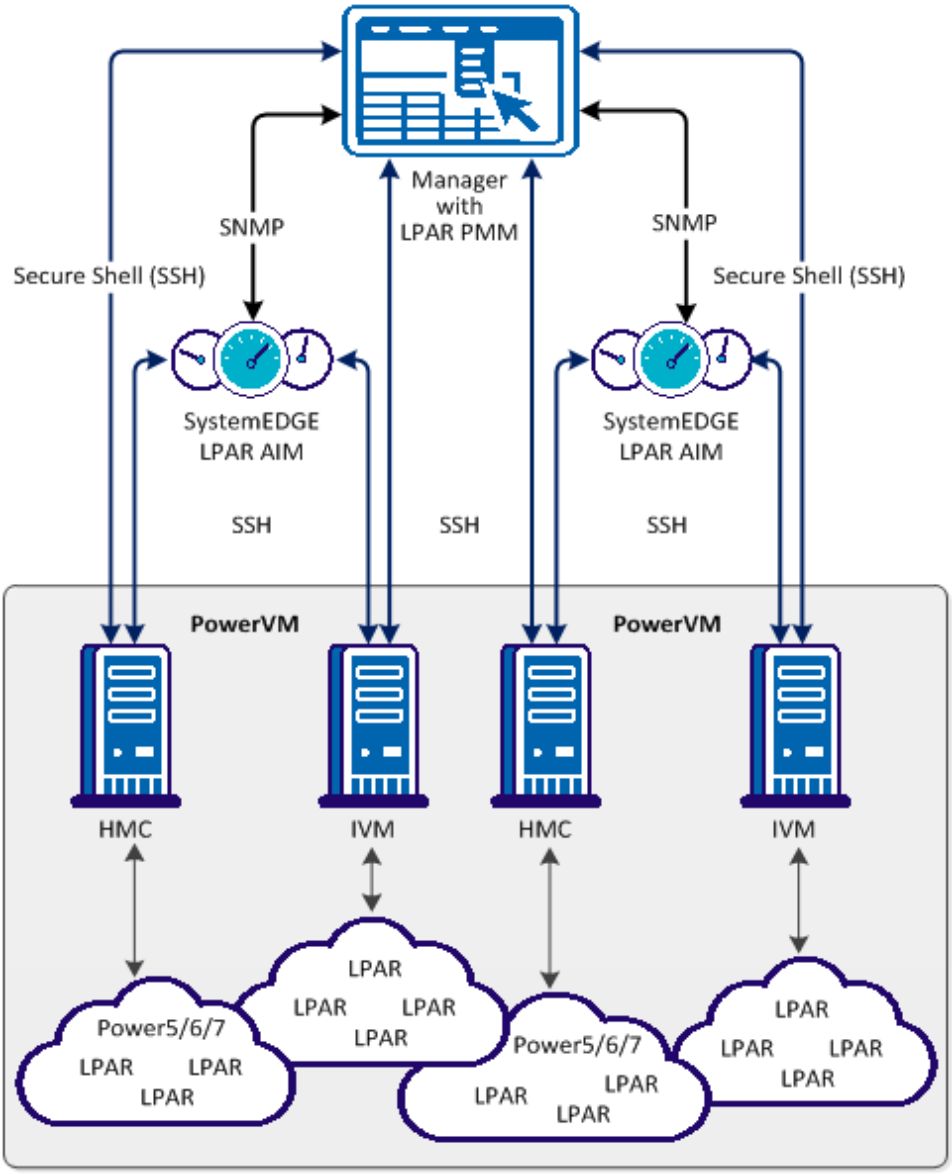

After the installation, configure your environment by adding the required connection information for each required HMC/IVM and Virtual I/O Server. Use *one* of the following methods:

- Administration tab of the user interface
- 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.

IBM PowerVM Configuration Use Cases

The following use cases describe the handling of LPAR AIM instance entries for managed PowerVM environment in the Administration tab:

- You add an HMC Server and an LPAR AIM instance.

  The AIM discovers:

  - Power Systems associated with the HMC.
  - Virtual I/O Servers associated with the Power systems. The AIM applied the default VIOS credentials specified when adding the HMC.

  Important! If you do not specify the default VIOS credentials for an HMC Server, provide the VIOS credential for *each* VIOS in the Virtual I/O Servers panel to complete the configuration of the discovered VIOS. If the default VIOS credentials do not apply to a particular VIOS, you can overwrite the credentials in the Virtual I/O Servers panel.

- Preferred AIM

  Two AIMs can manage one HMC. When second AIM is added it becomes the redundant AIM, the AIM added first becomes the preferred AIM. The status of the HMC under the redundant AIM becomes Suspended. This status reflects that the HMC is managed by the preferred AIM. You can change the preferred AIM in the HMC/IVM Servers panel.

- Dual HMC feature supports configurations where a P System is associated with two HMC servers.

  P-Server and associated HMC servers are one atomic management entity and as such have to be managed by one AIM. The dual HMC configuration is only supported in the scope of one AIM. For example, one Power System, P1 is connected to two HMC servers, HMC1 and HMC2. Both HMC servers are managed by one AIM, AIM1.

- Dual HMC failover

  If the preferred HMC fails, the redundant HMC starts managing your System automatically. The redundant HMC becomes the current one. However, when the preferred HMC becomes available, the current HMC is not changed. To manage your System by the preferred HMC again, change the current HMC in the LPAR AIM Server panel on the Administration, Configuration tab.

  Note: If a preferred HMC fails, the redundant HMC manages your System. After the failover, you can change the current HMC for your System manually.

- You do not specify the default VIOS credentials of enter incorrect Virtual I/O Server credentials.

  The user interface displays a message about the failure of the operation. When you try to apply incorrect Virtual I/O Server credentials (), the Virtual I/O Server changes to "authentication failed" state. The managed system instance changes from "pending VIOS" state to "out-of-date" due to connection problems.

- You add a managed system instance without Virtual I/O Servers.

  The managed system instance appears in the instances table in "ready" state.

- You add a P-server to a managed system.

  The new P-server and VIOS are discovered automatically. If VIOS credentials match the default VIOS credentials, no configuration is required. If the VIOS credentials do not match the default VIOS credentials, set the VIOS credentials on AIM instances ().

- You remove a Virtual I/O Server in "invalid configuration" state from a managed system.

  The LPAR AIM removes the corresponding record from the instance table and the managed system changes to 'ready' state.

- You remove a Virtual I/O Server in "ready" state from a managed system.

  The LPAR AIM removes the corresponding record from the instance table.

- You remove a managed system instance with one or two Virtual I/O Servers.

  The managed system instance and associated Virtual I/O Server entries disappear from the instances table.

- The IBM PowerVM administration pane displays status information through icons and tooltips.

  Detailed tooltips become visible when you hover the cursor over warning and error icons.

  The following icons can appear:

   Discovery in progress

   No polling

   Error

   Warning

   Disabled

   Unknown

Add an HMC or an IVM Server Connection to the Manager

You can add an HMC or an IVM Server connection using the Administration tab of the Virtual Assurance user interface. Also, you can use the following authentication methods while adding HMC or an IVM Server connection to the Manager.

- Password Authentication: Requires only password for authentication
- Certificate Authentication: Requires public and private keys for authentication

To use certificate authentication, complete the following steps before you add an HMC or an IVM Server connection to the Manager.

1. Configure HMC Server for Certificate Authentication
2. Configure VAIM server for Certificate authentication

Configure HMC Server for Certificate Authentication

1. Generate a pair of public and private keys without passphrase.
2. Convert the certificates to OpenSSH fomat.
3. Add the public key certificate to ‘authorized_keys2’ file of the HMC server as follows :
   a. Login to HMC as ‘hmcsuperadmin’ user
   b. Execute the following command :

      mkauthkeys --add "ssh-rsa AAAAB3Nza….”
   c. To verify if the certificate was added successfully, use ‘cat’ command as follows : cat .ssh/authorized_keys2

Configure VAIM server for Certificate authentication

Place the public and private key certificates that are converted to OpenSSH format in the following locations:

- VAIM server machine:

  <VAIM Install Directory>\CA\VirtualAssurance\bin
- AIM server machine:

  <SystemEDGE Install Directory>\plugins\calpara

Note: Copy the above files in both locations even if VAIM and AIM servers are located in the same machine.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.

2. Select IBM PowerVM from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed HMC and IVM Servers, associated Virtual I/O Servers, and the LPAR AIM Servers.

3. Click  (Add) on the HMC/IVM Servers pane toolbar.

   The New HMC/IVM Server dialog appears.

4. Enter the required connection data (server name, user, password), specify the preferred AIM, enable Managed Status (checkbox).

5. Select the required authentication method and complete the required information.

   Note: The preferred AIM field is active only if you specify more than one AIM instance for a given HMC or IVM server.

6. (Optional) Specify the Virtual I/O Servers Default Credentials.

   The default VIOS credentials apply to newly discovered VIO servers.

   Important! If you do not specify the default VIOS credentials for an HMC Server, provide the VIOS credential for each VIOS in the Virtual I/O Servers panel to complete the configuration of the discovered VIOS. If the default VIOS credentials do not apply to a particular VIOS, you can overwrite the credentials in the Virtual I/O Servers panel.

7. Click OK.

   If the network connection has been established successfully, the Server is added to the top right HMC/IVM Servers pane with a green status icon. CA Virtual Assurance discovers the HMC/IVM Server automatically.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Manager Connection to the Server Fails

Symptom:

After I have added a server connection under Administration, Configuration, the validation of the connection to the server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used server connection data is still valid. If necessary, update the connection data.
- Verify, if the server system is running and accessible.
- Verify, if all services that are required for the connection are running properly on the server system.

To update the server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 server cannot be established, continue with the next procedure.

To verify if the server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. To find out whether the server has a valid DNS entry and IP address, verify the output of these commands.

   If the server is not in the DNS, add the server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>

   Enter the correct IP address and server name and save the file. For example:

   `192.168.50.50 myServer`
4. Change to the CA Virtual Assurance user interface, Administration tab, Configuration, Server pane, and click  (Validate) in the upper-right corner.

   If the server credentials and connection data are correct and you can ping the server, the connection can still fail. In this case, it is possible that the server causes the problem. If the connection to the server cannot be established, continue with the next procedure.

To verify, if all services that are required for the connection are running properly on the server system:

1. To access the server, contact the system administrator.
2. Log in to the server system.
3. Verify, if all services that are required for the connection are running properly.
4. If necessary, start or restart the service.
5. Change to the CA Virtual Assurance user interface, server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the server connection.

   If the connection to the server fails, verify the validity of the data you gathered according to the requirements for this scenario.

   Work with the administrator or support to fix the server connection problem.

Add the LPAR AIM Instance

After adding an HMC or an IVM Server connection to the CA Virtual Assurance manager, add an AIM instance to manage the new Server. CA Virtual Assurance then discovers the PowerVM environment.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.

2. Select IBM PowerVM from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed HMC and IVM Servers, associated Virtual I/O Servers, and the LPAR AIM Servers.

3. Click  (Add) on the LPAR AIM Servers pane toolbar.

   The New LPAR AIM Server dialog appears.

4. Select the LPAR AIM Server from the drop-down list.

   The list of discovered LPAR AIM Servers appears. If you have installed the LPAR AIM on the local system, the name of the local system appears in the list too.

5. Select the HMC or IVM Server from the drop-down list.

   CA Virtual Assurance populates the HMC/IVM Server drop-down list with the HMC and IVM Servers listed in the HMC/IVM Servers pane. You can only manage those HMC or IVM Servers for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

   A new AIM instance for the selected Server is added. If the instance is not in an error or in a stopped state, CA Virtual Assurance starts to discover the associated PowerVM environment:

   - For each HMC server, the AIM discovers all Power Systems and the Virtual I/O servers.
   - For each IVM server, the AIM discovers a Power System that the IVM manages.

   When the discovery process is complete, you can start managing your PowerVM environment.

The Administration Tab shows an aggregated state for all the P Systems and VIO Servers the AIM discovered. To view their individual configuration state press the Show Managed Systems  icon.

Change the Preferred HMC for the Managed Power System

If your Power System uses a dual HMC, you can change the preferred HMC.

Important! Verify that one LPAR AIM manages both the primary and the redundant HMC servers.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.

2. Select IBM PowerVM from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed HMC and IVM Servers, associated Virtual I/O Servers, and the LPAR AIM Servers.

3. Click  (Configure Managed/VIO Servers) that is associated with the HMC Server.

   The IBM PowerVM dialog with Managed/VIO Servers appears.

4. Click  (Switch the preferred HMC) under the Actions row and confirm.

   The redundant HMC is set as the preferred HMC.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the number of managed objects that are related to virtual and physical resources in your environment. You can move the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Virtual Assurance adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance 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 AIM:

- Verify that the AIM Server is accessible.
- Verify that SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping servername`

2. Verify that the output of the commands has a valid DNS entry and IP address for the AIM server.

   If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 servername`

   Enter the correct IP address and AIM server name. For example:

   `192.168.50.51 myAIM`

4. Click (Validate) in the upper-right corner of the 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 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, AIM Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the AIM Server connection.

   If the error status remains unchanged, verify that the data you gathered is according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Virtual Assurance discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Virtual Assurance manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Group in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand IBM PowerVM group.

   The managed HMC and IVM Servers appear.
3. Expand the HMC or IVM Server entry.

   The managed systems appear.

CA Virtual Assurance is now ready to manage the added PowerVM environment with its virtual infrastructure.

calpara.xml File Overview

The main purpose of the calpara.xml file is to store LPAR AIM configuration data, such as: persistent data and default values. The settings for monitoring can be adjusted to specific environments.

This document is intended for system administrators, who are familiar with xml format. We recommended caution when changing this file. To change the calpara.xml file, stop SystemEDGE before and start SystemEDGE again after changing the file.

Important! If you require Monitoring Threshold, Lag, or Severity adjustment, change only the Default Values. Modify Poll Group and DisableOutOfDate setting only when CA Support asks you to do so.

The calpara.xml file is at the following location:

```
<SystemEDGE_InstallDir>\plugins\calpara\calpara.xml
```

Persistent Data

Persistent data is available the next time the AIM starts up. This data can change during the AIM lifetime and is user settable, by SNMP set requests.

The following list presents the examples of persistent data:

- Instances
- Systems
- Partitions
- Slots
- Poll Groups

Instances

For each configured Instance in the Instance Table (lparAimInstanceTable) a section similar to the following example is stored:

```
<ManagedInstance>
    <InstIndex>7</InstIndex>
    <SerialNr>1010101</SerialNr>
    <ServerName>vios1.company.com</ServerName>
    <ServerType>vios</ServerType>
    <RowStatus>1</RowStatus>
</ManagedInstance>
```

ServerType

Specifies one of the following server types: hmc, vios, or ivm.

RowStatus

Specifies the status as active (1) or notInService (2).

Systems

For each managed Power System in the System Table (lparAimStatSysTable) a section similar to the following example is stored inside the related ManagedInstance section:

```
<System>
    <MonitorIndices>530091,530092,530093,530094,530095</MonitorIndices>
</System>
```

MonitorIndices

Stores the indices of the SystemEDGE monitors that the AIM created to monitor the operational status, CPU, and Memory usage of the Power System.

Note: If the AIM does not manage the Power System any more, the corresponding monitors are deleted.

Partitions

For each managed Logical Partition in the Partition Table (lparAimStatLPTable) four corresponding entries similar to the following examples are stored in the Partitions section inside the related ManagedInstance section:

```
<Partitions>
    ...
    <LparIndex>7</LparIndex>
    <LparId>7</LparId>
    <LparName>LPAR12345</LparName>
    <MonitorIndices>530141,530142,530143,530144,530145</MonitorIndices>
    ...
</Partitions>
```

Note: If the AIM does not manage the Power System any more, the corresponding monitors are deleted.

Slots

For each Physical Slot in the Slot Table (lparAimStatSlotTable) four corresponding entries similar to the following examples are stored in the Slots section inside the related ManagedInstance section:

```
<Slots>
    ...
    <SlotIndex>3</SlotIndex>
    <DRCName>U787B.001.DNWFF77-P1-C3</DRCName>
    <DRCIndex>553713666</DRCIndex>
    <SlotName>C3</SlotName>
    ...
</Slots>
```

The LPAR AIM uses this data to detect any Slot-related change right after start-up, potentially causing to send a corresponding SNMP trap.

Poll Groups

A Poll Group is a set of related commands, all executed at the same poll interval. For each Poll Group in the Poll Table (lparAimPollTable) three corresponding entries in the respective Poll Group section are stored. The following example shows the Basic Poll Group:

```
<Basic>
    <PollDefault>5</PollDefault>
    <PollSpecific>30,30</PollSpecific>
    <PollInstances>4,6</PollInstances>
</Basic>
```

PollDefault

Stores the default poll interval (in minutes) to apply to all instances, except those instances listed in PollInstances (listing the indices of the instances).

PollSpecific

Stores a list of poll intervals (in minutes) to apply one-to-one to the corresponding list of instances stored in PollInstances.

Note: Initially PollSpecific and PollInstances are empty.

Default Values

The following section describes default values stored in calpara.xml file that specify lags, thresholds, and severities, the AIM uses when creating new SystemEDGE monitors.

Important! Default values cannot change during the AIM lifetime and are not user settable.

```
<LowestPollInterval>5</LowestPollInterval>
<DisableOutOfDate>0</DisableOutOfDate>
<MonitorIndexStart>530001</MonitorIndexStart>
<SysAliveSev>fatal</SysAliveSev>
<CpuLagValue>3</CpuLagValue>
<CpuThresh1Val>95</CpuThresh1Val>
<CpuThresh1Sev>warning</CpuThresh1Sev>
<CpuThresh2Val>98</CpuThresh2Val>
<CpuThresh2Sev>critical</CpuThresh2Sev>
<MemLagValue>2</MemLagValue>
```

```
<MemThresh1Val>95</MemThresh1Val>
<MemThresh1Sev>warning</MemThresh1Sev>
<MemThresh2Val>98</MemThresh2Val>
<MemThresh2Sev>critical</MemThresh2Sev>
```

LowestPollInterval

Stores the lowest allowed poll interval (in minutes).

DisableOutOfDate

Specifies whether the data status (lparAimInstDataStatus) outOfDate is excluded.

Note: Set the variable to 1 to disable the data status becoming outOfDate(7) if any command execution failed.

Default: 0

MonitorIndexStart

Specifies the index of the first SystemEDGE monitor the AIM creates after startup.

Note: When creating new monitors, the AIM always searches for the next free index being equal or greater.

SysAliveSev

Specifies the severity of the SystemEDGE monitor that the AIM creates for monitoring the operational status of a Power System or a Logical Partition.

Valid values: ok, warning, minor, major, critical, fatal.

Note: Changing this value has no effect on existing monitors.

CpuThresh1Val and CpuThresh2Val

Specify the threshold values of the two SystemEDGE monitors that the AIM creates for monitoring the CPU usage of a Power System or a Logical Partition.

Limits: 0 to 100.

Note: Changing this value has no effect on existing monitors.

CpuThresh1Sev and CpuThresh2Sev

Specify the severities of the two SystemEDGE monitors that the AIM creates for monitoring the CPU usage of a Power System or a Logical Partition.

Valid values: ok, warning, minor, major, critical, fatal.

Note: Changing this value has no effect on existing monitors.

MemThresh1Val and MemThresh2Val

Specify the threshold values of the two SystemEDGE monitors the AIM creates for monitoring the Memory usage of a Power System or a Logical Partition.

Limits: 0 to 100.

Note: Changing this value has no effect on existing monitors.

MemThresh1Sev and MemThresh2Sev

Specify the severities of the two SystemEDGE monitors that the AIM creates for monitoring the Memory usage of a Power System or a Logical Partition.

Valid values: ok, warning, minor, major, critical, fatal.

Note: Changing this value has no effect on existing monitors.

CpuLagValue and MemLagValue

Specify the lag values of the SystemEDGE monitors that the AIM creates for monitoring the CPU and Memory usage of a Power System or a Logical Partition. The lag value specifies the number of consecutive poll intervals (Basic Poll Group) that the monitor reaches the threshold before the monitor changes its status.

Note: Changing this value has no effect on existing monitors.

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://System:SerialNu mber/Total	Alive	critical	ok

mon Index	monOID	mon ObjClass	monObjInstance	mon ObjAttribute	mon Severity	mon CurrState
530002	lparAimStatSys CPUUsage PerMil.1	CPU	lpar://System:SerialNumber/Total	PercentUsed	warning	ok
530003	lparAimStatSys CPUUsage PerMil.1	CPU	lpar://System:SerialNumber/Total	PercentUsed	minor	ok
530004	lparAimStatSys MemoryUsage PerMil.1	Memory	lpar://System:SerialNumber/Total	PercentUsed	warning	warning
530005	lparAimStatSys MemoryUsage PerMil.1	Memory	lpar://System:SerialNumber/Total	PercentUsed	minor	minor
530006	lparAimStatLP Status.1.1	System	lpar://System:SerialNumber/lpar01/Total	Alive	critical	critical
530007	lparAimStatLPCPU Usage.1.1	CPU	lpar://System:SerialNumber/lpar01/Total	PercentUsed	warning	ok
530008	lparAimStatLPCPU Usage.1.1	CPU	lpar://System:SerialNumber/lpar01/Total	PercentUsed	minor	ok
530009	lparAimStatLP MemoryUsage.1.1	Memory	lpar://System:SerialNumber/lpar01/Total	PercentUsed	warning	ok
530010	lparAimStatLP MemoryUsage.1.1	Memory	lpar://System:SerialNumber/lpar01/Total	PercentUsed	minor	ok
530011	lparAimStatLP Status.1.2	System	lpar://System:SerialNumber/lpar02/Total	Alive	critical	critical
530012	lparAimStatLPCPU Usage.1.2	CPU	lpar://System:SerialNumber/lpar02/Total	PercentUsed	warning	ok
530013	lparAimStatLPCPU Usage.1.2	CPU	lpar://System:SerialNumber/lpar02/Total	PercentUsed	minor	ok
530014	lparAimStatLP Memory Usage.1.2	Memory	lpar://System:SerialNumber/lpar02/Total	PercentUsed	warning	ok

IBM PowerVM Management

This section describes the IBM PowerVM management operations that you can perform from the Resources page.

The Resources page lets you view events and perform management operations on LPARs. Expand the IBM PowerVM group in the Explore pane to list the following objects:

- HMC/IVM servers
- PowerVM systems
- Logical partitions (LPARs)

View Resource Summary and Events

CA Virtual Assurance displays the Summary in the right-hand pane. The Summary page provides resource properties at the following levels in the object hierarchy:

- PowerVM Server
- LPAR

Performance Chart pane displays the utilization with available metrics and options. Use appropriate filter settings to display the required performance charts:

- CPU
- Memory
- Other metrics

General Information pane includes the following properties:

- Name, Item Type, Type (pSeries)
- Quantity characteristics of CPU and memory
- Number of LPARs and available processing units
- Serial Number

Overview pane displays information about:

- CPU state
- Memory state
- Operating state
- Health state
- Propagated Health state
- Collection Engine state

The Summary tab lets you view information associated with that object for example, total memory, operating system, number of CPUs, IP address, overall CPU and memory usage and events associated with the resource. Click the Configuration tab on the Usage panel to configure threshold limits.

Control Power Status for Logical Partitions

You can control the status of logical partition by performing one of the following operations:

- Activate
- Restart
- Shutdown
- Delete

You can perform any of these operations on one or multiple logical partitions simultaneously.

Follow these steps:

1. Select the managed machine on which you want to perform a status operation in the Explore pane.
2. Right-click the partition, select Management. You can also click Quick Start and click the related link of power control. Select *one* of the following:

   Activate

   Activates the selected logical partition that is currently powered off or suspended.

   Restart

   Shuts down the guest operating system and restarts it.

   Shutdown

   Shuts down the selected logical partition. You can only shut down a logical partition that is currently powered on.

   Delete

   Deletes the selected logical partition permanently. You can only delete a logical partition if it is shut down.

   A confirmation dialog appears.
3. Click OK.

   The status operation occurs, and a confirmation message appears. Refresh the interface to view the new logical partition status. An event appears confirming the result of the operation.

Activate Logical Partition

Activate a Logical Partition to commit resources to the partition and start the installed operating system. You can only activate a partition when it is not running.

To activate a Logical Partition

1. Right-click a partition on the Explore pane and select Management, Activate.

   The Activate Logical Partition dialog appears.

2. Complete the following fields and click OK.

   Profile

   Specifies the partition profile used to activate the partition.

   Keylock

   Specifies the key lock position. Key Lock establishes the power-on and power-off modes allowed for the system. CA Virtual Assurance supports the following valid keylock modes:

   Do Not Override

   The LPAR uses the keylock mode specified in the selected profile.

   Normal

   The LPAR starts up as normal. Use this option to perform most everyday tasks.

   Manual

   Consider security impact, when you set the key lock position to Manual.

   Boot mode

   Specifies the boot mode. Select a boot mode and select the Activate check box only if you want to use a boot mode that is different from the one specified in the selected profile. The system uses this boot mode to start the operating system on the logical partition unless you specify otherwise when activating the partition profile. CA Virtual Assurance supports the following valid boot modes:

   Do Not Override

   The LPAR uses the boot mode specified in the selected profile.

   Normal

   The LPAR starts up as normal. Use this option to perform most everyday tasks.

   Open Firmware

   The LPAR boots to the open firmware prompt. This option is used by service personnel to obtain additional debug information.

3. Click the Events tab for the partition.

An event should appear confirming the result of the operation.

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 Group 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 Processors page appears.
4. Specify whether to allocate partial processor units or dedicated processors and the minimum, desired, and maximum processors units. 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 adapter, 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 Disks 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.

    Note: If a physical fibre channel port that supports NPIV is selected, a virtual fibre channel server adapter and virtual fibre channel client adapter are created for the partition.

13. Click Next.

    The Summary page appears.

14. Verify the Summary and click Add Computer.

    The logical partition is created.

Delete Logical Partition

You can delete a partition from the managed system that is no longer required. When you delete a logical partition, all hardware resources are returned to the primary partition. You can only delete a partition that is powered off.

To delete a Logical Partition

1. Right-click a partition in the Explore pane and select Management, Delete.

   A confirmation dialog appears.

2. Click OK.

   A message appears confirming that the request was submitted.

3. Click the Summary tab for the partition.

   An event should appear confirming the result of the operation. The deletion is unsuccessful if the partition is not powered off. If the deletion was successful, the partition disappears from the Explore pane after you refresh the interface.

Restart Logical Partition

You can restart a partition that is already running. Restarting a partition shuts it down and starts the operating system again.

Note: A Logical Partition must be in the Running or Open Firmware state to restart.

To restart a Logical Partition

1. Right-click a partition on the Explore pane and select Management, Restart.

   The Restart Logical Partition dialog appears.

2. Select one of the following restart types using the Type drop-down list and click OK:

   Immediate

   Shuts down the logical partition immediately. The HMC/IVM ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

   Operating System Shutdown

   Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

   Operating System Shutdown Immediate

   Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

3. Click the Summary tab for the partition.

   An event should appear confirming the result of the operation.

Shut Down Logical Partition

Shutting down a partition shuts down the operating system. A partition must be in the Running or Open Firmware state to shut down.

To shut down a Logical Partition

1. Right-click a partition on the Explore pane and select Management, Shut Down.

   The Shut Down Logical Partition page appears.

2. Select one of the following shutdown types using the Type drop-down list and click OK:

   Immediate

   Shuts down the logical partition immediately. The HMC/IVM ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

   Operating System Shutdown

   Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

   Operating System Shutdown Immediate

   Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

3. Click the Summary tab for the partition.

   An event should appear confirming the result of the operation.

Configuring CPU and Memory

You can configure memory shares allocated to a virtual machine to adjust its allocated resources. When you add resources, the appropriate amount of unassigned memory or CPU shares must be available for the operation to succeed. If values exist for the minimum and maximum allowed memory or CPU shares, any resource allocation change must stay within these limits. You can edit VM CPU and memory allocation using the Quick Start link on the Resources tab. You can also use create and schedule policy with specific VM resource allocation actions.

Important! For Dynamic LPAR operations, like adding or removing CPU and Memory, install AIX version 5.2 or 5.3 or 6.0 or higher on each LPAR system. Alternatively, run AIX resource control daemon IBM.DRM on the LPAR system.

Configure CPU

To configure VM CPU allocation

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Configuration, Configure Processor...

   The Configure Processor Resource Allocation dialog appears.

3. Select *one* of the following Adjustment Types:

   Dynamic Adjustment

   Updates the running VM.

   Profile Update

   Updates the active profile. The VM must be restarted to pick up the changes from the profile.

   Dynamic Adjustment and Update Profile

   Updates both the running VM and the active profile.

4. Edit the corresponding fields and click Ok.

   A confirmation message appears.

Configure Memory

To configure VM Memory allocation

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Configuration, Configure Memory...

   The Configure Memory Resource Allocation dialog appears.

3. Select one of the following Adjustment Types:

   Dynamic Adjustment

   Updates the running VM.

   Profile Update

   Updates the active profile. The VM must be restarted to pick up the changes from the profile.

   Dynamic Adjustment and Update Profile

   Updates both the running VM and the active profile.

4. Edit the corresponding fields and click Ok.

   A confirmation message appears.

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.

How to Configure Hyper-V Management

The following diagram provides an overview of the required actions. The diagram includes troubleshooting strategies for connection problems.

How to Configure the Hyper-V Server Management Components

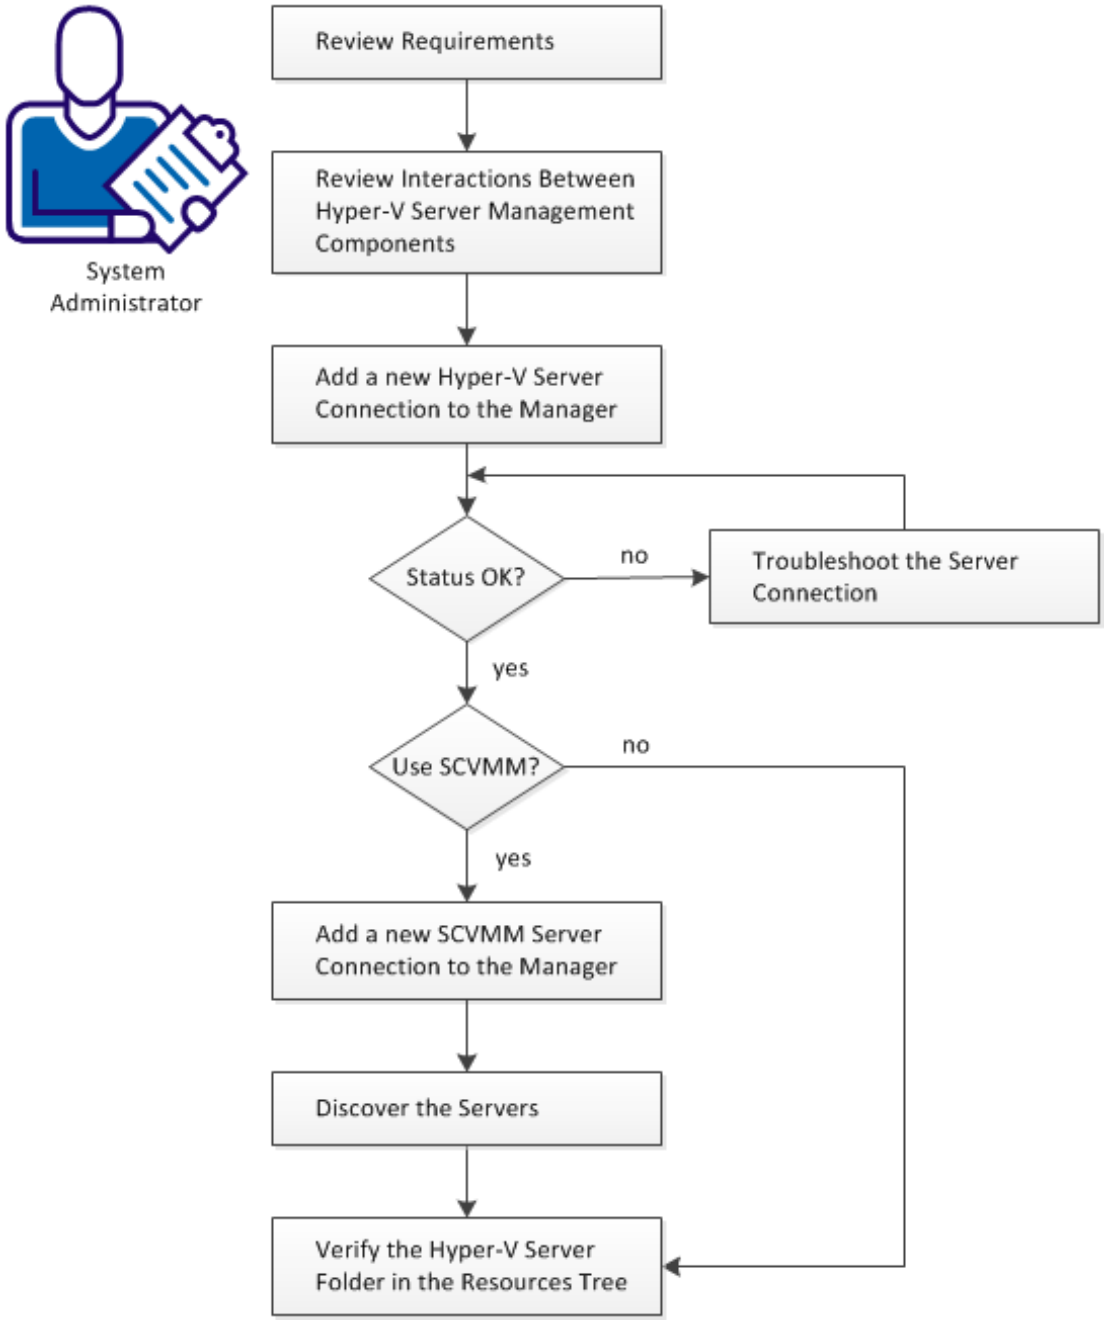

Review Hyper-V Requirements

Review the following requirements before you start configuring the Hyper-V management components of CA Virtual Assurance:

- 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 Hyper-V Servers.
- You can access a CA Virtual Assurance manager installation that includes the Hyper-V Platform Management Module (PMM), Hyper-V Application Insight Module (AIM), and Monitoring Agent (CA SystemEDGE).
- You can access the CA Virtual Assurance user interface.
- Verify that the Hyper-V AIM is installed at the Hyper-V Server.
- You have valid credentials available (user name and password) to access the Hyper-V Server that you want to manage.
- You have verified that the Hyper-V Server runs properly.
- Verify that the SNMP settings on the CA Virtual Assurance manager and the Hyper-V Server 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 the Hyper-V Server that you want to use.

More information:

Interactions Between Hyper-V Server Management Components (see page 397)
Apply Required Settings for Using Microsoft Hyper-V (see page 396)
Add a New Hyper-V Server Connection to the Manager (see page 398)
Discover the Servers (see page 405)
Verify the Hyper-V Server Folder in the Resources Tree (see page 406)
(Optional) Add the SCVMM Management Instance to the CA Virtual Assurance Manager (see page 401)

Apply Required Settings for Using Microsoft Hyper-V

Verify prerequisites and apply the following settings for Microsoft Hyper-V management:

To apply required settings for using Microsoft Hyper-V

1. Verify that SystemEDGE, Advanced Encryption, and the Hyper-V AIM are installed on the Hyper-V Server. You can only assign one AIM for each managed Hyper-V Server.
2. Disable the local User Access Control (UAC) on the Hyper-V Server.
3. Disable the network UAC by setting the following registry value:
   `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\LocalAccountTokenFilterPolicy,1 (REG_DWORD)`
4. Enable remote WMI firewall exception by running the following command from the command prompt:

   ```
   netsh advfirewall firewall set rule group="Windows Management Instrumentation (WMI)" new enable=yes
   ```

5. Verify that the user configured in the management components for server access is a member of the "Distributed COM Users" group.

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

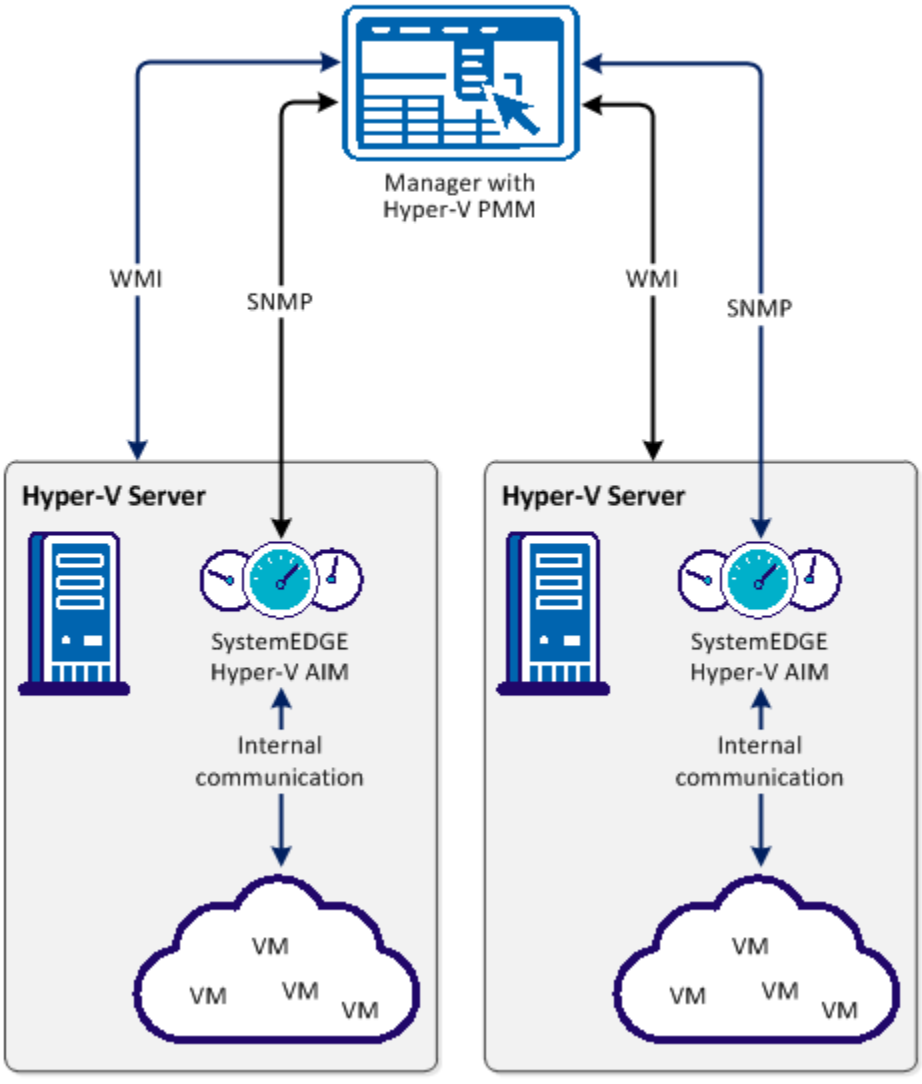

You can configure Hyper-V management by adding connection information. Use *one* of the following methods:

- Administration tab of the user interface
- NodeCfgUtil.exe utility on the AIM Server

More information:

Add a New Hyper-V Server Connection to the Manager (see page 398)
Discover the Servers (see page 405)
Verify the Hyper-V Server Folder in the Resources Tree (see page 406)
(Optional) Add the SCVMM Management Instance to the CA Virtual Assurance Manager (see page 401)

Add a New Hyper-V Server Connection to the Manager

You can add a Hyper-V 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 Hyper-V Server from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed Hyper-V Servers.

3. Click  (Add) on the Hyper-V Servers pane toolbar.

   The New Hyper-V Server dialog appears.

4. Enter the required connection data (server name, user, password) and click OK.

   If the network connection has been established successfully, the Hyper-V Server is added to the top right Hyper-V Servers pane with a green status icon. CA Virtual Assurance discovers the Hyper-V Server automatically.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the Hyper-V 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 Hyper-V Server Connection (see page 399).

More information:

Discover the Servers (see page 405)
Hyper-V Server Connection Failed (see page 399)
Verify the Hyper-V Server Folder in the Resources Tree (see page 406)
(Optional) Add the SCVMM Management Instance to the CA Virtual Assurance Manager (see page 401)

Hyper-V Server Connection Failed

Symptom:

After I have added a new Hyper-V Server connection under Administration, Configuration, the validation of the connection to the Hyper-V Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Hyper-V Server connection data (server name, user, password) is still valid. If necessary, update the connection data.
- Verify, if the Hyper-V Server system is running and accessible.

To update the Hyper-V Server connection data

1. Click (Add) or (Edit) that is associated with the failed connection.

   The New or Edit Hyper-V Server dialog appears.

2. Add the valid server name, user, and password. 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 Hyper-V Server cannot be established, continue with the next procedure.

To verify, if the Hyper-V Server system is running and accessible

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Hyper-V Server Name>`
   `ping <IP Address of Hyper-V Server>`

2. Verify the output of the commands to find out whether the Hyper-V Server has a valid DNS entry and IP address.

   If the Hyper-V Server is not in the DNS, add the Hyper-V Server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the Hyper-V 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 <Hyper-V Server Name>`

   Enter the correct IP address and Hyper-V Server name. For example:

   `192.168.50.50 myHyper-V`

4. Click  (Validate) in the upper-right corner.

   If the connection to the Hyper-V Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

   Work with the Hyper-V administrator or VMware support to fix the Hyper-V Server connection problem.

(Optional) Add the SCVMM Management Instance to the CA Virtual Assurance Manager

The following diagram provides an overview about the required actions. The diagram includes corresponding troubleshooting strategies in case of connection problems.

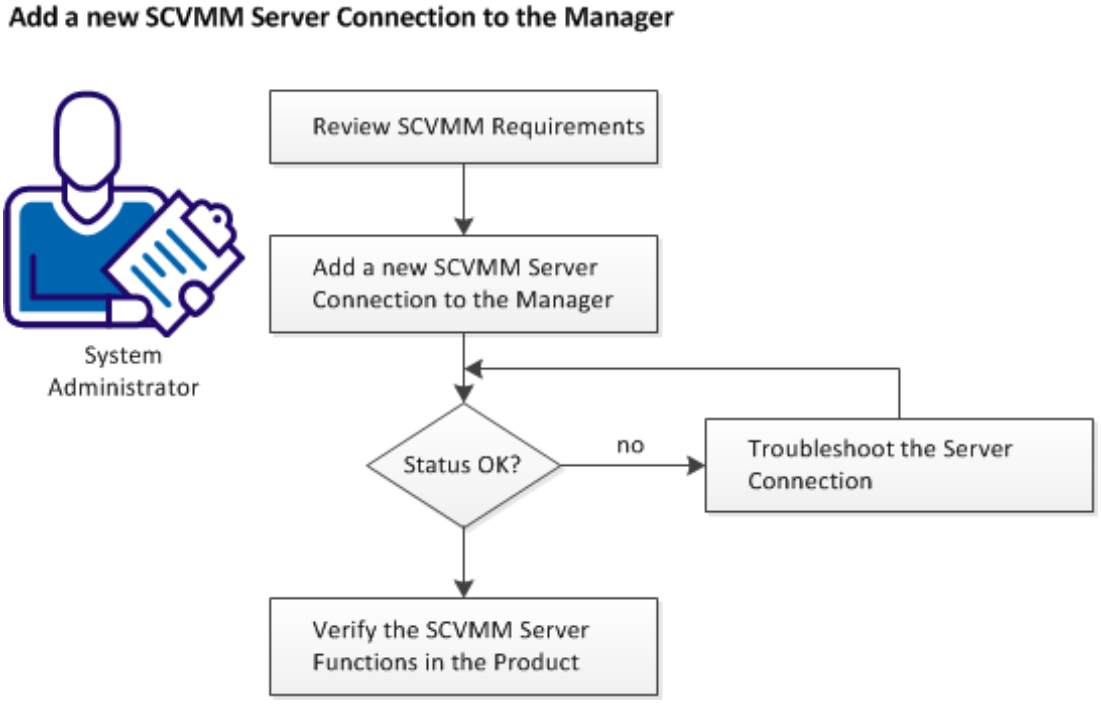

Follow these steps:

Apply Required Settings for Using Microsoft SCVMM (see page 402)
Add a New SCVMM Server Connection to the Manager (see page 403)
SCVMM Server Connection Failed (see page 404)

Apply Required Settings for Using Microsoft SCVMM

CA Virtual Assurance optionally integrates with Microsoft System Center Virtual Machine Manager (SCVMM) for Hyper-V provisioning. SCVMM is not required for Hyper-V monitoring and management. Instead of SCVMM, you can also use local templates (Hyper-V Server bound) for VM provisioning.

When using the optional SCVMM integration, verify the following prerequisites and apply the required settings:

To apply required settings for using Microsoft SCVMM

1. Verify that the SCVMM Server, all potential Hyper-V target hosts for VM provisioning, and the CA Virtual Assurance manager running the Hyper-V PMM are members of the same Active Directory domain.
2. Verify that Hyper-V target hosts for VM provisioning are configured in CA Virtual Assurance and SCVMM. CA Virtual Assurance does not perform a SCVMM discovery.
3. Verify that SCVMM has Windows Remote Management (WinRM) configured.
4. Run the following command on the command line of the SCVMM server to configure WinRM:

   ```
   winrm quickconfig
   ```

5. Allow unencrypted HTTP or enable HTTPS in addition to the basic WinRM configuration on the SCVMM server.

   Run the following command to allow unencrypted HTTP traffic:

   ```
   winrm set winrm/config/service @{AllowUnencrypted="true"}
   ```

   To enable HTTPS, obtain an SSL certificate for your SCVMM server, install it, and run the following command:

   ```
   winrm quickconfig -transport:https
   ```

Depending on the expected utilization of the SCVMM server in your environment, VM provisioning can fail due to inadequate parameter settings for the Quota Management for Remote Shells on the SCVMM server. Affected parameters are:

MaxShellsPerUser

Specifies the maximum number of shells per user.

Default: 5

MaxConcurrentUsers

Specifies the maximum number of concurrent users who can open shells.

Default: 5

If you expect more than one provisioning job at a time, you can increase the parameter settings on the SCVMM server as follows:

```
winrm set winrm/config/winrs @{MaxShellsPerUser="number"}
winrm set winrm/config/winrs @{MaxConcurrentUsers="number"}
```

Example

```
winrm set winrm/config/winrs @{MaxShellsPerUser="30"}
winrm set winrm/config/winrs @{MaxConcurrentUsers="10"}
```

See also the Quota Management for Remote Shells article from Microsoft.

Add a New SCVMM Server Connection to the Manager

You can add a SCVMM 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 SCVMM Server from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed SCVMM Servers.

3. Click  (Add) on the SCVMM Servers pane toolbar.

   The New SCVMM Server dialog appears.

4. Enter the required connection data (server name, user, and password) and click OK.

   If the network connection has been established successfully, the SCVMM Server is added to the top right SCVMM Servers pane with a green status icon. CA Virtual Assurance discovers the SCVMM Server automatically.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the SCVMM Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

SCVMM Server Connection Failed

Symptom:

After I have added a new SCVMM Server connection under Administration, Configuration, the validation of the connection to the SCVMM Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used SCVMM Server connection data (server name, user, password) is still valid. If necessary, update the connection data.
- Verify, if the SCVMM Server system is running and accessible.

To update the SCVMM Server connection data

1. Click (Add) or (Edit) that is associated with the failed connection.

   The New or Edit SCVMM Server dialog appears.

2. Add the valid server name, user, and password. 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 SCVMM Server cannot be established, continue with the next procedure.

To verify, if the SCVMM Server system is running and accessible

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   ```
   nslookup <SCVMM Server Name>
   ping <IP Address of SCVMM Server>
   ```

2. Verify the output of the commands to find out whether the SCVMM Server has a valid DNS entry and IP address.

   If the SCVMM Server is not in the DNS, add the SCVMM Server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the SCVMM 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 <SCVMM Server Name>

   Enter the correct IP address and SCVMM Server name. For example:

   `192.168.50.50 mySCVMM`

4. Click  (Validate) in the upper-right corner.

   If the connection to the SCVMM Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

   Work with the SCVMM administrator or Microsoft support to fix the SCVMM Server connection problem.

Discover the Servers

After adding a new Hyper-V Server connection and an optional SCVMM connection to the CA Virtual Assurance manager, discover the Hyper-V Server and the SCVMM Server. CA Virtual Assurance then discovers the entire Hyper-V environment with all its virtual components.

Verify, that the SNMP credentials on the CA Virtual Assurance manager and on the Hyper-V and SCVMM Servers are consistent. If necessary, update the SNMP configurations accordingly.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Resources, Data Center.

   The data Center page appears.

2. Click Discover System from Quick start in the right pane.

   The right pane refreshes and displays the Discover System wizard.

3. Enter the required data and click Finish.

   CA Virtual Assurance discovers the system.

More information:

Verify the Hyper-V Server Folder in the Resources Tree (see page 406)

Verify the Hyper-V Server Folder in the Resources Tree

After a successful configuration and discovery, the new Hyper-v Server is listed in the Resources Explore pane under the Hyper-V Server folder.

Follow these steps:

1. Click Resources, Explore.

   The resources tree appears.

2. Expand Microsoft Hyper-V Server.

   The managed SCVMM Servers appear.

3. Expand the new SCVMM Server entry.

   The managed Hyper-V server infrastructure appears.

CA Virtual Assurance is now ready to manage the added Hyper-V environment with its virtual infrastructure.

Chapter 7: Hyper-V Management

Hyper-V Server lets you manage your Hyper-V servers and virtual machines. The Hyper-V Server is the central location from which you can view all virtual resources and perform management operations such as start, stop, delete, and so on.

This section describes the management operations that you can perform on Hyper-V Server resources from the Resources page. The Resources page lets you view basic information and details about the following objects:

- Hyper-V Servers
- Virtual Machines

Click Resources, open the Explore pane, and select one of the Hyper-V resources; then click Summary for the resource.

The Summary page lets you view information associated with that object (for example, Hyper-V Server, or virtual machines on a Hyper-V server) and events associated with the resource.

The Details page lets you view other detailed resource information, such as system properties, software, hardware, performance, and so on.

Right-click menus on the Explore pane let you perform management and policy tasks.

This section contains the following topics:

More Information

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..

More Information

Provision Machine: Microsoft Hyper-V (see page 674)

Manage VM Status (Hyper-V)

You can control the status of Hyper-V Server virtual machines by performing one of the following VM operations:

- Start
- Stop
- Pause
- Restart
- Shut Down
- Save

Note: The available options differ based on the current state of the machine, for example when the if machine is powered off state , Turn off VM option is not available.

You can perform any of these operations on multiple VMs simultaneously.

To control VM status

1. Select the virtual machine on which you want to perform a status operation in the Explore pane.

Right-click the VM, select Management. You can also click Quick Start and click the related link of power control. Select one of the following:

Start

Starts the virtual machine and boots the guest operating system. You can only power on a virtual machine that is currently powered off or suspended.

Stop

Powers off the virtual machine. You can only power off a virtual machine that is currently powered on or suspended.

Pause

Pauses the virtual machine and saves its current state. All activity is suspended until you resume the machine.

Restart

Shuts down the guest operating system and restarts it.

Shutdown

Shuts down the guest operating system. You can only shut down a virtual machine that is currently powered on.

Save

Saves the current status of the virtual machine. In other platforms, this option is similar to Suspend.

A confirmation dialog appears.

1. Click OK.

   The status operation occurs, and a confirmation message appears. Refresh the interface to view the new VM status. An event appears confirming the result of the operation.

Delete a Virtual Machine

When you delete a virtual machine from Hyper-V Server, the virtual machine is deleted from the virtual disk.

Note: You can only delete a VM that is in the power off state, otherwise the delete option is not available.

To delete a virtual machine

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Delete.

   The Delete Hyper-V VM dialog appears with options to delete additional components.

3. Click OK.

   A message appears confirming the request submission.

4. Click the Summary tab for the virtual machine.

   An event appears confirming the result of the operation. If successful, the virtual machine is deleted from the virtual disk, and the virtual machine disappears from the Explore pane after you refresh the interface.

Rename a Virtual Machine

You can rename an existing virtual machine from Hyper-V Server.

Note: The Rename VM option is available only when machine is in powered Off or Saved state.

To rename a virtual machine

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Rename.

   The Rename VM dialog appears.

3. Enter the New VM Name and Click OK.

   The message appears confirming the request submission.

4. Click the Summary tab for the virtual machine.

   An event appears confirming the result of the operation. If successful, the virtual machine is renamed after you refresh the interface.

Create Action and Rules

You can create actions and rules based on predetermined policies for different types of resources such as a virtual machine.

To create Actions and Rules for a virtual machine

1. Select Resources, Explore, Data Center.
2. Open the Policy tab and the Actions subtab.
3. Click + (Add) to create a new Action.
4. Select the appropriate items from the drop-down menus and follow the instructions in the user interface to complete the Action.
5. Select the Rules subtab and click + (Add) to create a new Rule.

   The Rule/Template Identification and Evaluation page dialog appears.

   A wizard guides you through the creation process. Assign an Action from the available Actions list to this rule.

   For more information about Actions and Rules, see Create an Action (see page 611) and Create a Rule.

Edit Startup and Shutdown Actions

You can edit the actions to start up and shut down a virtual machine.

To edit Startup and Shutdown Actions

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Configuration, Startup and Shutdown Actions.

   The Startup and Shutdown Actions dialog appears.

3. The Startup and Shutdown Actions dialog contains the following fields:

   Start Action

   Specifies the action to perform when the Hyper-V Server starts. Select one of the following from the drop-down list:

   - Always

     Always starts the VM when the Hyper-V Server starts.

   - Auto

     Automatically starts the VM when the Hyper-V Server starts, if the VM is shutdown in running mode.

   - None
     Does not start the VM when the Hyper-V Server starts.

   Start Delay

   Adjust the delay (in seconds) to start a VM after the Hyper-V Server starts.

   Shutdown Action

   Specifies the action to perform when the virtual machine shuts down. Select one of the following from the drop-down list:

   - Off

     Turns off the VM before Hyper-V Server shuts down.

   - Save

     Saves (Suspends) the VM before Hyper-V Server shuts down.

   - Shutdown

     Shuts down the VM before Hyper-V server shuts down.

   Recovery Action

   Specifies the action to regain the previous details of a virtual machine when the Hyper-V Server fails. Select one of the following from the drop-down list:

   - None

Does not take a specific action when the Hyper-V Server starts after the server fails.

- Restart

  Restarts the VM when Hyper-V Server starts after the VM server fails.

- Revert

  Reverts the VM with the latest snapshots when the Hyper-V Server starts after the server fails.

4. Click Ok after you edit all the details. A confirmation message appears.

Edit VM CPU and Memory Allocation

You can edit the number of CPU and memory shares allocated to a virtual machine to adjust its allocated resources. When you add resources, the appropriate amount of unassigned memory or CPU shares must be available for the operation to succeed. If values exist for the minimum and maximum allowed memory or CPU shares, any resource allocation change must stay within these limits.

You can also create and schedule policy with specific VM resource allocation actions.

To edit VM CPU and memory allocation

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Configuration, Resource Allocation, CPU and Memory.

   The CPU and Memory Resource Allocation dialog appears.

3. Adjust the number of CPUs, reserved CPU percentage, and CPU Limit percentage.

4. Adjust the memory shares allocated to the virtual machine and click Ok after you edit all the details.

   A confirmation message appears.

Hyper-V Management Actions

The following action types are available for use with Hyper-V Server:

- Delete Machine (see page 649)
- Change Machine State (see page 621)
- Configure Power (see page 634)
- Configure CPU/Memory (see page 625)

You can use these action types to create new actions that automate the configuration of startup and shutdown options for Hyper-V Server, and other operations when assigned rule criteria are met. You can also schedule these actions to occur at specific times.

For more information about using actions and rules to create automation policy, see the "Policy" chapter.

Red Hat Enterprise Virtualization

CA Virtual Assurance introduced kernel-based virtual machine support. The k*ernel-based virtual machine (KVM)* is a hardware-assisted virtualization infrastructure for the Linux kernel. The CA KVM AIM is implemented as a multi-instance, remote AIM. The CA KVM AIM enables RHEV monitoring. *Red Hat Enterprise Red Hat Enterprise Virtualization (RHEV)* is an enterprise virtualization product that is based on the KVM hypervisor.

How to Configure the Red Hat Enterprise Virtualization Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

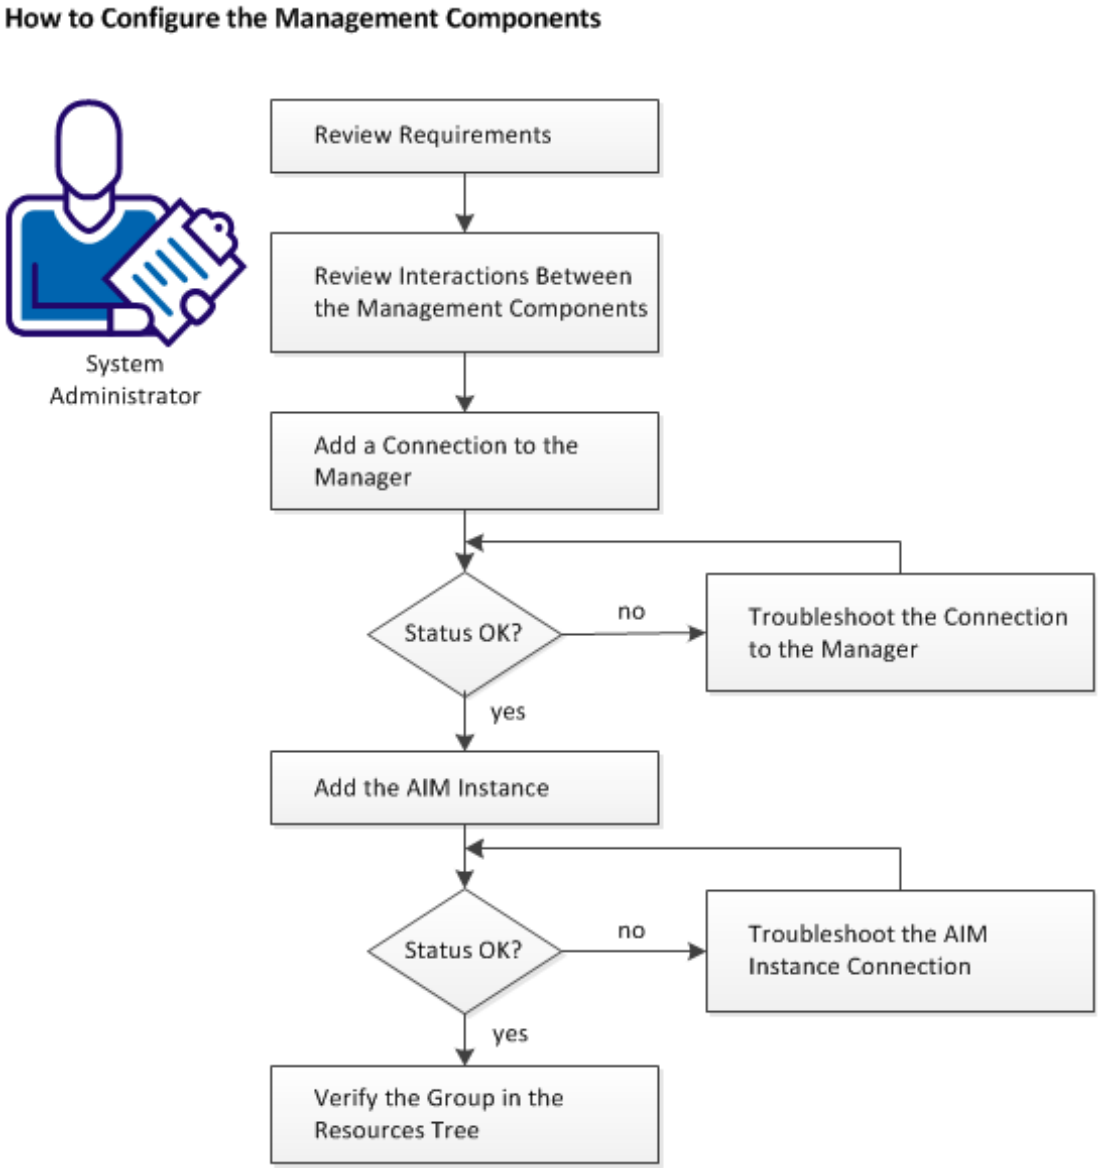

Follow these steps:

Review Requirements

Review the following requirements before configuring the management components of CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Virtual Assurance and SystemEDGE.
  - You can access a CA Virtual Assurance manager installation that includes:
  - Platform Management Module (PMM)
  - Application Insight Module (AIM)
  - Monitoring Agent (SystemEDGE)
- You can access the CA Virtual Assurance user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote AIM Servers that you want to use.

Interactions Between RHEV Management Components

The following diagram illustrates how the components involved in RHEV monitoring interact. SystemEDGE and the KVM AIM run on a Windows Server. The AIM communicates with one or more RHEV managers using REST API.

Interaction Between KVM Management Components

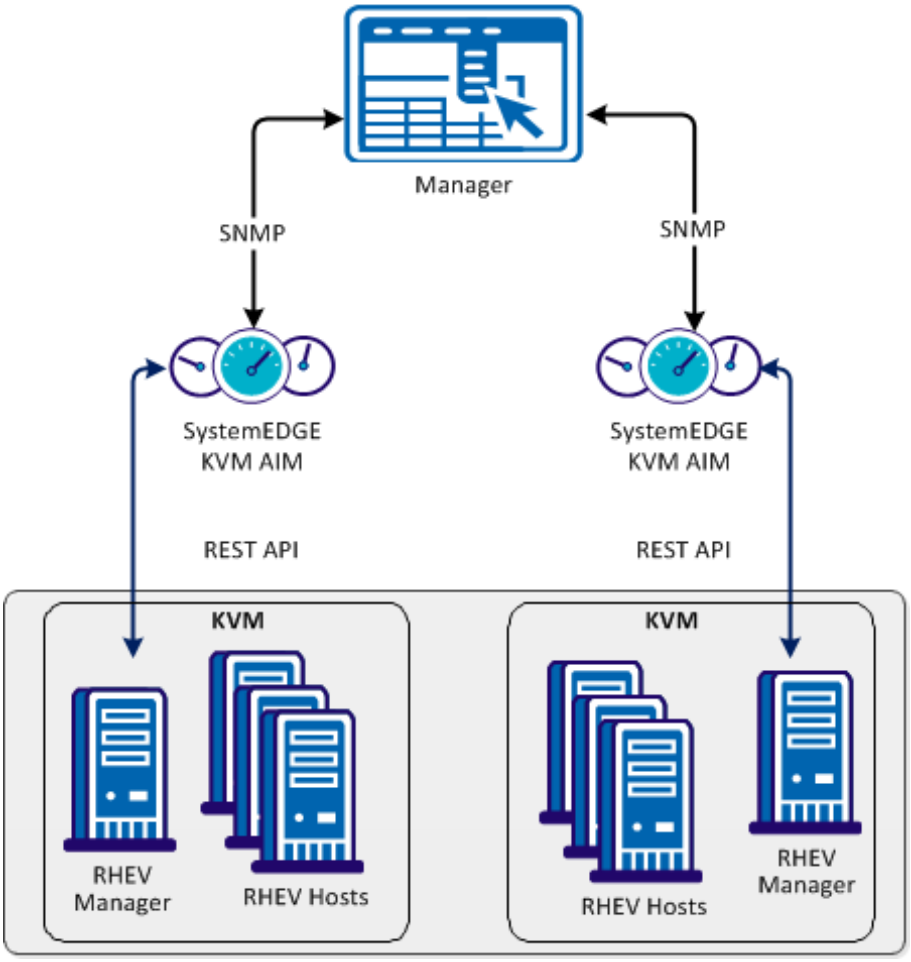

Add a Red Hat Enterprise Virtualization Connection to the Manager

You can add a Red Hat Enterprise Virtualization 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 Red Hat Enterprise Virtualization from the Provisioning section in the left pane.

3. Click  (Add) on the Registered Red Hat Enterprise Virtualization Servers pane toolbar.

   The Add Red Hat Enterprise Virtualization Server dialog appears.

4. Enter the required connection data (server name, user, password, ISO Library credentials, port), specify the preferred AIM, enable Managed Status (checkbox).

   Note: The ISO Library contains ISO images for provisioning. Without the ISO image, the provisioning does not work.

5. Click OK.

   If the network connection has been established successfully, the Server is added to the top right pane with a green status icon.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Manager Connection to the Server Fails

Symptom:

After I have added a server connection under Administration, Configuration, the validation of the connection to the server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used server connection data is still valid. If necessary, update the connection data.
- Verify, if the server system is running and accessible.
- Verify, if all services that are required for the connection are running properly on the server system.

To update the server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 server cannot be established, continue with the next procedure.

To verify if the server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. To find out whether the server has a valid DNS entry and IP address, verify the output of these commands.

   If the server is not in the DNS, add the server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>`

   Enter the correct IP address and server name and save the file. For example:

   `192.168.50.50 myServer`
4. Change to the CA Virtual Assurance user interface, Administration tab, Configuration, Server pane, and click  (Validate) in the upper-right corner.

   If the server credentials and connection data are correct and you can ping the server, the connection can still fail. In this case, it is possible that the server causes the problem. If the connection to the server cannot be established, continue with the next procedure.

To verify, if all services that are required for the connection are running properly on the server system:

1. To access the server, contact the system administrator.
2. Log in to the server system.
3. Verify, if all services that are required for the connection are running properly.
4. If necessary, start or restart the service.
5. Change to the CA Virtual Assurance user interface, server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the server connection.

   If the connection to the server fails, verify the validity of the data you gathered according to the requirements for this scenario.

   Work with the administrator or support to fix the server connection problem.

Add the Discovered Red Hat Enterprise Virtualization AIM Instance

After adding a Red Hat Enterprise Virtualization connection to the CA Virtual Assurance manager, add an AIM instance to manage the Red Hat Enterprise Virtualization environment.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.
2. Select Red Hat Enterprise Virtualization from the Provisioning section in the left pane.

3. Click  (Add) on the Discovered Red Hat Enterprise Virtualization AIM Instances pane toolbar.

   The Add Red Hat Enterprise Virtualization AIM Instance appears.

4. Select the RHEV AIM Server from the drop-down list.

   The list of discovered RHEV AIM Servers appears.

5. Select the RHEV Server from the drop-down list.

   CA Virtual Assurance populates the RHEV Server drop-down list with the RHEV Servers listed in the Registered Red Hat Enterprise Virtualization pane. You can only manage those RHEV Servers for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

   A new AIM instance for the selected Server is added. If the instance is not in an error or in a stopped state, CA Virtual Assurance starts to discover the associated environment. When the discovery process is complete, you can start managing your Red Hat Enterprise Virtualization environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the number of managed objects that are related to virtual and physical resources in your environment. You can move the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Virtual Assurance adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance 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 AIM:

- Verify that the AIM Server is accessible.
- Verify that SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping servername`

2. Verify that the output of the commands has a valid DNS entry and IP address for the AIM server.

   If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 servername`

   Enter the correct IP address and AIM server name. For example:

   `192.168.50.51 myAIM`

4. Click  (Validate) in the upper-right corner of the 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 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, AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

   CA Virtual Assurance validates the AIM Server connection.

   If the error status remains unchanged, verify that the data you gathered is according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Virtual Assurance discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Virtual Assurance manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Red Hat Enterprise Virtualization Group in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand Red Hat Enterprise Virtualization group.

   The managed Red Hat Enterprise Virtualization resources appear.

CA Virtual Assurance is now ready to manage the Red Hat Enterprise Virtualization environment that was configured.

How to Prepare Linux template for KVM Provisioning

CA Virtual Assurance supports customized provisioning of new virtual machines (VM) running the following operating systems:

- Red Hat Enterprise Server 6.0
- SUSE Linux Enterprise Server 11

Customization options include hostname, password, domain, or network configuration.

The following diagram illustrates how a system administrator prepares Linux template for VM provisioning.

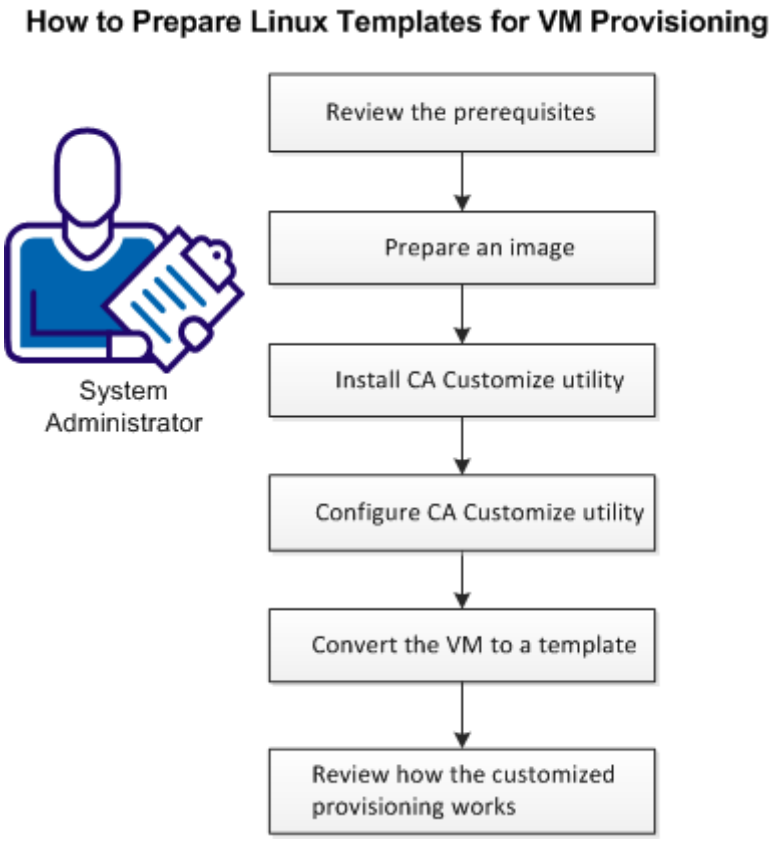

Follow these steps:

Prerequisites for Customized VM Provisioning (see page 427)
Prepare a Linux Image (KVM) (see page 427)
Install CA Customize Utility (see page 428)
Configure CA Customize Utility (see page 429)
Convert the VM to a Template (see page 429)
How the Customized Provisioning Works (see page 430)

Prerequisites for Customized VM Provisioning

To customize the Linux guest, one needs direct access to the file system or console.

Ensure that the following prerequisites are met for the RHEV environment:

- Each RHEV data center uses a local ISO library on the RHEV manager system.
- Each machine has SFTP access enabled.
- The RHEV manager has SSH access enabled.

Prepare a Linux Image (KVM)

Before you create a template containing the Linux operating system, prepare the image by following this procedure. The specific steps may differ based on the Linux Distribution.

Follow these steps:

1. Install the Linux operating system on a new virtual machine from scratch.
2. Install RHEV Guest Tools inside the virtual machine.
3. Apply any customizations like user accounts, policy, applications, hotfixes, that you would like to apply on the new virtual machines.

This Linux image is ready for further customization using the CA Customize utility.

Install CA Customize Utility

CA Customize utility enables CA Virtual Assurance to change the virtual machine settings externally. The guest utility monitors CD drive on the OS start. If a special ISO is connected, the following actions are executed:

1. A set of commands customizes the guest.
2. The guest system is marked as customized.
   The system cannot be modified again until someone resets this state.
3. The system is halted to indicate that the customization succeeded.

To install correct ca-customize guest utility:

1. Find this utility at:
   - Valid for Red Hat Enterprise Server 6.0

     `<InstallationRoot>\Utilities\linuxCustomization\rh6`
   - Valid for SUSE Linux Enterprise Server 11

     `<InstallationRoot>\Utilities\linuxCustomization\sles11`
2. Transfer this executable file to the following location on a hard drive of the VM being prepared:

   `/usr/bin/ca-customize`
3. (Optional) Provide your own version of ca-customize script to support other guest systems that we do not support.
4. Enable executable bit of the ca-customize utility:

   `chmod 755 /usr/bin/ca-customize`

Configure CA Customize Utility

You can set up the template for Linux provisioning. To customize the guest, use the available scripts. You can also use your own scripts to allow further setup.

Follow these steps:

1. Disable the network interfaces so that the network does not affect the customization process.

   Note: The network is enabled automatically during the customization.

2. Override the default CDROM device name if needed using the */etc/ca-customize.conf* file.

   CD_DEVICE=/dev/cdrom

   Defines the device name that is used for CD drive.

   Default: /dev/cdrom

3. Set up the automatic start at the end of the boot process.

   - (Valid for SUSE Linux) Create or modify the */etc/init.d/after.local* file:

     ```
     #!/bin/bash
     [ -e /etc/ca-customized ] || /usr/bin/ca-customize
     ```

   - (Valid for Red Hat Linux) Add the following line to the */etc/rc.local* file:

     ```
     [ -e /etc/ca-customized ] || /usr/bin/ca-customize
     ```

4. Shut down the system.

Convert the VM to a Template

The template allows you to create any number of customized virtual machines.

Follow these steps:

1. Shut down the VM.
2. To convert the shutdown virtual machine to an RHEV template, use the RHEV Administration portal.

The template appears in CA Virtual Assurance and can be used for customized provisioning.

Once these steps have been performed, the new template can be used to create any number of new customized virtual machines.

How the Customized Provisioning Works

The following steps represent the Customized VM Provisioning Workflow.

1. The platform management service provisions new Linux VM.
2. The platform management service prepares new ISO using customization parameters and attach it to new VM.
3. The platform management service starts the VM.
4. The VM detects that customization ISO is attached. The VM applies the customization changes.
5. If the customization is successful, the VM shuts down. The PMM detects that the VM is stopped. The platform management service starts VM again and finishes provisioning.
6. If the customization failed, the VM is not halted. The platform management service takes the following actions:
    a. Returns a provisioning failure
    b. Sets the provisioning job in exception state

Customization Log

A successful customization is stored in the */etc/ca-customized* file. This file contains a list of the customization changes.

If the customization fails, the logs are stored in the */etc/ca-customized.tmp* file.

How to Prepare Windows Templates for KVM Provisioning

CA Virtual Assurance supports customized provisioning of new virtual machines (VM) running Windows 2003 R2 Server (32 bit and 64 bit), Windows 2008 (32 bit and 64 bit) or Windows 2008 R2 Server (64 bit). Customization options include a number of settings. For example, changing the built-in Administrator account password, computer name, and the network configuration.

The following diagram illustrates how a system administrator prepares Windows templates for KVM provisioning.

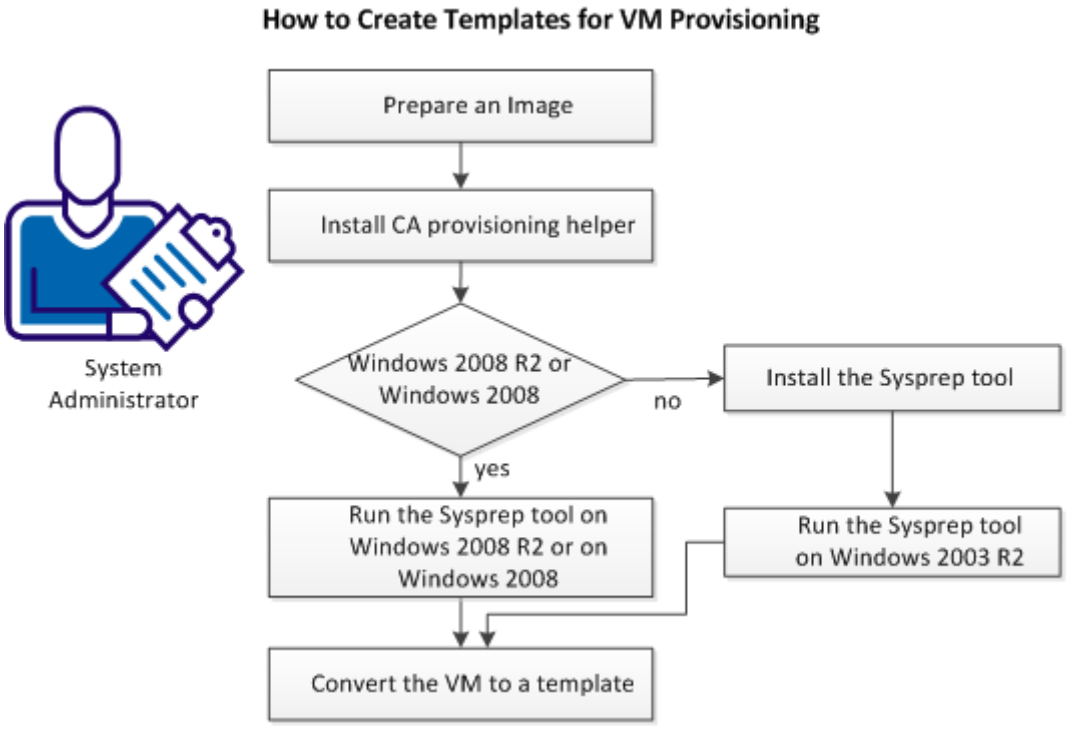

Follow these steps:

1. Prepare an image. (see page 432)
2. Install CA provisioning helper. (see page 321)
3. (Valid on Windows 2003 R2) Install the Sysprep tool. (see page 322)
4. Depending on your operating system select *one* of the following actions:
   - Run Sysprep tool on Windows 2003 R2. (see page 322)
   - Run Sysprep tool on Windows 2008 R2. (see page 322)
5. Convert the VM to a template. (see page 433)

Prerequisites for RHEV Environments

Ensure that the following prerequisites are met for the RHEV environment:

- Each RHEV data center uses a local ISO library on the RHEV manager system.
- Each machine has SFTP access enabled.
- The RHEV manager has SSH access enabled.

Prepare a Windows Image

When creating a template containing the Windows operating system, prepare the image by following this procedure. Follow the steps to enable CA Virtual Assurance provisioning operations to customize the template. The specific steps differ based on the Windows version.

Follow these steps:

1. Install the Windows operating system on a new virtual machine from scratch.
2. Install RHEV Guest Tools inside the virtual machine.
3. Apply any customizations like user accounts, policy, applications, hotfixes, and so on, that you would like to apply on the new virtual machines.
4. (Valid on Windows 2003) Blank out the built-in Administrator account password.

   Note: If the Administrator password is not empty, SysPrep is unable to set a new password during provisioning and the existing password remains.

Install CA provisioning helper

CA provisioning helper enables CA Virtual Assurance to change the virtual machine settings externally.

Follow these steps:

1. Find this utility at <InstallationRoot>\Utilities\Sysprep\CAProvisioningHelper.exe
2. Transfer this executable file to any location on a hard drive of the VM being prepared.
3. Execute CA provisioning helper once from the command line.

Install the Sysprep Tool

Install the Sysprep tool from the Windows installation CD-ROM.

The Sysprep Tool

The Microsoft provided Sysprep tools to generalize, freeze and shut down the readily configured Windows installation. The following sections describe how to use the Sysprep tool for Windows 2003 R2 and Windows 2008 R2 in detail.

Run the Sysprep Tool on Windows 2003 R2

After you configure the Sysprep tool installation, run the Sysprep tool.

Follow these steps:

1. Locate and open the following CAB file:

   `\SUPPORT\TOOLS\DEPLOY.CAB`

2. Select all files contained in the CAB file and copy them to the following location: %SystemDrive%\Sysprep (normally C:\Sysprep).

   Note: Do not change the directory name.

3. Change to the Sysprep directory and run:

   ```
   sysprep -quiet -reseal -mini -forceshutdown
   ```

Run the Sysprep Tool on Windows 2008 R2

The regular Windows installation process installs all files to perform the SysPrep process. After you configure the Windows installation, perform the following steps:

1. Generate a valid XML response file using the Windows Automated Installation Kit (WAIK) for Windows Server 2008 R2. WAIK is available from the Microsoft Web site.

   Note: The way provisioning requires a dummy unattended response file, or it cannot shut down. The content of the response file is irrelevant, since the provisioning process replaces it, but the file must follow the SysPrep-specific XML schema.

2. Name the generated XML file “sysprep.xml” and place it into the Sysprep directory:

   %SystemRoot%\system32\sysprep

3. Run the following command:

   ```
   sysprep /generalize /oobe /shutdown /unattend:sysprep.xml
   ```

Convert the VM to a Template in RHEV

To convert the shutdown virtual machine to an RHEV template, use the RHEV Administration portal.

The template appears in CA Virtual Assurance and can be used for customized provisioning.

Once these steps have been performed, the new template can be used to create any number of new customized virtual machines.

Manage VM Status (KVM)

You can control the status of virtual machines by performing one of the following operations:

- Discover
  - Server
  - Network
- Start
- Suspend
- Shutdown
- Destroy

To control VM status:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Right-click a VM, select Management and one of the following options:

   Discover

   Discovers a server or network.

   Start

   Starts a VM on the specified RHEV host.

   Suspend

   Suspends a running VM on the specified RHEV host and saves its current state. All activity is suspended until you resume the VM.

   Shutdown

   Shuts down a running VM on the specified RHEV host.

   Destroy

   Removes a VM.

   A corresponding wizard appears.

3. Fill in the required information and proceed to the next step.

4. Submit.

   The status operation occurs, and a confirmation message appears. Refresh the interface to view the new VM status. An event appears confirming the result of the operation.

Provision a RHEV Virtual Machine

You can provision virtual machines by performing the following procedure. Ensure that you prepare a Windows template for VM provisioning.

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Right-click the Red Hat Enterprise Virtualization group, select Provisioning, Provision RHEV Virtual Machine.

   A provisioning wizard appears.

3. Fill in the required information:

   VM Name

   Defines the new VM name.

   Template

   Specifies the Windows provisioning template.

   Administrator Password

   Defines the administrator password for the new VM.

   Product Activation Key

   Defines the Windows 2003 Activation Key.

   Full Name

   Defines the full VM name.

4. (Optional) Fill in the additional information (Workgroup, Memory, CPUs, VM Host, Organization). If you want to use a static IP address, disable the DHCP and provide the IP address, mask, and default gateway.

   Note: The Memory and CPUs settings depend on the Windows provisioning template used.

5. Submit.

   The confirmation message appears.

6. Refresh the Jobs panel to view the progress.

   An event appears confirming the result of the operation.

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 ensures 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.

How to Configure the Solaris Zones Management Components

The Solaris Zone PMM provides health monitoring, management, and provisioning of Solaris Zone environments.

Follow these steps:

Review Requirements (see page 437)
Interaction Between Solaris Zones Management Components (see page 439)
Add a Solaris Zones Connection to the Manager (see page 440)
Manager Connection to the Server Fails (see page 440)
Add the Zones AIM Servers (see page 442)
Troubleshoot the AIM Instance Connection (see page 443)
Verify the Solaris Zones Group in the Resources Tree (see page 446)

Review Requirements

Review the following requirements before configuring the management components of CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Virtual Assurance and SystemEDGE.
  - You can access a CA Virtual Assurance manager installation that includes:
  - Platform Management Module (PMM)
  - Application Insight Module (AIM)
  - Monitoring Agent (SystemEDGE)
- You can access the CA Virtual Assurance user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote AIM Servers that you want to use.

More information:

Requirements for Solaris Zones Management (see page 438)

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 privileges to execute the following commands:
  - zlogin
  - zoneadm
  - zonecfg
- From the global zone, the user must have the permission to log in to individual Solaris Zones with zlogin and run the following commands:
  - uname -a
  - sar
  - prstat
  - netstat
- Using the user interface or the NodeCfgUtil.exe utility on the Managed Node where the Solaris Zones AIM resides, add this user name and the corresponding password to CA Virtual Assurance during configuration. You can specify the additional feature in the user name field when using NodeCfgUtil.exe utility and the syntax is as follows:

  `cassh://`*`ZoneHost`*`:`*`sshPort`*`?authMethod=[Password|PublicKey]&username=`*`nonRootUser`*`[&sudo][&sshPublicKeyFile=`*`publicKeyFileName`*`][&sshPrivateKeyFile=`*`privateKeyFileName`*`]`

  authMethod=[Password|Publickey]

  Specifies the type of authentication method. The Default authentication method is password.

  username

  Specifies the username to log in to the Zone host.

  sudo

  Specifies the AIM to run sudo after the user logs in as defined in the username parameter.

  sshPublicKeyFile=*publicKeyFileName* sshPrivateKeyFile=*privateKeyFileName*

  Specifies the AIM to use sudo before running any command on the Zone host.
- Create a resource pool from Explore, Management, Create Resource Pool to allocate resources to zones, projects, and applications on a Solaris Zones server. Assign it to a zone during zone creation.

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

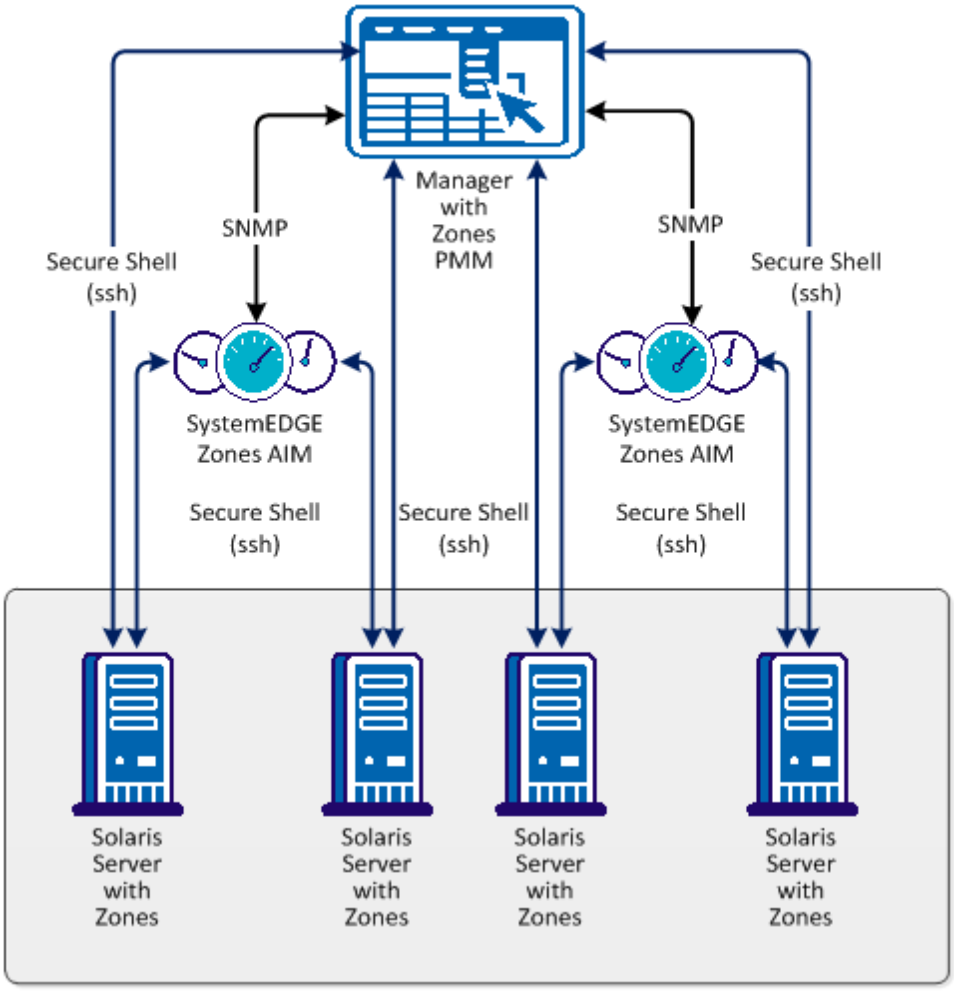

To add the required connection information for each Solaris Zones Server, use the Administration tab of the user interface or the NodeCfgUtil.exe utility on the managed node. The connection information is written to the configuration file on the managed node. The AIM polls the configuration file and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

Add a Solaris Zones Connection to the Manager

You can add a Solaris Zones 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 Solaris Zone from the Provisioning section in the left pane.

3. Click  (Add) on the Solaris Zone Servers pane toolbar.

   The Add Solaris Zone Server dialog appears.

4. Enter the required connection data (server name, user, password, port), specify the preferred AIM, enable Managed Status (checkbox).

5. Click OK.

   If the network connection is established successfully, the Server is added to the top right pane with a green status icon.

   Note: If the connection fails, the Validation Failed dialog appears. Click Yes, CA Virtual Assurance adds the Server to the list with a red status icon. If you click No, nothing is added.

Manager Connection to the Server Fails

Symptom:

After I have added a server connection under Administration, Configuration, the validation of the connection to the server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used server connection data is still valid. If necessary, update the connection data.
- Verify, if the server system is running and accessible.
- Verify, if all services that are required for the connection are running properly on the server system.

To update the server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 server cannot be established, continue with the next procedure.

To verify if the server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. To find out whether the server has a valid DNS entry and IP address, verify the output of these commands.

   If the server is not in the DNS, add the server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>`

   Enter the correct IP address and server name and save the file. For example:

   `192.168.50.50 myServer`
4. Change to the CA Virtual Assurance user interface, Administration tab, Configuration, Server pane, and click  (Validate) in the upper-right corner.

   If the server credentials and connection data are correct and you can ping the server, the connection can still fail. In this case, it is possible that the server causes the problem. If the connection to the server cannot be established, continue with the next procedure.

To verify, if all services that are required for the connection are running properly on the server system:

1. To access the server, contact the system administrator.
2. Log in to the server system.
3. Verify, if all services that are required for the connection are running properly.
4. If necessary, start or restart the service.
5. Change to the CA Virtual Assurance user interface, server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the server connection.

   If the connection to the server fails, verify the validity of the data you gathered according to the requirements for this scenario.

   Work with the administrator or support to fix the server connection problem.

Add the Zones AIM Servers

After adding a Solaris Zone connection to the CA Virtual Assurance manager, add an AIM instance to manage the Solaris Zone environment.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.
2. Select Solaris Zone from the Provisioning section in the left pane.
3. Click (Add) on the Zone AIM Servers pane toolbar.

   The New Zone AIM Server dialog appears.
4. Select the AIM Server from the drop-down list.

   CA Virtual Assurance populates the Instance drop-down list with the Zone Servers listed in the Registered Solaris Zones pane. You can only manage those Zone Servers for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first. After discovery, the AIM server appears in the drop-down list.
5. Select the Instance from the drop-down list and click OK.

   A new AIM instance for the selected Server is added.

   The AIM on the AIM Server is now configured to collect data from the specified Zone Server. If the instance is not in an error or in a stopped state, CA Virtual Assurance starts to discover the associated environment. When the Discovery process is complete, you can start managing the Solaris Zone environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the number of managed objects that are related to virtual and physical resources in your environment. You can move the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Virtual Assurance adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance 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 AIM:

- Verify that the AIM Server is accessible.
- Verify that SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping` *`servername`*

2. Verify that the output of the commands has a valid DNS entry and IP address for the AIM server.

   If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 servername

   Enter the correct IP address and AIM server name. For example:

   192.168.50.51 myAIM

4. Click (Validate) in the upper-right corner of the 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 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, AIM Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the AIM Server connection.

   If the error status remains unchanged, verify that the data you gathered is according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered AIM instances in the network, the status icons of several instances show (Disabled). This AIM instance is not managed.

This status appears, if CA Virtual Assurance discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Virtual Assurance manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Solaris Zones Group in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand Solaris Zone group.

   The managed Solaris Zone resources appear.

CA Virtual Assurance is now ready to manage the configured Solaris Zone environment.

Solaris Zones Management

Solaris Zones 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.

This section describes the management operations that you can perform on Solaris Zones resources from the Resources page. The Resources page lets you view basic information and details about the following objects:

- Solaris Zones servers
- Solaris Zones

Click Resources, open the Explore pane, and select one of the resources; then click Summary for the resource. The Summary page lets you view information associated with the resource (for example, zones, resource pools, and disks on a Solaris Zones server or network interfaces and projects on a zone) and events associated with the resource.

Note: If you select alert as Normal using the Configuration button in the Usage panel, the zone is shown in the normal state (in green) even if the CPU or memory is in the critical or warning state. Similarly, if you select alert as warning, the zone is always shows in the warning state.

The components tree displays only those resource pools that are used by a zone. Inactive resource pools are not listed in this panel.

The Details page lets you view other detailed resource information, such as system properties, software, hardware, performance, and so on.

Other pages may be available to perform resource management tasks. Right-click menus on the Explore pane also let you perform management and policy tasks.

More Information

Delete a Zone (see page 452)
Clone a Zone (see page 451)
Create Resource Pool (see page 449)
Control Zone Status (see page 450)
Available Solaris Zones Actions (see page 452)

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, Whole Root, 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.

Create Resource Pool

You can create a resource pool for use in allocating resources to zones, projects, and applications on a Solaris Zones server. After you create a resource pool, you can assign it to a zone during zone creation.

To create a resource pool

1. Right-click a Solaris Zones server on the Explore pane and select Management, Create Resource Pool.

   The Create Resource Pool dialog appears.

2. Complete the following fields and click OK:

   Name

   Identifies the resource pool name.

   Min CPU Shares

   Identifies the minimum number of CPU shares that the pool must have at any time.

   Max CPU Shares

   Identifies the maximum number of CPU shares that the pool can have.

   Processor Set Name

   Identifies the pool's processor set name.

   Scheduler Type

   Identifies the type of scheduling to use when allocating resources. Select FSS to use Fair Share Scheduling to allocate resources based on workload importance (the number of CPU shares specified for a project or task).

   The pool is created, and a confirmation message appears.

3. Click the Summary tab for the Zones server on which you created the pool, and select Resource Pools in the Show drop-down list to verify that the pool was created.

Control Zone Status

You can perform stop, reboot, start, and uninstall operations to control the status of a zone. You cannot perform these operations for global zones or when a zone is in the installed state.

To control zone status

1. Right-click a zone on the Explore pane, and select Management and one of the following options:

   Start

   Starts the zone, putting it into a running state. You can only start a zone that is in the installed state.

   Stop

   Halts the zone by resetting it to the installed state. Halting the zone stops all processes, removes network interfaces, and performs other operations to remove the zone's existing application environment and virtual platform. After halting a zone, you must start the zone to re-initiate the environment. You can only halt a zone that is currently running.

   Reboot

   Halts the zone and boots it again. You can only reboot a zone that is currently running. When you reboot a zone, the server gives it a new zone ID.

   Delete

   Deletes a zone. For more information, see the section Delete a Zone.

   Install

   Installs a native or branded zone which enters the configured state when the installation completes. Installing a zone opens a dialog that asks you for the archive path of a branded zone. If you install a native zone, leave this field empty. In case of a branded zone, provide the archive path.

   Note: You get an error message when you try to install a branded zone with no archive path parameter entered or a native zone with archive path parameter entered.

   Uninstall

   Uninstalls all the files under the zone's root file system. You can only uninstall a zone that is not currently running (installed state). You should uninstall a zone before you delete it.

   Clone

   Clones a zone. For more information, see the section Clone a Zone

   A confirmation dialog appears.

2. Click OK.

A message appears confirming that the request was submitted.

3. Click the Summary tab for the zone host.

   An event should appear confirming the result of the operation.

Note: The zone status shows incomplete, if the current operation is in progress and has not yet completed.

More Information

Delete a Zone (see page 452)
Clone a Zone (see page 451)

Clone a Zone

Cloning a zone lets you configure and install a new zone by copying the data from an existing zone. The zone that you are cloning must be halted for the Clone operation to be available. You cannot perform this operation for global zones or when a zone is in the configured or running state.

To clone a zone

1. Right-click a zone on the Explore pane and select Management, Clone.

   The Cloning pane appears.

2. Complete the following fields in the Target pane and click Clone.

   Name

   Specifies the name of the zone to which you want to copy the cloned information.

   Path

   Specifies the file path of the zone to which you want to copy the cloned information.

   A confirmation message appears.

3. Click the Summary tab for the zone host.

   An event confirming the result of the operation appears in the dashboard. The cloned zone appears in the Explore pane under its containing host when the operation completes.

Delete a Zone

You can delete a non-global zone from a Solaris Zones server. The zone must be shut down before you delete it.

If the zone is in the installed state, this operation will uninstall and then delete the zone. If the zone is in any other state such as running, an error message displays.

To delete a zone

1. Right-click a zone on the Explore pane and select Management, Delete.

   A confirmation dialog appears.

2. Click OK.

   A message appears confirming the deletion.

3. Click the Summary tab for the zone host.

   An event should appear confirming the result of the operation. The deleted zone should disappear from the Explore pane when the operation completes.

Available Solaris Zones Actions

The following action types are available for use with Solaris Zones:

- Clone Zone Machine (see page 622)
- Delete Zone Machine (see page 650)
- Provision Zone Machine (see page 677)

You can use these action types to create new actions that automate zone operations when assigned rule criteria are met. You can also schedule these actions to occur at specific times.

For more information about using actions and rules to create automation policy, see the chapter "Policy."

VMware vCloud

VMware vCloud Director lets you build secure, multitenant clouds by pooling virtual infrastructure resources into virtual data centers and exposing them to users. CA Virtual Assurance supports VMware vCloud Director management.
vCloud Director resources depend on underlying vSphere resources such as CPU, memory, storage, or vNetwork Distributed Switches to run virtual machines. You can use these underlying vSphere resources to create virtual machines and vApps in vCloud.

A *vCloud Organization* is a unit of administration that represents a collection of users, groups, and computing resources. Associated virtual datacenters provide the required computing resources. After users authenticate at the organization level, they can create, use, and manage virtual machines or vApps.

A *virtual datacenter (vDC)* provides virtual computing resources to a vCloud organization. You can provision, run, and store virtual systems in a virtual datacenter. A vCloud organization can have multiple virtual datacenters.

Organizations provide *catalogs* to store vApp templates and media files. The members of an organization can use the vApp templates and media files in the catalog to create their own vApps.

How to Configure the vCloud Director Management Components

Follow these steps:

Interactions Between vCloud Management Components (see page 455)
Review vCloud Requirements (see page 453)
Add a vCloud Director Connection to the Manager (see page 456)
Troubleshoot the vCloud Server Connection (see page 457)
vCloud Server Connection Failed (see page 458)
Add the AIM Instance for the vCloud Server (see page 460)
Troubleshoot the vCloud AIM Instance Connection (see page 461)
Verify the VMware vCloud Folder in the Resources Tree (see page 465)

Review vCloud Requirements

Review the following requirements before you start configuring the vCloud Director management components of CA Virtual Assurance:

- 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, VMware vSphere, and VMware vCloud.
- You can access a CA Virtual Assurance manager installation that includes the VMware Platform Management Module (PMM), vCloud 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 vCloud Director server that you want to manage.
- You have found out which protocol (HTTP or HTTPS) and port to use for accessing the vCloud Director through web services. Default: HTTPS, Port 443

- You have verified that the vSphere environment and the vCloud Director run properly.
- If the VMware PMM and vCloud 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 vCloud AIM Servers that you want to use.

More information:

Add a vCloud Director Connection to the Manager (see page 456)
Add the AIM Instance for the vCloud Server (see page 460)
Verify the VMware vCloud Folder in the Resources Tree (see page 465)

Interactions Between vCloud Management Components

The following diagram illustrates how the components involved in vCloud Director management interact. SystemEDGE and the vCloud AIM run on a Windows server. The AIM communicates with one or more remote vCloud Director servers to manage the virtual environment. The vCloud AIM collects the data for an entire view of the virtual resources associated with the vCloud Director. The underlying vSphere environment provides the required resources to run virtual machines and vApps.

Interaction Between vCloud Director Management Components

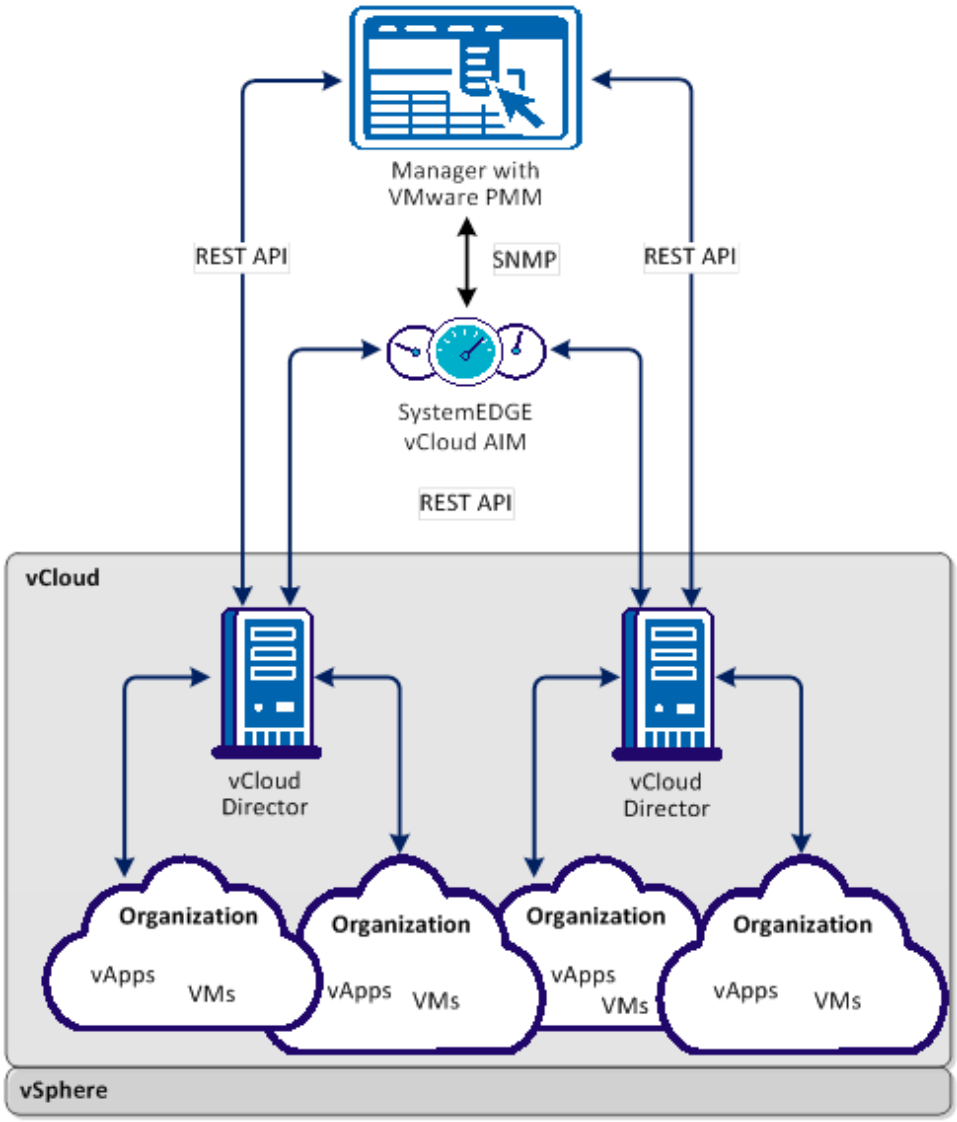

You can configure vCloud management through the Administration tab of the user interface.

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.

Add a vCloud Director Connection to the Manager

You can add a vCloud Director 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 vCloud Server from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed vCloud Servers, associated vCloud AIM Servers, and the AIM Instance for the vCloud Server.

3. Click  (Add) on the vCloud Servers pane toolbar.

   The Add vCloud Server dialog appears.

4. Enter the required connection data (server name, username, password, protocol, port), specify the preferred AIM, enable Managed Status (checkbox), and click OK.

   When specifying the username, you can use the following syntax to consider user roles and access levels:

   - System Administrator (Full access): administrator@System
   - Limit operation at the organization level and the role assignment (Organizational Access) *username@organization_name*

   If the network connection has been established successfully, the vCloud Server is added to the top right vCloud Servers pane with a green status icon. CA Virtual Assurance discovers the vCloud Server automatically.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the vCloud 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 vCloud Server Connection (see page 457).

More information:

Troubleshoot the vCloud Server Connection (see page 457)
Add the AIM Instance for the vCloud Server (see page 460)
Verify the VMware vCloud Folder in the Resources Tree (see page 465)

Troubleshoot the vCloud Server Connection

The vCloud Server connection has failed. Follow the troubleshooting information indicated in the following diagram:

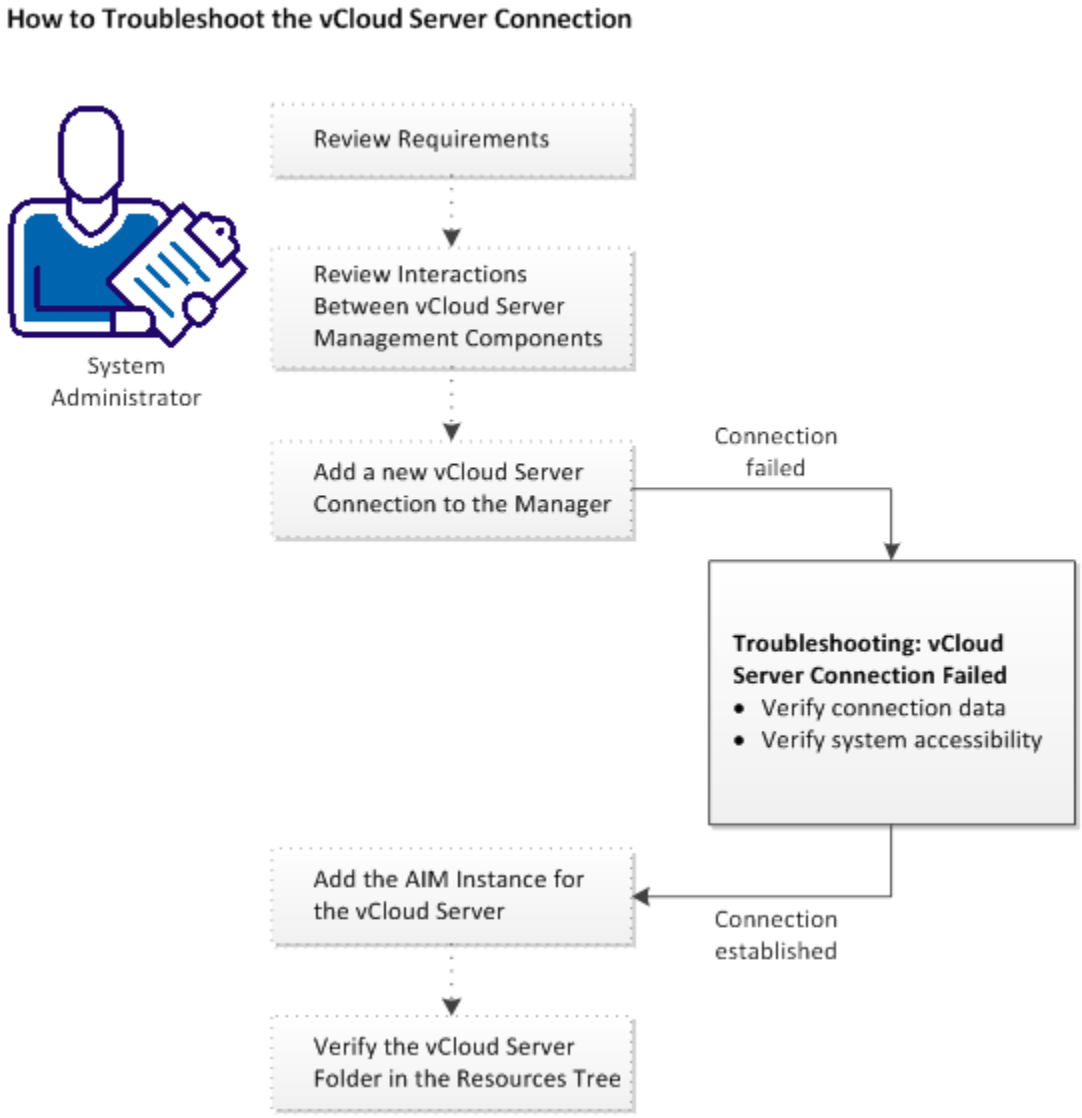

Follow these steps:

vCloud Server Connection Failed (see page 458)
Add the AIM Instance for the vCloud Server (see page 460)
Verify the VMware vCloud Folder in the Resources Tree (see page 465)

vCloud Server Connection Failed

Symptom:

After I have added a new vCloud Server connection under Administration, Configuration, the validation of the connection to the vCloud Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used vCloud Server connection data (server name, user, password, protocol, port) is still valid. If necessary, update the connection data.
- Verify, if the vCloud Server system is running and accessible.

To update the vCloud Server connection data

1. Click (Add) or (Edit) that is associated with the failed connection.

   The New or Edit vCloud Server dialog appears.

2. Add the valid server name, user, password, protocol, and port. Specify the preferred AIM. 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 vCloud Server cannot be established, continue with the next procedure.

To verify, if the vCloud Server system is running and accessible

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   ```
   nslookup <vCloud Server Name>
   ping <IP Address of vCloud Server>
   ```

2. Verify the output of the commands to find out whether the vCloud Server has a valid DNS entry and IP address.

   If the vCloud Server is not in the DNS, add the vCloud Server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the vCloud 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 <vCloud Server Name>

   Enter the correct IP address and vCloud Server name. For example:

   `192.168.50.50 myvCloud`

4. Click  (Validate) in the upper-right corner.

   If the connection to the vCloud Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

   Work with the vCloud administrator or VMware support to fix the vCloud Server connection problem.

Add the AIM Instance for the vCloud Server

After adding a new vCloud Server connection to the CA Virtual Assurance manager, add a vCloud AIM instance to manage the new vCloud Server. CA Virtual Assurance then discovers the entire vCloud environment with all its virtual components, such as Organizations, vApps, 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 vCloud Server from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed vCloud Servers, associated vCloud AIM Servers, and the AIM Instances for managed vCloud Servers.

3. Click  (Add) on the vCloud AIM Servers pane toolbar.

   The New vCloud AIM Server dialog appears.

4. Open the vCloud AIM Server drop-down list.

   The list of discovered vCloud AIM Servers appears. If you have installed the vCloud AIM on the local system, the name of the local system appears in the list too.

5. Select a vCloud AIM Server from the drop-down list.

   CA Virtual Assurance populates the vCloud Server drop-down list with the vCloud Servers listed in the vCloud Servers pane. That is, you can only manage those vCloud Servers for which your CA Virtual Assurance manager has a valid connection established.

6. Select the vCloud Server you want to manage and click OK.

   A new AIM instance for the selected vCloud Server is added. If the instance is not in an error or stopped state, CA Virtual Assurance starts to discover the associated vCloud environment. When the discovery process is complete, you can start managing the virtual resources of vCloud.

More information:

Troubleshoot the vCloud AIM Instance Connection (see page 461)
Verify the VMware vCloud Folder in the Resources Tree (see page 465)

Troubleshoot the vCloud AIM Instance Connection

The vCloud 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.

Error - Unable to connect to the AIM. Check the network configuration.

No Polling - The CA Virtual Assurance manager does not poll this AIM instance.

Disabled - This instance is not managed.

Follow the troubleshooting information indicated in the following diagram:

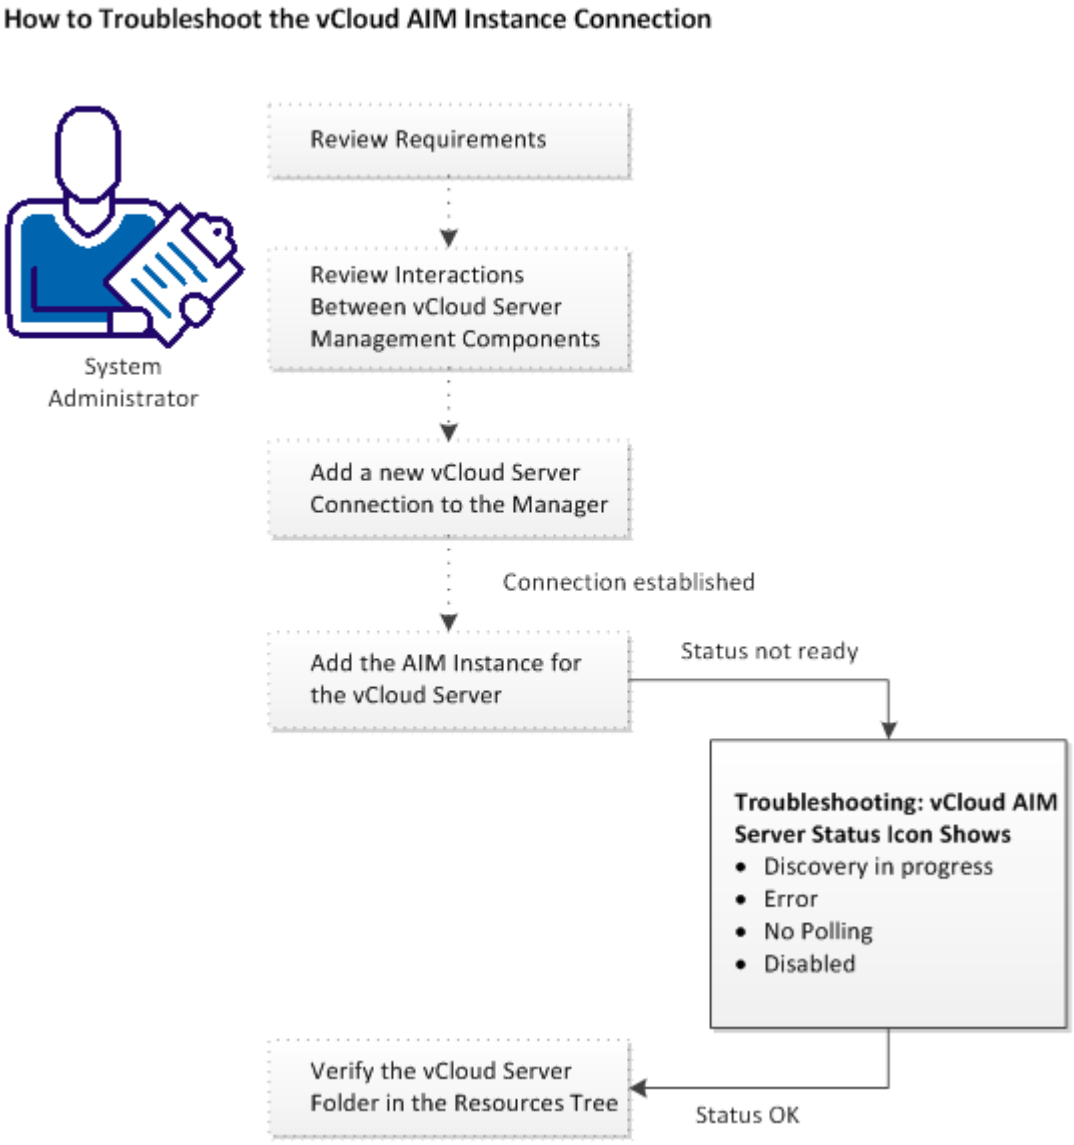

More information:

vCloud AIM Instance Status Icon Shows Discovery in Progress (see page 463)
vCloud AIM Instance Status Icon Shows Error (see page 463)
vCloud AIM Instance Status Icon Shows No Polling (see page 464)
vCloud AIM Instance Status Icon Shows Disabled (see page 465)

vCloud AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I have added a vCloud AIM instance for a vCloud Server under Administration, Configuration, the status icon shows  (Discovery in Progress).

Solution:

Wait until the discovery process of the vCloud environment has completed. The discovery duration depends on the number of managed objects related to virtual resources in vCloud. 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 vCloud Server folder to the Resources tree. Then you can start managing vCloud and its entire virtual infrastructure.

vCloud AIM Instance Status Icon Shows Error

Symptom:

After I have added a vCloud AIM instance for a vCloud 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 vCloud AIM:

- Verify, if the vCloud AIM Server is accessible.
- Verify, if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify, if the vCloud AIM server system is accessible

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping` *`servername`*

2. Verify the output of the commands to find out whether the vCloud AIM server has a valid DNS entry and IP address.

   If the vCloud AIM server is not in the DNS, add the vCloud AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the vCloud 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 servername

   Enter the correct IP address and vCloud AIM server name. For example:

   `192.168.50.51 myvCloudAIM`

4. Click (Validate) in the upper-right corner of the vCloud 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 vCloud 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, vCloud AIM Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the vCloud 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.

vCloud AIM Instance Status Icon Shows No Polling

Symptom:

After I add a vCloud AIM instance for a vCloud Director under Administration, Configuration, the status icon shows (no polling).

Solution:

No specific actions are required for the associated instance. This icon informs you that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular vCloud Director, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

vCloud AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered vCloud AIM instances in the network, the status icons of several instances show  (Disabled). This vCloud AIM instance is not managed.

This status appears, if CA Virtual Assurance has discovered a vCloud AIM with the following relationships:

- The vCloud AIM is configured for a vCloud Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a vCloud Server that has not been configured in the vCloud Servers pane.

Solution:

To change the status of the AIM instance to ready, do one of the following:

- Add the missing vCloud Server connection to the CA Virtual Assurance manager.
- Edit the existing vCloud Server connection and change its managed status to enabled.

Verify the VMware vCloud Folder in the Resources Tree

After a successful configuration and discovery, the new vCloud Server is listed in the Resources Explore pane under the VMware vCloud folder.

Follow these steps:

1. Click Resources, Explore.

   The Resources tree appears.

2. Expand VMware vCloud.

   The managed vCloud Director Servers appear.

3. Expand the new vCloud Director Server entry.

   The managed vCloud infrastructure appears: Organizations, vApps, VMs, ...

CA Virtual Assurance is now ready to manage the added vCloud environment with its virtual infrastructure.

Remote and Multi-instance vCloud Director Support

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

Each vCloud Director is uniquely associated with one Preferred vCloud AIM that you specify during configuration. Setting the preferred AIM indicates which AIM should be used for polling if multiple AIMs manage one vCloud Director.

vCloud Folder Structure

After a successful configuration of the connection to a vCloud Director server, CA Virtual Assurance discovers the vCloud Director environment that consists of organizations, vApps, and virtual machines. When the discovery completes, the VMware vCloud folder appears in the Explore pane of the Resources tab. You can expand the folder and can manage your vCloud environment.

The following diagram shows the object hierarchy underneath the VMware vCloud folder.

```
..
└──VMware vCloud
    └──VMware vCloud Director Servers
        └──VMware vCloud Organization
            └──VMware vApps
                └──VMware Virtual Machines
```

The VMware vCloud folder represents the service level at the top. The vCloud service can consist of multiple VMware vCloud Director servers. Each vCloud Director usually has multiple organizations with vApps and virtual machines.

At the organization level, you can provision vApps based on templates stored in a catalog.

vApp Support in vCloud

The vApps concepts in vCloud and vSphere environments are similar. Both represent an application object that can be operated on as a single entity. Usually, a vApp contains multiple VMs, each with its own purpose to the complete vApp application or service that it provides to the end user. Operations that are performed on the vApp are also performed on all VMs in the vApp. For example, both types define start and stop orders for all VMs in a vApp and define CPU and memory resource limits that all VMs in the vApp can use.

The purpose of vApps in vCloud is to be able to define an application or service once as a template, and make it accessible to multiple organizations through the organizations catalog. vCloud stores its data in the vCloud database which is different from the vCenter Server database.

Important! Do not operate on the VMs defined in vCloud directly from a vCenter Server. Those operations can cause the vCloud database to become out-of-sync with the actual defined VM. CA Virtual Assurance provides a limited set of operations for those VMs which appear under vCloud and vCenter Server so that the databases do not become out-of-sync.

Differences between vCloud and vSphere vApps

- A vCloud vApp does not provide the ability for a nesting hierarchy. vSphere vApps can contain other vApps and Resource Pools.
- In vCloud, CPU and memory resource limits are defined through Virtual Data Centers (vDC), and the vApp is mapped to one of those Virtual Data Centers.

  In vSphere, vApp resource limits are defined on the vApp itself.
- vCloud vApps can contain VMs that are defined on many different vCenter Servers and ESX Hosts.

  VMs in vSphere vApps are limited to VMs in a particular data center and cluster.
- vCloud vApps have lease limits. You can define a runtime and storage limit on the vApps. When the runtime limit is reached, a vCloud vApp can no longer be used. When the storage limit is reached, the vApp is deleted from the vCloud or is moved to the Expired Items folder, depending on organization lease policy.

  vSphere vApps remain in existence until a user manually deletes them.
- vCloud vApps are created from vApp templates. vApp templates are created by importing a VM from a vCenter Server or by importing an OVF package. vApps are created by deploying the template to the cloud for the organization on which the template was created. After deployment, additional VMs can be moved into the vApp.

  vSphere vApps are created by defining a vApp with the CPU and memory resource limits desired. Then VMs for the data center where the vApp is defined can be moved into the vApp.

vCenter Server as Resource Pool Provider for vCloud

You can configure the role of vCenter Server to serve as the resource pool provider for vCloud. In such cases, vCenter Server provides the compute and memory resources for vCloud to create VMs. The resource pool appears in vCloud as Provider vDC.

As a result of this configuration, the VMs of this resource pool appear in the CA Virtual Assurance Explore pane in the vCenter object hierarchy and in the vCloud object hierarchy. The Summary panel of such a VM shows the same information under vCloud as under the vCenter Server:

- Performance Chart
- General Information
- Overview (Status information of monitored resources)
- CPU and Memory Usage (Threshold configuration supported in vCenter Server only)
- Disk Usage

The set of operations that you can apply to these VMs is limited in both vCloud and vCenter Server environments. The limited set of operations prevents vCenter and vCloud from being out of synchronization. For example, you cannot power off a VM under the vCenter Server while the parent vApp of that VM is running in vCloud. You can only power off such a VM by first powering off the vApp in vCloud.

Valid VM operations are as follows:

- Deploy Monitoring Software
- Manage Automation Rules
- Configure Server Metrics Collection
- Configure Threshold Settings

If a VM is created in vCloud without a connection to a vCenter Server, the Summary pane shows the following information only:

- Item Type
- Name
- Operating Status

vCloud Organizations

A vCloud organization is a unit of administration for a collection of users, groups, and computing resources. Organizations provide catalogs to store vApp templates and media files. The members of an organization can use the vApp templates and media files in the catalog to create their own vApps.

A virtual data center (vDC) provides virtual computing resources to a vCloud organization. You can provision, run, and store virtual systems in a virtual data center. A vCloud organization can have multiple virtual data centers.

Provision vApp from Template

From the vCloud Organization level, you can provision vApps from templates which are stored in catalogs.

Follow these steps:

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Expand the VMware vCloud folder.

   The vCloud folder structure appears.

4. Right-click the organization object.

   The Provisioning pop-up menu appears.

5. Click Provision vApp from Template.

   The Provision new vApp from Template dialog appears.

6. Specify Name, vApp Template, Deployment lease, and Storage lease. Click OK.

   CA Virtual Assurance creates a vApp in the organization.

Operations on vApps in vCloud

From the vCloud Organization level, you can provision vApps from templates which are stored in catalogs.

Follow these steps:

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Expand the VMware vCloud folder.

   The vCloud folder structure appears.

4. Right-click a vApp object.

   The Management pop-up menu appears.

5. Select one of the available operations.

   Power On vApp

   Powers On the vCloud vApp.

   Power Off vApp

   Powers Off the vCloud vApp.

   Reset VApp

   Resets the vCloud vApp.

   Suspend vApp

   Suspends the vCloud vApp.

   Resume vApp

   Resumes the vCloud vApp.

   Clone vApp

   Creates a vCloud vApp from an existing vApp.

   Move vApp

   Moves a vCloud vApp to another virtual datacenter.

   Delete vApp

   Deletes the vCloud vApp.

   Modify vApp Lease

   Modifies the deployment and storage lease.

   Specify the required parameter values and click OK.

6. Click Events to verify the new status of the vApp.

   The list of events appears.

VMware vSphere and vCenter Server

CA Virtual Assurance manages VMware vSphere and vCenter Server virtual environments. The vCenter Server is the central component which CA Virtual Assurance and the vCenter AIM use to access the vSphere environment. SystemEDGE and the vCenter AIM run on the CA Virtual Assurance manager server or on an arbitrary Windows server.

CA Virtual Assurance provides connection and operational support for all VMware vCenter Server operations. The manager 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.

The vCenter AIM communicates with one or more remote vCenter Server instances through web-services. The AIM communicates with the manager through SNMP. If more than one vCenter AIM is available to manage a vCenter Server, you can specify your preferred vCenter AIM during configuration or you can let the manager choose it on its own.

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

Monitored vSphere and vCenter Server Resources

The vCenter AIM detects the logical and physical relationships between the components in a vSphere environment. The AIM provides a view of the entire virtualized environment and manages the following resource types and properties:

Datacenter

A datacenter serves as a container for your hosts, virtual machines, resource pools, or clusters. Depending on their virtual configuration, 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 infrastructure.

Datastore

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.

ESX Host

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

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.

Physical NIC

Specifies a physical Ethernet adapter on an ESX Server.

Resource Pool

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.

vApp

A vApp is a specific resource pool which treats a collection of VMs as a single unit. vApp uses the Open Virtualization Format. The Open Virtualization Format (OVF) is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels that are 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

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

Virtual Disk

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 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.

VMware Cluster/High Availability/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.

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.

vNetwork Standard 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 that are connected to virtual switch uplinks.

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.

How to Configure the vCenter Server Management Components

Review Requirements

Review the following requirements before you start configuring the vCenter Server management components of CA Virtual Assurance:

- 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 know 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 verified that the new vSphere environment and its vCenter Server are running properly.
- If the VMware 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 the SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote vCenter AIM Servers that you want to use.

More information:

Add a New vCenter Server Connection to the Manager (see page 477)
Verify the vCenter Server Folder Appearance in the Resources Tree (see page 487)
Review Interactions Between vCenter Server Management Components (see page 474)
Add the AIM Instance for the vCenter Server (see page 481)

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 other virtual resources. Based on thresholds, SystemEDGE and the vCenter AIM determine the status of a monitored resource and propagate 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.

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 that are automatically controlled by thresholds, status, and values that are 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 that are 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 does not specify any limitations. The required network connections are based on TCP/IP, SNMP, and web services.

Interaction Between vCenter Server Management Components

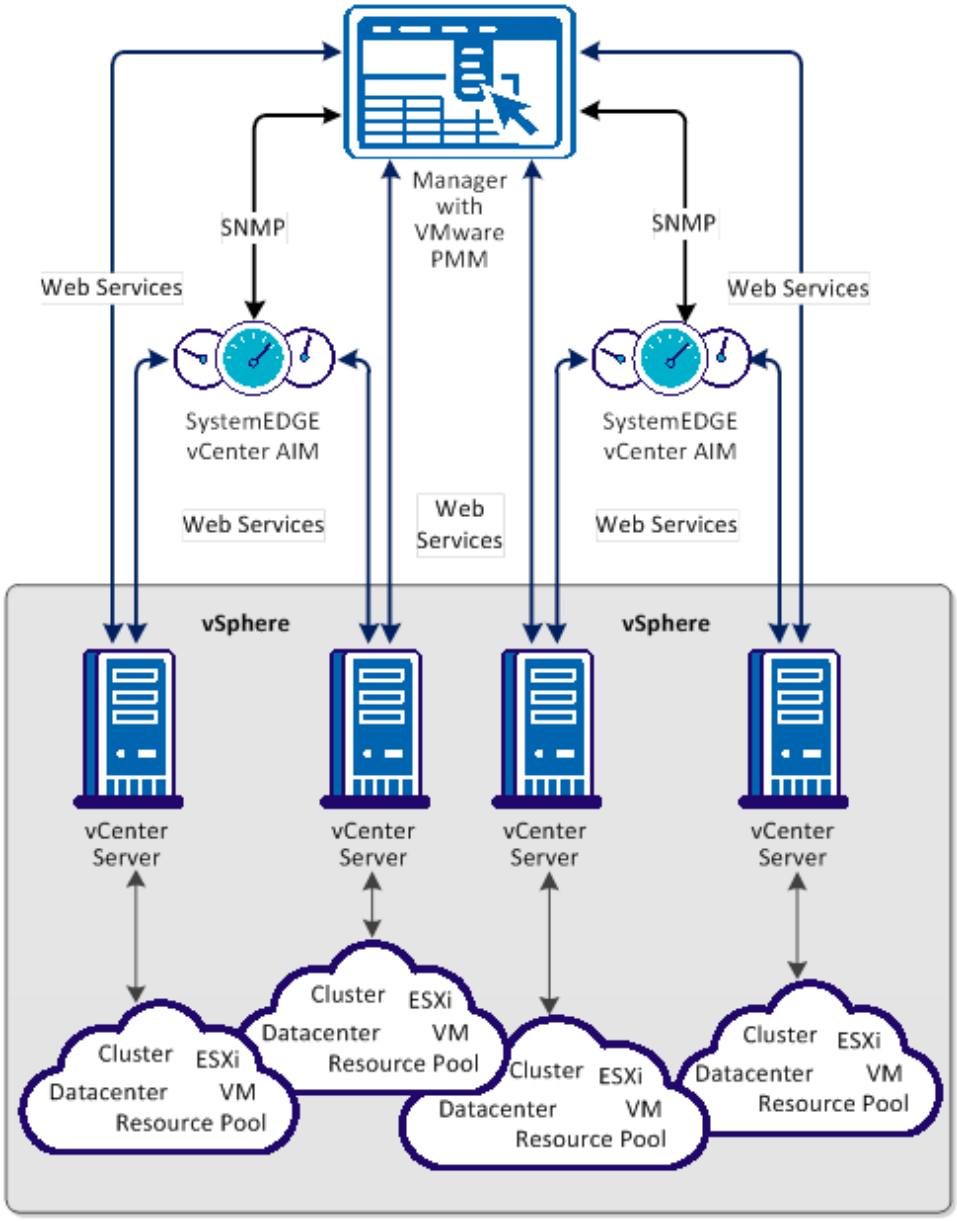

When you have configured the CA Virtual Assurance components successfully, CA Virtual Assurance discovers the new vSphere environment. After a successful discovery, the vCenter Server of the vSphere environment and its virtual infrastructure appear in the Resources tree of the CA Virtual Assurance Explore pane. You can then manage the new vSphere environment.

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:

Add a New vCenter Server Connection to the Manager (see page 477)
Verify the vCenter Server Folder Appearance in the Resources Tree (see page 487)
Add the AIM Instance for the vCenter Server (see page 481)

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. When this option is configured the Reservation Manager end user can access the VMware VM using a URL instead of using remote desktop connection.

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, web client protocol, web client port, web client user, web client user password), specify the preferred AIM, 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 478).

More information:

Verify the vCenter Server Folder Appearance in the Resources Tree (see page 487)
Add the AIM Instance for the vCenter Server (see page 481)
Troubleshoot the vCenter Server Connection (see page 478)

Troubleshoot the vCenter Server Connection

The vCenter Server connection has failed. Follow the troubleshooting information indicated in the following diagram:

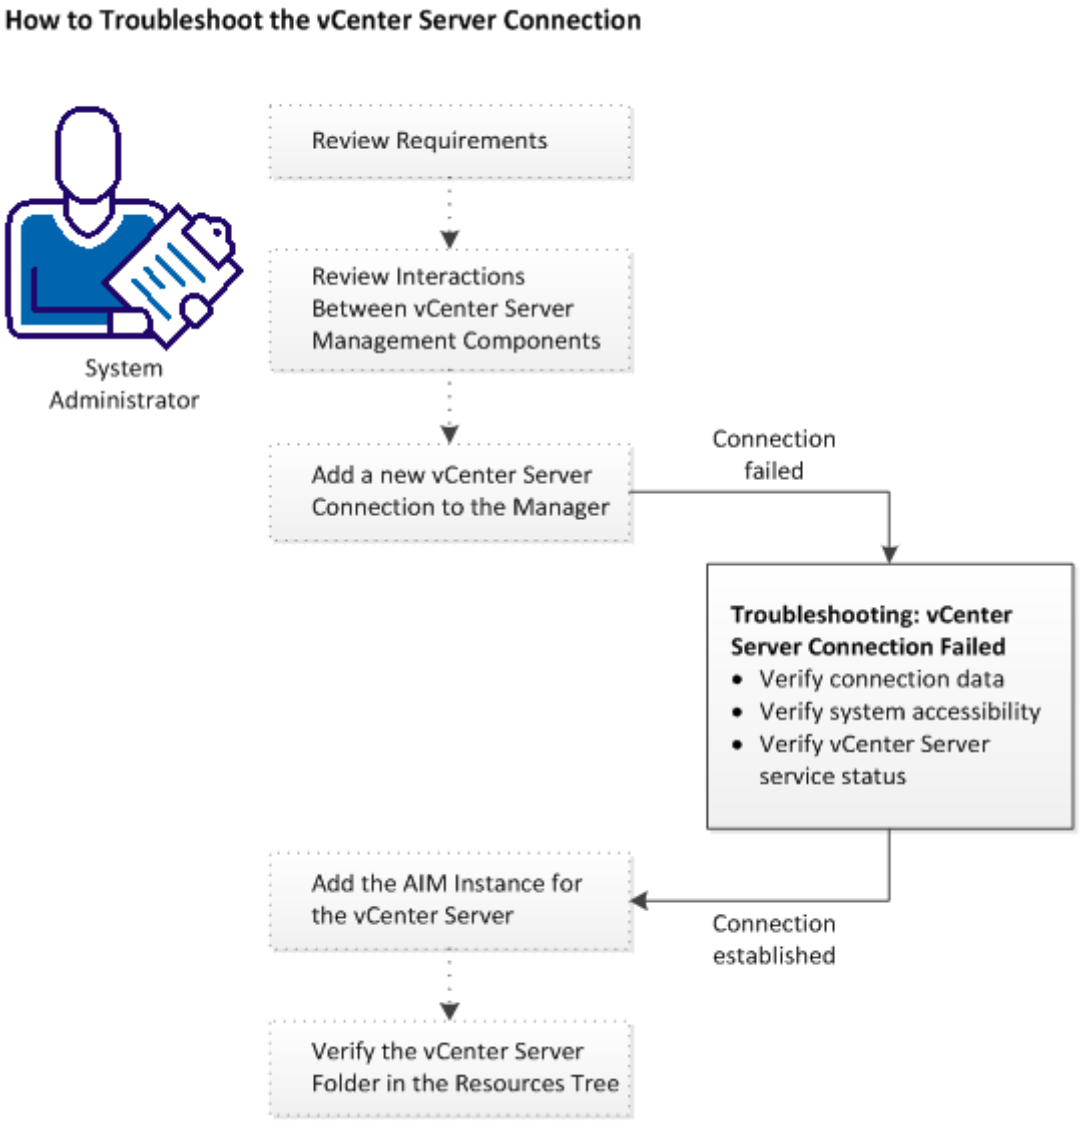

Follow these steps:

Verify the vCenter Server Folder Appearance in the Resources Tree (see page 487)
vCenter Server Connection Failed (see page 479)
Add the AIM Instance for the vCenter Server (see page 481)

vCenter Server Connection Failed

Symptom:

After I have added a 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, port, web client protocol, web client port (optional), web client user(optional), web client user password(optional). 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 other virtual components.

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 487)
Troubleshoot the vCenter AIM Instance Connection (see page 482)

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.

 Error - Unable to connect to the AIM. Check the network configuration.

 No Polling - The CA Virtual Assurance manager does not poll this AIM instance.

 Disabled - This instance is not managed.

Follow the troubleshooting information indicated in the following diagram:

How to Troubleshoot the vCenter AIM Instance Connection

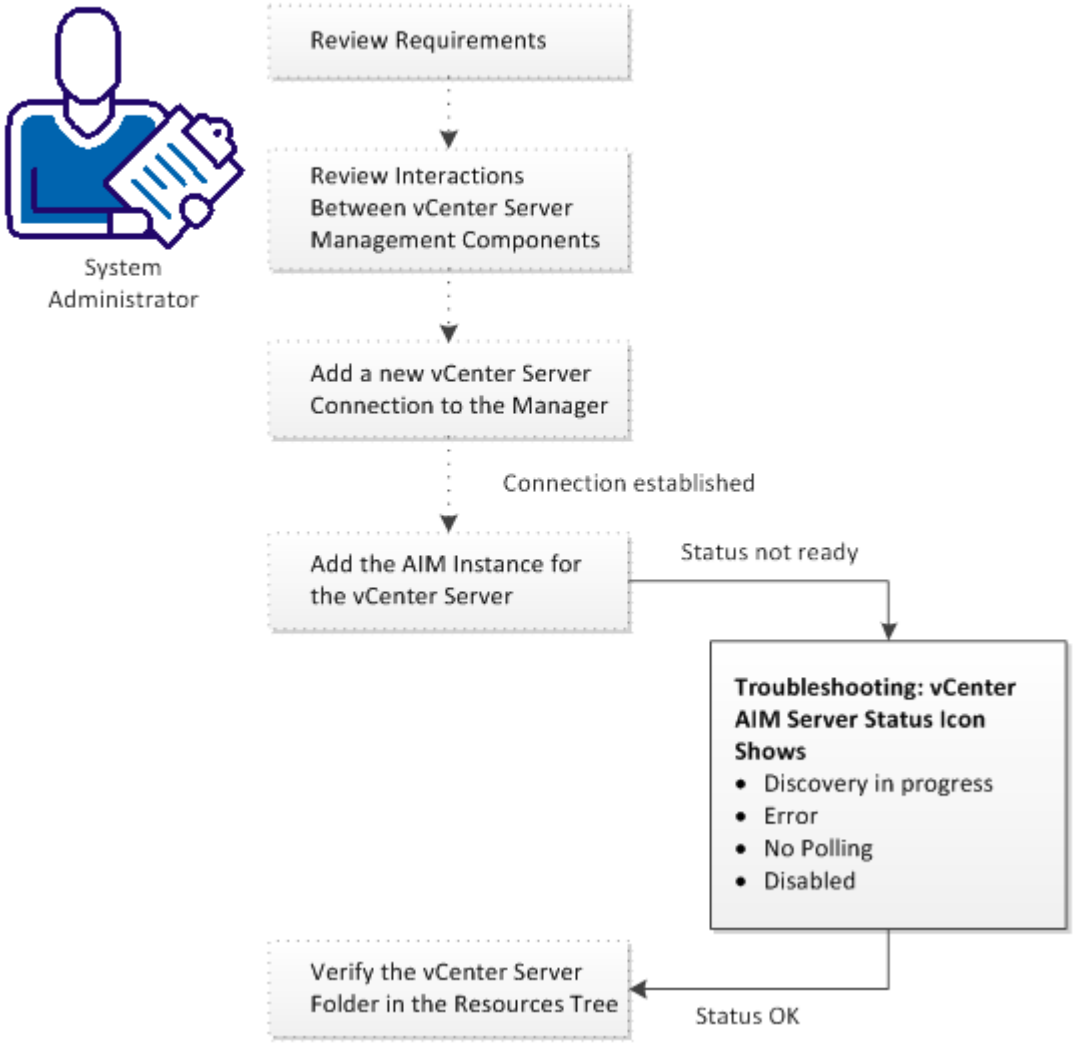

vCenter AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add 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 number of managed objects that are 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 finishes, 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 add 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` *`servername`*

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 host 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 servername

   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 No Polling

Symptom:

After I add a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows (no polling).

Solution:

No specific actions are required for the associated instance. This icon informs you that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular vCenter Server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

vCenter AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance discovers 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 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.

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.

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**productname*\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 also must 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**productname*\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>
       ...
   ```

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 page in the user interface). User-scoped authentication is disabled by default.

This scenario conforms to a full installation and 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).

Device Management for VMs

Device management includes the following actions:

- Add or Remove Virtual Disk (see page 489)
- Add or Remove Virtual Network Interface (see page 491)

Add or Remove Virtual Disk

You can dynamically add or remove virtual disks from a VM. The following disks can be added:

- A new disk from the same or another datastore
- An existing disk from the datastore
- Adding an existing disk from another datastore

Follow these steps:

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Find and right-click a virtual machine on the Explore pane and select Configuration,

To add a virtual disk

1. Click Add New Disk.

   The Add Disk dialog appears.

2. Enter the new disk details according to your needs.

   A message prompts for confirmation.

3. Click Ok.

   A message appears confirming that the new disk is added.

To remove a virtual disk

1. Click Delete Disk.

   The Delete Disk dialog appears.

2. Select the hard drive and whether to delete data.

   A message prompts for confirmation.

3. Click Ok.

   A message appears confirming that the disk is deleted.

Add or Remove Virtual Network Interface

You can dynamically add or delete a virtual network interface from an existing VM.

Follow these steps:

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Find and right-click a virtual machine on the Explore pane and select Configuration,

To add a virtual network interface

1. Click Add New Virtual Network Interface.

   The Add New Virtual Network Interface dialog appears.

2. Enter the new network interface details.

   A message prompts for confirmation.

3. Click Ok.

   A message appears confirming that the new card is added.

To remove a virtual network interface

1. Click Delete Virtual Network Interface.

   The Delete New Virtual Network Interface dialog appears.

2. Select the Network Interface you want to delete.

   A message prompts for confirmation.

3. Click Ok.

   A message appears confirming that network interface is deleted.

Fault Tolerance for Virtual Machines

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

Fault Tolerance Properties of Virtual Machines

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.

Monitor Fault Tolerance

To monitor fault tolerance properties

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Expand the VMware vCenter Server folder and the ESX server object.

   A list of VMs appears.

4. (Optional) Select the ESX host.

   The following FT attributes appear in the Summary tab:

   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

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

   Total Secondary VMs

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

   Powered on Primary VMs

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

   Powered on Secondary VMs

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

5. (Optional) Select the VM.

   The following FT properties display in the Summary tab.

   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.

Manage Fault Tolerance

You can control the fault tolerance properties of the VMs.

To manage fault tolerance properties of VMs

1. Select a VM in the Explore pane.

   The General Information pane appears on the right side, displaying the Fault Tolerance Status of the VM.

2. Right-click a VM, select Management, and select one action from the drop-down menu. The following actions for managing fault tolerance are available:

   - Turn Off Fault Tolerance
   - Enable Fault Tolerance
   - Disable Fault Tolerance
   - Migrate Secondary VM

3. Provide information and or confirmation for a selected action.

   Confirmation message appears.

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 (see page 496)
- Adding vRAM (see page 497)

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

Dynamically Add or Remove vCPU

You can dynamically add or remove CPU to VMs that have been provisioned. If hot plug is enabled for the VM, you can dynamically add vCPUs during runtime.

Note: To add or remove vCPU, the virtual machine must be turned off.

CA Virtual Assurance verifies the following VM properties:

- ESX license (ESX Level)
- Maximum supported vCPUs (ESX Level)
- Hot plug enabled (VM Level)

Examples

- If ESX license allows 8 CPUs (Enterprise Plus) AND Max Support vCPUs is 8 AND Hot Plug is DISABLED then you can add: 1, 2, 4, 8 CPUs
- If ESX license allows 8 CPUs (Enterprise Plus) AND Max Support vCPUs is 8 AND Hot Plug is ENABLED then you can add: 1, 2, 3, 4, 5, 6, 7, 8 CPUs

To add or remove vCPU

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane, and select Configuration, Add/Remove vCPU.

   The Modify vCPU dialog appears.

3. Adjust the number of CPUs according to your needs.

   A message prompts for confirmation.

4. Click Ok.

   A message appears confirming the modification.

Dynamically Add or Remove Memory

You can dynamically add or remove memory to VMs that have been provisioned. If hot plug is enabled for the VM, you can dynamically add memory during runtime.

Note: To add or remove memory, the virtual machine must be turned off.

To add or remove memory

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Find and right-click a virtual machine on the Explore pane, and select Configuration, Add/Remove Memory.

   The Modify Virtual Memory dialog appears.

4. Adjust the memory according to your needs.

   A message prompts for confirmation.

5. Click Ok.

   A message appears confirming the modification.

Logical Volumes in Virtual Machines

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

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 value 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 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.

Note: To set the SSRM memory allocation to unlimited configure the resource pool property “Allow memory over commitment” to a very high value for example, 999 (days). This indirectly sets the SSRM memory allocation to unlimited, and passes through to VMware to determine available memory based on the physical limits of the underlying resources(ESX server or cluster).

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 then you upgrade the host, the number of shares does 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 reservations to specify the *minimum* acceptable amount of CPU or memory, not the amount that 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 reservations for virtual machines, do not commit all resources. Plan to leave an appropriate portion unreserved, because when you move closer to 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.

Edit VM CPU and Memory Allocation

You can edit the number of CPU and memory shares allocated to a virtual machine to adjust its allocated resources. When you add resources, the appropriate amount of unassigned memory or CPU shares must be available for the operation to succeed. If values exist for the minimum and maximum allowed memory or CPU shares, any resource allocation change must stay within these limits.

You can also create and schedule policy with specific VM resource allocation actions.

To edit VM CPU and memory allocation

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Configuration, Resource Allocation.

   The Resource Allocation section appears.

3. Adjust the number of CPU and memory shares allocated to the virtual machine and click Save for each value that you edit.

   A confirmation message appears.

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 to help System Administrators 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 VMs. Using shares allows dynamic allocation of CPU resources. Each VM is allocated a specified number of 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 ensures dynamic actions that sustain business continuity. Using policy actions ensures that resources are allocated to virtual machines that are in need and deallocated when the need is gone.

How to Use Policy Actions to Identify Performance Issues

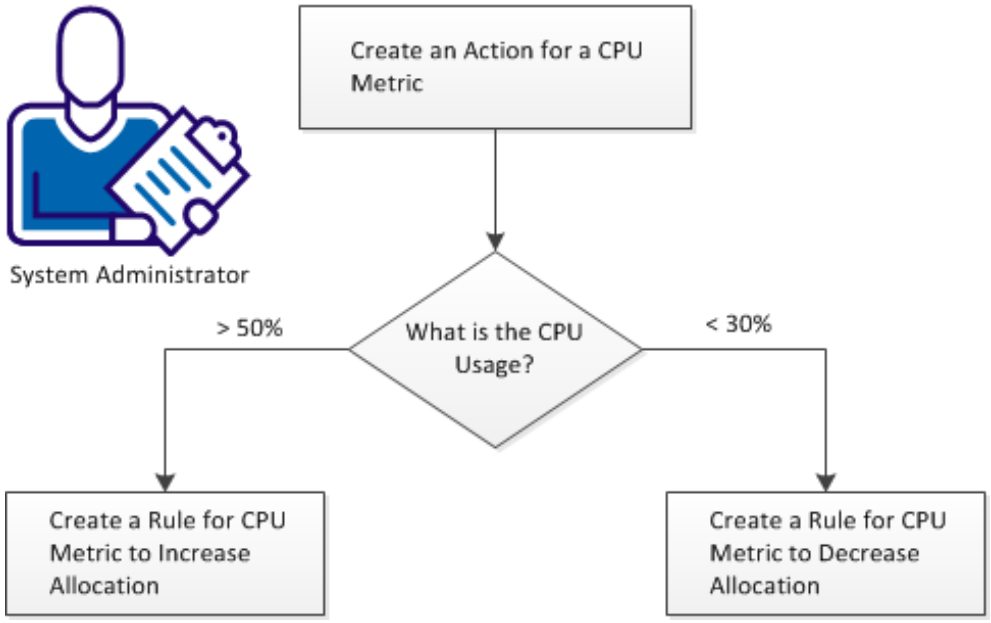

To identify and address performance issues using policy actions, follow these steps:

1. Create an action for CPU metric. (see page 502)
2. If CPU usage is more than 50 percent, create a rule for CPU metric to increase allocation. (see page 503)
3. If CPU usage is less than 30 percent, create a rule for CPU metric to decrease allocation. (see page 503)

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 the upper right side bar to add an action.
5. Enter the name of the Action.
6. Select Resource Configuration from the Category drop-down list.
7. Select Configure Shares from the Type drop-down list.
8. In the VC Server field, leave the entry as "%VCServer%" to apply this action on any VM across any VC Server.
9. In the 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 the 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 ensures dynamic resource allocation when the usage exceeds the threshold.

Follow these steps:

1. Click on the Resources tab, Policy, Rules.
2. Click '+' on the upper right side bar to add a rule.
3. Enter the 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. This rule decreases the CPU shares when the CPU usage is less than 30 percent.

Follow these steps:

1. Click on the Resources tab, Policy, Rules.
2. Click '+' on the upper right side bar to add a rule.
3. Enter the name of rule and click Next.
4. Select the action from the "Action Selection" list for the rule and click Next.
5. Enter the rule-based on metric where CPU usage is less than 30 percent to decrease the CPU shares of the VM.

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. The *Open Virtualization Format (OVF)* is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels that are 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.

A vApp is represented at the host and cluster level.

CA Virtual Assurance supports the following management operations on the vApp level:

- Discover
  - Server
  - Network
  - vCenter Server
- Capture Service
- Add Resource
- Clone vApp
- Power On vApp
- Power Off vApp
- Suspend vApp
- Delete from VMware vCenter
- Unregister from VMware vCenter
- Edit Sort Order

CA Virtual Assurance supports the following provisioning operations on vApps:

- Provision VMware VM
- Provision VMware vApp

Provision VMware vApp

You can create a vApp directly on the ESX host or cluster level, or as part of an existing resource pool or vApp.

Follow these steps:

1. From the host or clusters level in the Explore pane, right-click the ESX host or cluster.

   A pop-up menu opens.

2. Select Provisioning, Provision VMware vApp.

   The Create New vApp dialog appears.

3. Specify the following fields and click OK.

   Name

   Identifies the vApp.

   CPU Shares

   Specifies CPU shares for this vApp with respect to the total CPU resources of the parent host, resource pool, or vApp. Sibling vApps share resources according to their relative share values bounded by the reservation and limit. Specify a number of shares which expresses the appropriate proportional weight.

   For example, assume that you have vApp1 and vApp2 on a host and each has 1000 CPU shares. The weight is equal and each vApp can allocate 50 percent CPU time of the parent host. However, if vApp1 has 2000 CPU shares and vApp2 has 1000, then the weight is not equal. The total number is 3000 shares, and 1000 shares represent 33.3 percent, and 2000 shares represent 66.6 percent. So vApp1 can allocate 66.6 percent and vApp2 can allocate 33.3 percent of CPU time.

   CPU Reservation

   Specifies the guaranteed CPU allocation for this vApp.

   CPU Unlimited

   Disables the CPU limit setting. The actual limit is now set to the available physical resource.

   CPU Limit

   Specifies the upper limit for the CPU allocation for this vApp. You can usually accept the default.

   Memory Shares

Specifies memory shares for this vApp with respect to the total memory resources of the parent resource pool or vApp. Sibling vApps share resources according to their relative share values bounded by the reservation and limit. Specify a number of shares which expresses the appropriate proportional weight.

For example, assume that you have vApp1 and vApp2 on a host and each has 1000 memory shares. The weight is equal and each vApp can allocate 50 percent memory of the parent host. However, if vApp1 has 2000 memory shares and vApp2 has 1000, then the weight is not equal. The total number is 3000 shares, and 1000 shares represent 33.3 percent, and 2000 shares represent 66.6 percent. So vApp1 can allocate 66.6 percent and vApp2 can allocate 33.3 percent of memory.

Memory Reservation

Specifies the guaranteed memory allocation for this vApp.

Memory Unlimited

Disables the memory limit setting. The actual limit is now set to the available physical resource.

Memory Limit

Specifies the upper limit for the memory allocation for this vApp. You can usually accept the default.

The new vApp appears in the Explore pane.

Clone vApp

You can clone a vApp which is similar to clone a virtual machine.

Follow these steps:

1. From the host or clusters level in the Explore pane, select the vApp which you want to clone.
2. Right-click the vApp.

   A pop-up menu opens.
3. Select Management, Clone vApp.

   The Clone a vApp dialog appears.
4. Specify the following fields and click OK.

   Name

   Identifies the cloned vApp.

   Location

   Specify the appropriate location. Expand the object that is displayed on the pop-up menu and select the location.

   Datastore

   Specify the appropriate datastore from the drop-down menu.

   The cloned vApp appears in the Explorer pane.

More vApp Operations

CA Virtual Assurance supports the following additional operations on vApps:

- Power On
- Power Off
- Suspend
- Delete from VMware vCenter
- Unregister from VMware vCenter

Follow these steps:

1. From the host or clusters level in the Explore pane, select the appropriate vApp.
2. Right-click the vApp.

   A pop-up menu opens.
3. Select Management, and click the desired operation.

   A confirmation dialog appears.
4. Click OK.

   CA Virtual Assurance performs the selected operation.

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.

How to Import an OVF Package Using CA Virtual Assurance

This scenario provides information about importing OVF packages using CA Virtual Assurance. This information is meant to help System Administrators import OVF packages and deploy vApps that are specified in those OVF packages.
The *Open Virtualization Format (OVF)* is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels that are associated with it.

Follow these steps

Review Requirements (see page 510)
Provide Access to the OVF Package (see page 510)
dpmovf import Command--Import an OVF Package (see page 511)
Provide Custom Properties in Dialog Mode (see page 513)
Verify the Imported Objects in the Resources Tree (see page 513)

Review Requirements

Review the following requirements:

- You can access the CA Virtual Assurance user interface.
- You verified that the target vSphere environment and its vCenter Server are running properly.
- You verified that you can start the CMD window as administrator and the dpmovf.exe file is installed on the computer.

More information:

Add a New vCenter Server Connection to the Manager (see page 477)
Verify the vCenter Server Folder Appearance in the Resources Tree (see page 487)
Review Interactions Between vCenter Server Management Components (see page 474)
Add the AIM Instance for the vCenter Server (see page 481)

Provide Access to the OVF Package

To access the OVF packages from CA Virtual Assurance, do *one* of the following tasks:

- On the Manager, map the drive where the OVF package is located.
- Copy the OVF package on the Manager.

dpmovf import Command--Import an OVF Package

The dpmovf import command imports the OVF package and creates VMs or vApps. You can provide a custom properties file by using the -properties attribute. A custom properties file allows you to specify custom properties that are defined in the OVF package. The custom properties file contains a list of property keys and the corresponding property values.

Note: If you do not have a custom properties file, the properties.txt file is created in your working directory. The default directory is CA\ProductName\bin.

This command has the following format:

```
dpmovf import
-host vCenter_server
-user user_name
-password user_password
-name VM_VApp_name
-path OVF_file_path
-datacenter data_center
-datastore data_store
-resourcepool resource_pool
[-locale iso639value]
[-properties properties_file]-
```

-host vCenter_server

Specifies the name of the vCenter server host.

-user user_name

Specifies the user name to log in.

-password user_passsword

Specifies the user password to log in.

-name VM_VApp_name

Specifies the name of the VM or the vApp.

-path OVF_file_path

Specifies the OVF file path.

-datacenter data_center

Specifies the data center name.

-datastore data_store

Specifies the data store.

-resourcepool resource_pool

Specifies the resource pool.

-locale iso639value

(Optional) Specifies an IS0 639_3166 combination to override the default English output, for example, fr_FR for French. To use the locale of the command prompt, specify "native".

-properties properties_file

(Optional) Specifies the custom properties file path.

Example: Import the OVF file for CA Platform using CA Virtual Assurance

This example imports CA Platform OVF package and creates a vApp and VMs. The CA Platform OVF file is *CA Platform_v1_0_0_92c.ovf* and the file location is *D:\OVF\CA_Platform*. The username is *user123.* The following attributes for the vApp are specified: *my_datastore*, *my_datacenter*, and *my_resourcepool*. The custom properties are provided in the *custom_properties.txt* file.

```
dpmovf import -path "D:\OVF\CA_Platform\CA Platform_v1_0_0_92c.ovf" -name
"My_CA_Platform" -host my_host.company.com -user user123 -locale en-US -datastore
"my_datastore" -datacenter "my_datacenter" -resourcepool "my_resourcepool"
-properties "custom_properties.txt"
```

Provide Custom Properties in Dialog Mode

If the OVF file contains custom properties, you can edit custom properties in dialog mode. If you specified a custom properties file, you can overwrite the custom properties file in dialog mode.

Note: If you do not have a custom properties file, the properties.txt file is created in your working directory. The default directory is CA\ProductName\bin.

Follow these steps:

1. Type the custom property number for the custom property that you want to edit.
2. Type the custom property value.
3. Repeat steps 1 through 2 for all custom properties that you want to provide or edit.
4. Enter *any* of the following options:

   r

   Reads the properties file.

   w

   Overwrites all properties in the properties file.

   c

   Executes the import command.

   Note: Some of the provided properties are validated to verify if the conditions are met, or if provided values are valid.

CA Virtual Assurance deploys the OVF to vCenter and you can see vApps and VMs that are specified in the OVF file.

Verify the Imported Objects in the Resources Tree

After a successful import of vApps and VM the added instances are listed in the Resources Explore pane under the VMware vCenter Server folder.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand VMware vCenter Server.

   The imported objects appear.

The vApps that were specified in the OVF file are imported in your vCenter environment. CA Virtual Assurance is now ready to manage the added vApps and VM in your vSphere environment.

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.

Virtual Standard Switches and Virtual Distributed Switches in the vNetwork Panel

The vNetwork panel is available in the user interface at the VMware datacenter level and ESX host level. At the VMware datacenter level vNetwork indicates the Virtual Distributed Switches of that datacenter. At the ESX host level vNetwork indicates the associated Virtual Distributed Switches and Virtual Standard Switches.

More information:

vNetwork Standard Switches (vSwitch) (see page 514)
Distributed Virtual Switches (see page 515)
Properties (see page 517)
Actions (see page 521)
Monitor Distributed Virtual Switches Through Events (see page 522)

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.

You can expand the Virtual Standard Switch objects to see the associated ports and portgroups.

- The port groups contain associated virtual machines that use the portgroup.

Distributed Virtual Switches

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 Distributed Virtual Switch operates as a single virtual switch that spans across all hosts which are associated with that switch. A Distributed Virtual Switch 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 Distributed Virtual Switch is a network hub that VMs can use. A Distributed Virtual Switch 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 Distributed Virtual Switch and specify port configuration options for each member port. dvPort Groups define how a connection is made through the Distributed Virtual Switch 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 Distributed Virtual Switch.

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.

You can expand the Virtual Distributed Switch objects to see the associated portgroups and uplink groups.

- The port groups contain associated VMs that use the portgroup.
- The Uplink Groups list the physical uplink adapters.

Properties

The Properties pane displays the properties of Virtual Standard Switches or Virtual Distributed Switches.

vSwitch Properties

The following vSwitch properties indicate port number characteristics:

Number of Ports

Indicates the current number of ports of the Virtual Distributed Switch or Virtual Standard Switch.

Maximum Number of Ports

Indicates the maximum number of ports of the Virtual Distributed Switch.

Note: For Virtual Distributed Switches, this information is only available on the VMware datacenter level. On the ESX host level, it is not available.

Port Group Properties

The following port group property indicates the VLAN ID:

VLAN ID

Indicates the VLAN ID of a port group.

Port Properties

The following properties specify port characteristics:

VLAN ID

Indicates the VLAN ID of a port.

Type

Indicates the type of a port, for example, VMkernel Port or Service Port.

Network Properties

The following properties specify network characteristics of the virtual switch:

- IPv4 Address
- IPv6 Address
- MAC Address

Virtual Machine Counts

The following values provide statistical information about VMs associated with a port group.

- Powered On
- Powered Off
- Suspended
- Unknown

Policies

The following list contains default policies or enabled policies for Virtual Standard Switches or Virtual Distributed Switches.

Promiscuous Mode

Indicates whether all traffic is seen on the port.

MAC Address Changes

Indicates whether the Media Access Control (MAC) address can be changed.

Forged Transmit

Indicates if the MAC address is different from the MAC address of the virtual network adapter.

Traffic Shaping

Indicates whether traffic shaper is enabled on the port.

Average Bandwidth

Indicates the average bandwidth in bits per second if shaping is enabled on the port.

Peak Bandwidth

Indicates the peak bandwidth during bursts in bits per second if traffic shaping is enabled on the port.

Burst Size

Indicates the maximum burst size allowed in bytes if shaping is enabled on the port.

Network Failure Detection

Indicates whether network failure detection is enabled. Valid values are:

- false (1)
- true (2)

Notify Switches

Specifies whether to notify the physical switch if a link fails.

Fallback

Indicates if fallback is enabled.

Policy Inbound Frames

Indicates whether the teaming policy is applied to inbound frames.

Active Adapters

Displays a list of active network adapters used for load balancing.

Standby Adapters

Displays a list of standby network adapters used for failover.

vSwitch Properties

The following vSwitch properties indicate port number characteristics:

Number of Ports

Indicates the current number of ports of the Virtual Distributed Switch or Virtual Standard Switch.

Maximum Number of Ports

Indicates the maximum number of ports of the Virtual Distributed Switch.

Note: For Virtual Distributed Switches, this information is only available on the VMware datacenter level. On the ESX host level, it is not available.

Port Group Properties

The following port group property indicates the VLAN ID:

VLAN ID

Indicates the VLAN ID of a port group.

Port Properties

The following properties specify port characteristics:

VLAN ID

Indicates the VLAN ID of a port.

Type

Indicates the type of a port, for example, VMkernel Port or Service Port.

Network Properties

The following properties specify network characteristics of the virtual switch:

- IPv4 Address
- IPv6 Address
- MAC Address

Virtual Machine Counts

The following values provide statistical information about VMs associated with a port group.

- Powered On
- Powered Off
- Suspended
- Unknown

Actions

Use the appropriate actions to manage your Virtual Standard Switches and Virtual Distributed Switches. The following actions are available:

- Add vSwitch
- Update vSwitch
- Remove vSwitch
- Add Portgroup
- Update Portgroup
- Remove Portgroup
- Rename Portgroup

When you apply these actions, a dialog opens and prompts you to enter the required information. Possible fields are:

Switch Name

Specifies the switch name to perform the operation on.

NICs

(Optional) Specifies lists of physical NICs associated with the ESX host members.

Uplink Port Names

(Optional) Specifies a list of uplink port names to use.

Maximum Number of Ports

(Optional) Specifies the maximum number of ports for the Virtual Distributed Switch.

Bindtype

(Optional) Specifies the bind type of the port group. Valid values are:

earlyBinding

Assigns the ports when the VM binds to the portgroup. This type of binding ensures connectivity at all times, but permanently reserves the port. This binding type is the default.

lateBinding

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. LateBinding is configurable through vCenter.

ephemeral

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. Ephemeral binding is configurable through the ESX Host and vCenter.

Number of Ports

(Optional) Specifies the number of ports of the port group.

Portgroup Name

Specifies the port group name.

New Portgroup Name

Specifies the new port group name.

LAN ID *vlanid*

(Optional) Specifies an Integer value (*vlanid*) used for the virtual portgroup operations.

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.

VMware vCenter Provisioning and Common Use Cases

This section provides instructions how to provision virtual resources and to perform common use cases.

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.

9. Click Next.

10. 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.

11. 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.

Clone a Virtual Machine

Cloning a virtual machine creates a copy of the virtual machine that you can place anywhere in the same virtual machine farm. You can also customize the guest operating system when creating a clone. You can only clone a virtual machine when it is in the powered off state.

To clone a virtual machine

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click the virtual machine to clone on the Explore pane and select Management, Cloning.

   The Cloning pane appears.

3. Complete the following fields and click Clone:

   Name

   Specifies the VM clone name.

   Datastore

   Specifies the datastore under which to store the cloned VM. The datastore must be in the same form as the source VM.

   Custom Spec

   Specifies the guest operating system specification to use. You can select the default or a customization.

   Destination Resource Pool

   Specifies the pool from which the cloned VM obtains resources.

   A message appears confirming the request submission.

4. Click the Summary tab for the virtual machine.

   Verify that an event appears to confirm the operation. The clone appears in the Explore pane after the operation completes.

Manage VM Status (VMware)

You can control the status of vCenter Server virtual machines by performing one of the following VM operations:

- Power On
- Power Off
- Suspend
- Reset
- Shut Down

You can perform any of these operations on multiple VMs simultaneously.

To control VM status

1. Select the virtual machine on which you want to perform a status operation in the Explore pane.
2. Right-click the VM, select Management. You can also click Quick Start and click the related link of power control. Select *one* of the following:

   Power On

   Starts the virtual machine and boots the guest operating system. You can only power on a virtual machine that is currently powered off or suspended.

   Power Off

   Powers off the virtual machine. You can only power off a virtual machine that is currently powered on or suspended.

   Suspend

   Pauses the virtual machine and saves its current state. All activity is suspended until you resume the machine.

   Reset

   Shuts down the guest operating system and restarts it.

   Shutdown

   Shuts down the guest operating system. You can only shut down a virtual machine that is currently powered on.

   A confirmation dialog appears.

3. Click OK.

   The status operation occurs, and a confirmation message appears. Refresh the interface to view the new VM status. An event should appear confirming the result of the operation.

The following icons indicate VM status:

Indicates that the VM is in critical state.

Indicates that the VM is in warning state.

Indicates that the VM is in normal state.

Indicates that the VM is in unknown state.

Convert a Template to a Virtual Machine

You can convert a virtual machine template to a virtual machine. When you convert a template to a VM, the VM uses the template name and settings.

To convert a template to a virtual machine

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine template on the Explore pane and select Management, Convert to Virtual Machine.

   The Conversion page appears.

3. Select the ESX server and resource pool for the virtual machine and click Convert.

   A message appears confirming the request submission.

4. Click the Summary tab for the virtual machine template.

   Verify that an event appears to confirm the operation. After the operation completes, the template appears as a virtual machine on the Explore pane when you refresh the interface.

Convert a Virtual Machine to a Template

You can convert a powered off virtual machine to a template to use the virtual machine's configuration as a base for other virtual machines.

To convert a virtual machine to a template

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Convert to Template.

   A confirmation dialog appears.

3. Click OK.

   A message appears confirming the request submission.

4. Click the Summary tab for the virtual machine.

   Verify that an event appears to confirm the operation. After the operation completes, the virtual machine appears as a template in the Explore pane when you refresh the interface.

Create a Snapshot

Create a snapshot to preserve the current state of a virtual machine so that you can return to the same state at a later time. A snapshot preserves the entire state of the virtual machine, including memory contents, settings, and virtual disk state. You can create snapshots for virtual machines that are powered on, powered off, or suspended.

To create a snapshot

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots, Create New Snapshot.

   The Create New Snapshots dialog appears.

3. Enter the snapshot name and description, specify whether to enable capture memory, and click OK.

   A confirmation message appears.

4. Click Summary for the virtual machine.

5. Verify that an event confirms the operation.

   The snapshot appears in the Snapshots pane once the operation is complete.

Delete a Snapshot

You can delete a snapshot that you no longer need.

To delete a snapshot

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots, Modify Snapshots.

   The Snapshots pane appears and displays all existing snapshots for the virtual machine.

3. Select a snapshot and select Delete from the menu.

   A confirmation dialog appears.

4. Click OK.

   A message appears confirming the request submission.

5. Click the Summary tab for the virtual machine.

   Verify that an event appears to confirm the operation. The snapshot disappears from the Snapshots pane after the operation completes.

Delete all Snapshots

You can delete all existing snapshots for a virtual machine in one operation.

To delete all snapshots

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots.

   The Snapshots pane appears and displays all existing snapshots for the virtual machine.

3. Select Delete All from the menu.

   A confirmation dialog appears.

4. Click OK.

   A confirmation message appears.

5. Click Summary for the virtual machine.

   Verify that an event appears to confirm the operation. All snapshots disappear from the Snapshots pane after the operation completes.

Delete a Virtual Machine

When you delete a virtual machine from VMware vCenter Server, the virtual machine is deleted from the virtual disk.

To delete a virtual machine

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Delete from vCenter Server.

   A confirmation dialog appears.

3. Click OK.

   A message appears confirming the request submission.

4. Click the Summary tab for the virtual machine.

   An event should appear confirming the result of the operation. If successful, the virtual machine is deleted from the virtual disk, and the virtual machine disappears from the Explore pane after you refresh the interface.

Deploy a Virtual Machine from a Template

You can deploy a virtual machine from a template to use the template's settings to create and deploy a new virtual machine.

To deploy a virtual machine from a template

1. Right-click VMware vCenter Server in the Explore pane and select Provisioning, Provision VMware VM.

   VMware vCenter Provisioning dialog appears.

2. Specify the all required fields and select the appropriate VC Template.

   Click Next

3. Perform the remaining steps to specify the virtual hardware for that virtual machine. Click Finish.

   A message appears confirming the request submission.

4. Verify that an event confirms the operation.

   After the operation completes, the new virtual machine appears in the Explore pane when you refresh the interface.

Manage Cluster Services

You can control the status of the following services on VMware vCenter clusters:

HA

Allows automatic migration and restarting of VMs when a host fails.

DRS

Lets you manage hosts as a collection of resources. The DRS service migrates VMs to hosts and resources to VMs as necessary.

To manage cluster services

1. Select a VMware vCenter cluster on the Explore pane.

   The Overview pane appears on the right side, displaying the status of the HA and DRS services.

2. Select Enable or Disable from the drop-down menu.

   The status of the service changes.

Migrate a Virtual Machine

You can migrate a virtual machine to move it to another ESX host. You can migrate a powered off machine or migrate a powered on machine with VMotion. You cannot migrate a virtual machine that is suspended.

To migrate a virtual machine

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click the virtual machine to migrate on the Explore pane and select Management, Migration.

   The Migration pane appears.

3. Enter the destination ESX server and resource pool for the virtual machine, and click Migrate.

   Note: VM migration between ESX hosts is only supported when the VMs datastore/disk is shared between the two ESX hosts.

   A confirmation message appears.

4. Click the Summary tab for the virtual machine.

   Verify that an event appears to confirm the operation. The virtual machine appears in its migrated location in the Explore pane after the operation completes.

Monitor a Virtual Machine

You can monitor the status and the properties of VMs in detail.

To monitor virtual machines

1. Click Resources

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Expand the VMware vCenter Server folder and the ESX server object.

   A VM list appears.

4. Click the Summary tab.

   The right pane displays general information, FT properties, overview, CPU and memory usage, disk usage (logical volumes), and events.

   In the Overview panel, the disk states indicate the virtual hardware state of the virtual disk, as calculated by SystemEDGE, and based on monitors configured in the vCenter AIM. This information is based on true performance data of the virtual disk, in terms of reads and writes per second.

   In the Disk Usage panel, the disk states indicate the usage of the logical volumes as viewed through the guest operating system. The state is calculated by SystemEDGE, and is based on monitors configured in the vCenter AIM. This information is only valid when the VM and the guest operating system are running.

   The General Information panel provides details about the connection state of the VM. Valid connection state values are as follows:

   - Not connected
   - Connected
   - Orphaned

   The orphaned connection state can happen during Cluster failover situations. When a virtual machine has been marked as orphaned, states reflected under the Overview section are based upon data collected prior to becoming orphaned.

Monitor an ESX Server

You can monitor the status and the properties of ESX servers in detail.

To monitor ESX servers

1. Click Resources

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Expand the VMware vCenter Server folder and select an ESX server.

4. Click the Summary tab.

   The right pane displays general information, FT attributes, overview, CPU and memory usage, utilization, and events.

5. Click the vNetwork tab.

   The right pane displays a list of associated virtual standard switches (vSwitches) and virtual distributed switches (vDS).

6. Select a virtual switch from the list.

   The right pane displays the properties of the virtual switch.

Revert to a Snapshot

When you revert to a snapshot, you return the virtual machine to its exact state when the snapshot was taken.

To revert to a snapshot

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots.

   The Snapshots pane appears and displays all existing snapshots for the virtual machine.

3. Select a snapshot and select Revert from the menu.

   A confirmation dialog appears.

4. Click OK.

   A confirmation message appears.

5. Click the Summary tab for the virtual machine.

   Verify that an event appears to confirm the operation.

Unregister a Virtual Machine

When you unregister a virtual machine from the vCenter Server, the virtual machine still exists but is removed from VMware vCenter Server inventory.

To unregister a virtual machine

1. Click Resources

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Find and right-click a virtual machine on the Explore pane and select Management, Unregister from vCenter Server.

   A confirmation dialog appears.

4. Click OK.

   A message appears confirming the request submission.

5. Click the Summary tab for the virtual machine.

   An event should appear confirming the result of the operation. If successful, the virtual machine is removed from the vCenter inventory.

vCenter Automation and Policy Actions

The following action types are available for use with VMware vCenter Server:

- Add Disk (see page 616)
- Add Network Interface (see page 618)
- Configure Shares (see page 640)
- Configure CPU/Memory (see page 627)
- Configure Power (see page 636)
- Convert Template to Virtual Machine (see page 641)
- Convert Virtual Machine to Template (see page 643)
- Delete Machine (see page 651)
- Manage VM Snapshots (see page 659)
- Modify CPU (see page 667)
- Modify Memory (see page 668)
- Provision Machine (see page 680)
- Remove Disk (see page 683)
- Remove Network Interface (see page 684)
- Migrate Machine (see page 666)

You can use these action types to create new actions that automate vCenter power, resource allocation, and other operations when assigned rule criteria are met. You can also schedule these actions to occur at specific times.

For more information about using actions and rules to create automation policy, see the "Policy" section.

View Custom Specifications

Custom specifications are custom versions of guest operating systems that you are using on virtual machines. You can view all current custom specifications, the date of the last update, and the current version number.

To view custom specifications

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Find and select a VMware vCenter server.

   The server page appears in the right pane.

3. Click the Configuration tab and select the Customization Specifications submenu.

   The Customization Specifications section appears and displays existing customization specifications.

View General Information

CA Virtual Assurance displays General Information in the right hand pane and provides resource properties at the following levels in the object hierarchy:

- vCenter Server
- ESX Server
- Resource Pool
- VM

Resource properties include information about the following categories:

- Names, item types, versions
- Quantitative characteristics of CPU and memory
- Number of VMs and Resource Pools
- Current mode of the resource

Additionally, CA Virtual Assurance displays connection state, power state, and Fault Tolerance information about the VM level.

Valid Fault Tolerance status values are:

- Not Fault Tolerant
- Protected
- Not Protected (Starting)
- Not Protected (Need Secondary VM)
- Not Protected (Disabled)
- Not Protected (VM Not Running)

The Secondary Location values are:

- Not Available
- Total Secondary CPU Usage
- Total Secondary Memory

The General Information panel provides details about whether Fault Tolerance is configured, version, and various counts of supported FT VMs. The counts consider:

- Total Primary VMs
- Total Secondary VMs
- Powered On Primary VMs
- Powered 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.

More Information

Monitor an ESX Server (see page 534)
Monitor a Virtual Machine (see page 533)

Chapter 8: Monitoring Clusters and Virtual Desktops

This section contains the following topics:

Citrix XenDesktop Environments

CA Virtual Assurance monitors Citrix XenDesktop environments remotely. Citrix XenDesktop AIM provides statistical data and helps detect issues within Citrix XenDesktop environments. The monitoring includes but is not limited to desktops, controllers, machines, catalogs, hypervisor connections, and service statistics.

Interaction Between Citrix XenDesktop Management Components

The following diagram illustrates how the components involved in Citrix XenDesktop management interact. The AIM Server is a Windows Server on which SystemEDGE and the XenDesktop AIM run. The communication between the XenDesktop AIM and the Citrix XenDesktop Controller uses Windows Remote Management (WinRM). The communication between the XenDesktop AIM and Citrix XenDesktops in your environment uses WMI. CA Virtual Assurance can connect to multiple Citrix XenDesktop Controllers, and you gain an overall view of your Citrix XenDesktop environment.

Interaction Between Citrix XenDesktop Management Components

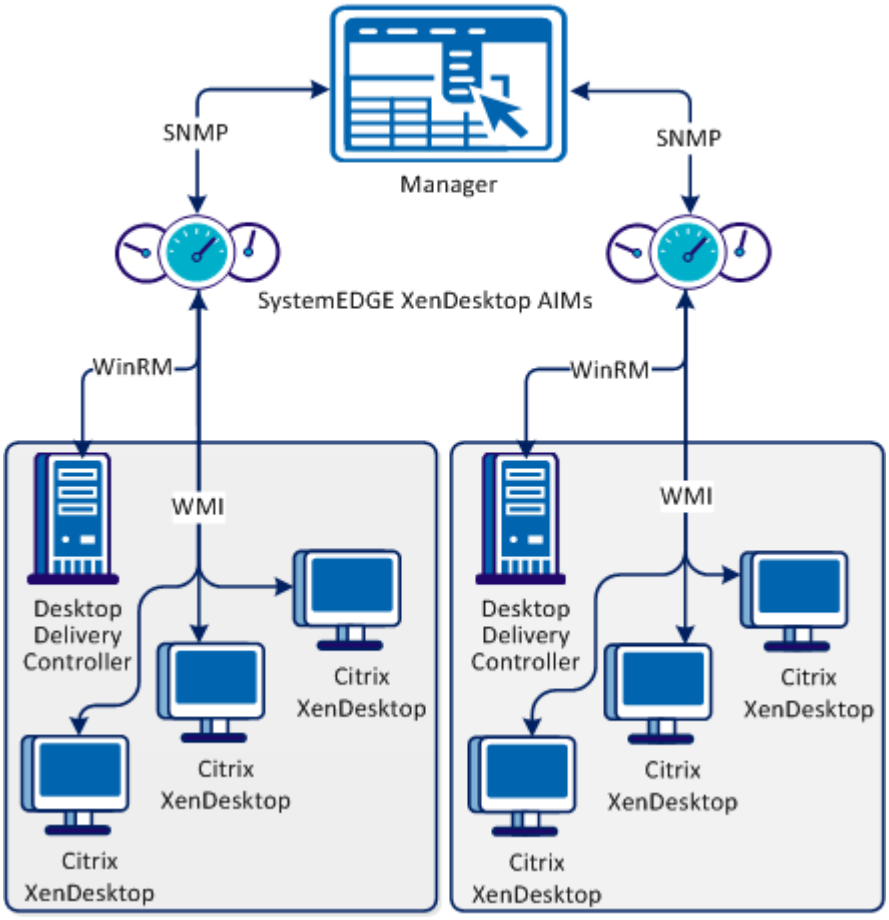

To add the required connection information for Citrix XenDesktop Controller, use the following method:

- NodeCfgUtil.exe utility on the AIM Server

The connection information is written to the configuration file on the managed node. The XenDesktop AIM polls the configuration file and starts monitoring your Citrix XenDesktop environment through the Citrix XenDesktop Controller or directly from Citrix XenDesktops.

Citrix XenDesktop Prerequisites

The listed prerequisites are required for installing XenDesktop AIM. Verify that the following components are installed on the machine where XenDesktop AIM is installed:

- Microsoft .NET Framework 4.0
- Windows Management Framework Core (Windows PowerShell 2.0, Windows Remote Management (WinRM) 2.0)

  Note: For more information about the Windows Management Framework, see the Microsoft 968929 Knowledge Base article.

Add Machine Name to the Trusted Hosts List

If a Citrix XenDesktop is in a different domain, add the machine name to the trusted hosts configuration settings for WinRM service on the AIM machine.

Use the following command:

```
set-Item wsman:\localhost\client\trustedhosts machine_dnsname
```

machine_dnsname

Specifies the list of full dns names of computers that XenDesktop AIM connects to.

IBM PowerHA

CA Virtual Assurance monitors IBM PowerHA, previously known as High Availability Cluster Multiprocessing (HACMP). CA Virtual Assurance monitors the clusters remotely, detects any failure, and provides details on alerts in the cluster and any other environmental problems.

Interaction Between IBM PowerHA Management Components

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

Interaction Between PowerHA Management Components

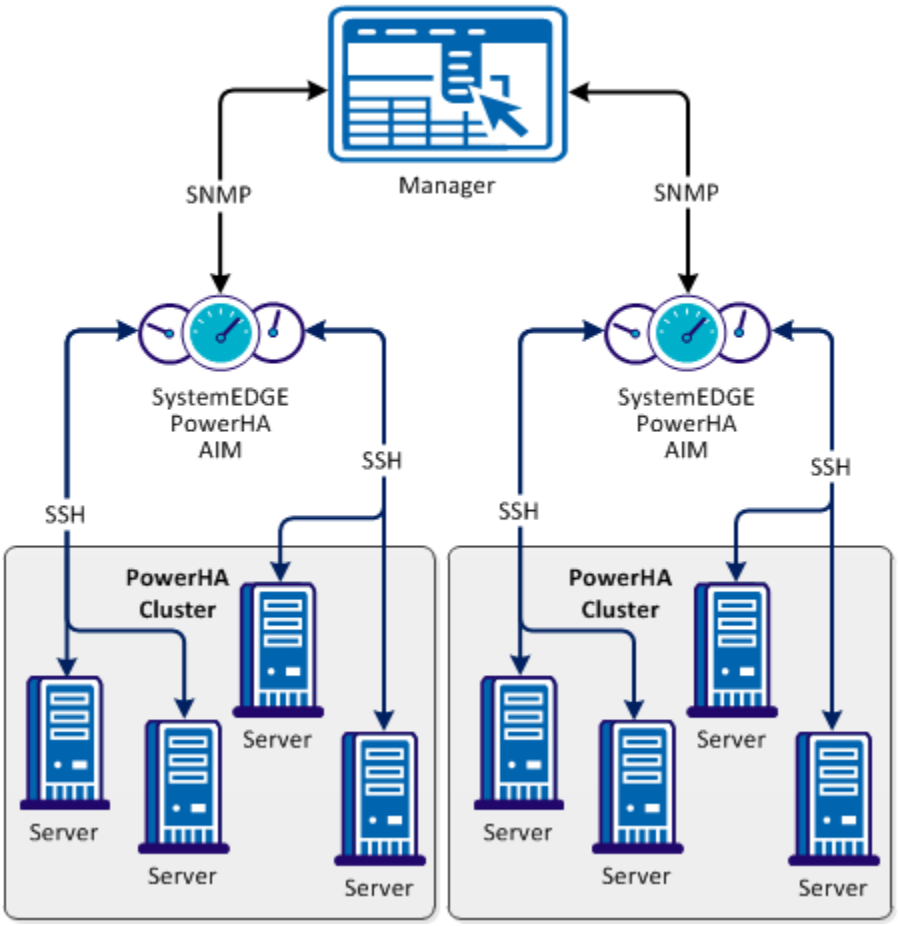

To add the required connection information for each required IBM PowerHA cluster, use the following method:

- NodeCfgUtil.exe utility on the AIM Server

The connection information is written to the configuration file on the managed node. The PowerHA AIM polls the configuration file and starts monitoring your IBM PowerHA environment through the master node.

Configure SSH

To monitor cluster nodes, configure SSH for remote access.

Follow these steps:

1. Install and run the SSH daemon on the cluster (nodes).
2. Configure local firewall to allow SSH connections.

Configure PowerHA AIM with NodeCfgUtil in Dialog Mode

The *NodeCfgUtil.exe* is a utility that lets you modify the AIM configurations. Use the utility in dialog mode to configure which nodes the appropriate AIM manages.

Follow these steps:

1. Log in as Administrator and open Windows Explorer on the computer on which the AIM is installed.
2. Change to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory, and start NodeCfgUtil.exe.

   NodeCfgUtil discovers and lists the installed AIMs in subsequent dialogs.
3. Enter *1* to add a new managed node.
4. Follow the on-screen instructions to complete the configuration. Each node requires a valid user name and password for authentication.
5. After the configuration, enter 0 to return to previous menus, or to exit the utility.

NodeCfgUtil writes a configuration file for PowerHA(hacmp.cfg) to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory. You can also use the NodeCfgUtil utility to edit or remove existing entries.

Example

The following example shows the Install Managed Node dialog for *mycluster* that has been successfully added to the configuration of the PowerHA AIM. The PowerHA AIM is a multi-instance AIM. You can repeat this procedure and can add more entities that you want to manage with this AIM.

```
****  Choose Managed Node  ****
1. Microsoft Cluster
2. IBM PowerHA
0. Go Back to Previous Menu
*******************************
Enter choice: 2
Enter following information for the IBM PowerHA Node...
(At any point to go back to the previous menu, Enter 'CTRL Q').
```

```
1. Cluster Name: mycluster
2. User Name: administrator
3. Password: ************
4. Port [default=22]:
CAAC1016 Authenticating, please wait...
CAAC1019 Authentication SUCCESSFUL.
CAAC1023 Added Node Successfully.
Press any key to continue...
```

Configure PowerHA AIM with NodeCfgUtil in Command Mode

The NodeCfgUtil.exe is a utility that lets you modify the AIM configurations. When you use the utility in command mode, you can only add managed nodes to an AIM configuration.

Note: Run NodeCfgUtil.exe as Windows Administrator.

This command has the following format:

```
(1) nodecfgutil -help
(2) nodecfgutil powerha -u user -p password -h cluster_name [-t port]
```

-help

Displays usage information about the console.

powerha

Specifies the virtual or physical environment.

-u user|usercertificate

Specifies the name of an administrative user or the user certificate, accordingly.

-p password

Specifies the password of that user.

- h cluster_name

Specifies the name of the cluster.

-t port

(Optional) Specifies the port number.

Default: 22

Follow these steps:

1. Open a command prompt on the system on which the AIM is installed.

   The command prompt appears.

2. Enter *one* of the following commands:

   ```
   (1) nodecfgutil -help
   ```

   ```
   (2) nodecfgutil powerha -u user -p password -h cluster_name [-t port]
   ```

   (1) Displays the usage information about the console.

   (2) Authenticates and stores the passed credentials for IBM PowerHA

   The utility writes a configuration file for IBM PowerHA(hacmp.cfg) to the S*ystemEDGE_InstallLpath**plugins**AIPCommon* directory.

CA IBM SystemEDGE PowerHA AIM Traps

CA SystemEDGE PowerHA AIM Trap Types

The following list provides the trap types for CA SystemEDGE PowerHA AIM. Refer to the MIB file for complete varbind descriptions.

hacmpAimInstanceAddedTrap

Sends a trap when a new Instance or Server is added.

Trap ID: 165800

hacmpAimInstanceRemovedTrap

Sends the trap when an Instance or Server is removed.

Trap ID: 165801

hacmpAimInstanceDataStatusChanged

Sends a trap when Instance or Server Data Status is changed.

Trap ID: 165802

hacmpAimNodeAddedTrap

Sends a trap when a node is added.

Trap ID: 165803

hacmpAimNodeRemovedTrap

Sends a trap when a node is removed.

Trap ID: 165804

hacmpAimResourceGroupAddedTrap

Sends a trap when a resource group is added.

Trap ID: 165805

hacmpAimResourceGroupRemovedTrap

Sends a trap when a resource group is removed.

Trap ID: 165806

hacmpAimResourceGroupMigration

Sends a trap when a resource group is migrated.

Trap ID: 165807

hacmpAimResourceAddedTrap

Sends a trap when Instance or Server a resource is added.

Trap ID: 165808

hacmpAimResourceRemovedTrap

Sends a trap when a resource is removed.

Trap ID: 165809

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.

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.

How to Configure Microsoft Cluster Service Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the Management Components

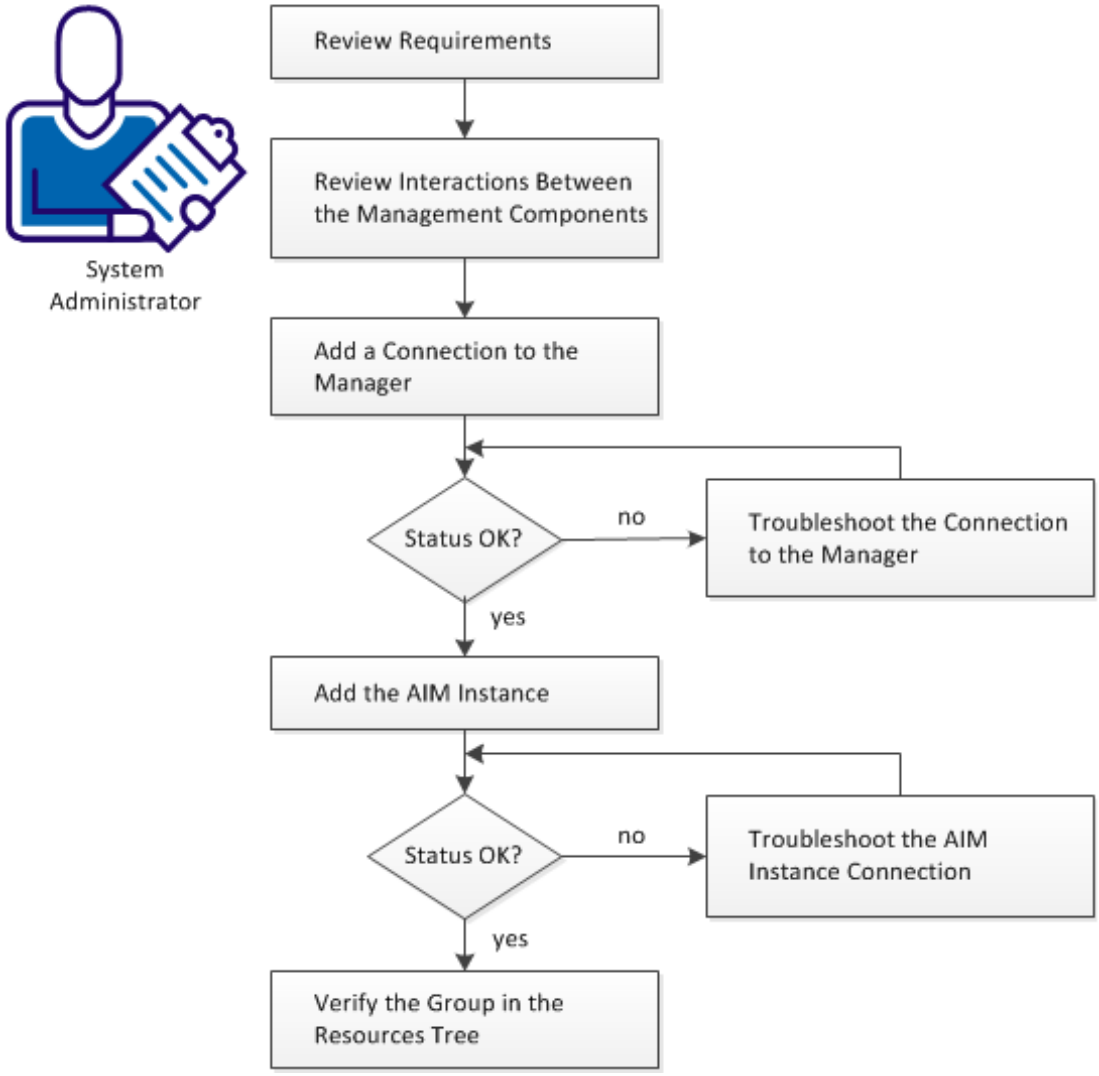

The Microsoft Cluster Service (MSCS) connects two or more servers together and shows them as a single computer to clients. Clustering helps you to have a fail-safe application. A cluster aware application such as Microsoft SQL Server runs on one node at a time. If that node goes down, another node takes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

Follow these steps:

Review Requirements

Review the following requirements before configuring the management components of CA Virtual Assurance:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Virtual Assurance and SystemEDGE.
  - You can access a CA Virtual Assurance manager installation that includes:
  - Platform Management Module (PMM)
  - Application Insight Module (AIM)
  - Monitoring Agent (SystemEDGE)
- You can access the CA Virtual Assurance user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Virtual Assurance manager discovered remote AIM Servers that you want to use.

Interactions Between MSCS Management Components

The following diagram illustrates how the components involved in MSCS monitoring interact. SystemEDGE and the MSCS AIM run on a Windows Server.

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.

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.

Interaction Between MSCS Management Components

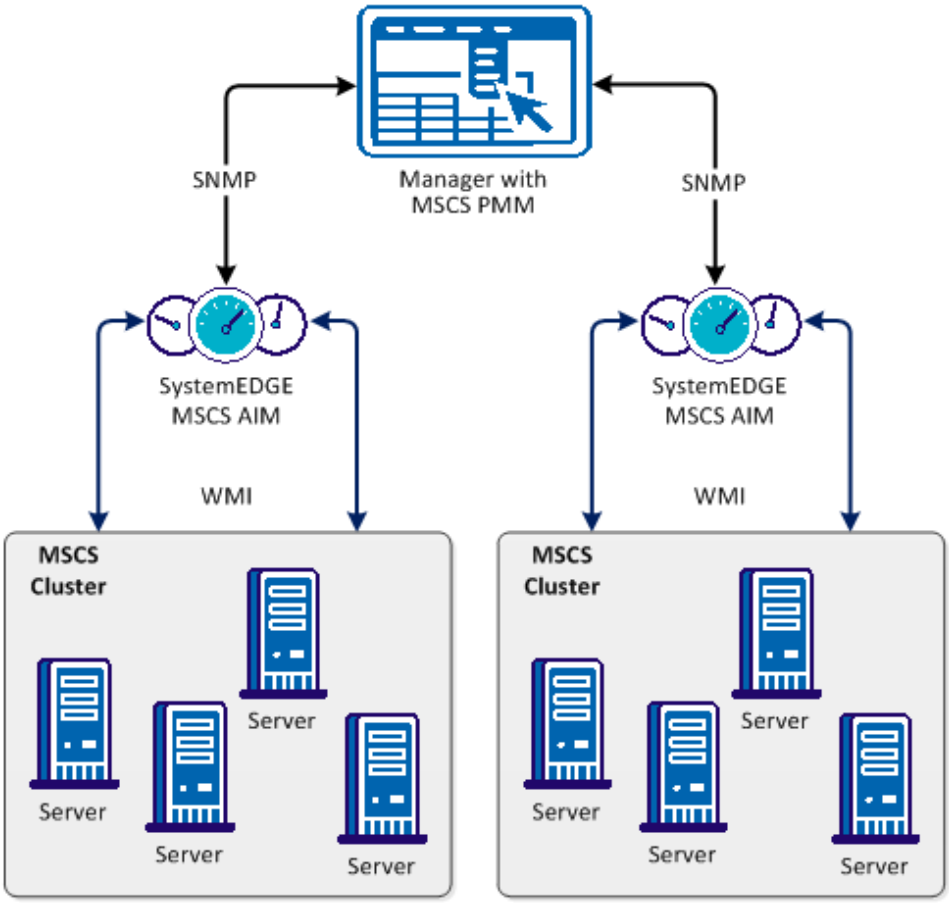

Add a Microsoft Cluster Service to the Manager

You can add a Microsoft cluster 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 Microsoft Cluster Services from the Provisioning section in the left pane.

3. Click  (Add) on the Microsoft Cluster Service pane toolbar.

   The Register New Cluster dialog appears.

4. Enter the required connection data (server name, user, password, port), specify the preferred AIM, enable Managed Status.

5. Click OK.

   The Microsoft Cluster is registered.

   When the network connection has been established successfully, the Server is added to the top right pane with a green status icon.

   Note: If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Manager Connection to the Server Fails

Symptom:

After I have added a server connection under Administration, Configuration, the validation of the connection to the server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used server connection data is still valid. If necessary, update the connection data.
- Verify, if the server system is running and accessible.
- Verify, if all services that are required for the connection are running properly on the server system.

To update the server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 server cannot be established, continue with the next procedure.

To verify if the server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. To find out whether the server has a valid DNS entry and IP address, verify the output of these commands.

   If the server is not in the DNS, add the server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>`

   Enter the correct IP address and server name and save the file. For example:

   `192.168.50.50 myServer`
4. Change to the CA Virtual Assurance user interface, Administration tab, Configuration, Server pane, and click  (Validate) in the upper-right corner.

   If the server credentials and connection data are correct and you can ping the server, the connection can still fail. In this case, it is possible that the server causes the problem. If the connection to the server cannot be established, continue with the next procedure.

To verify, if all services that are required for the connection are running properly on the server system:

1. To access the server, contact the system administrator.
2. Log in to the server system.
3. Verify, if all services that are required for the connection are running properly.
4. If necessary, start or restart the service.
5. Change to the CA Virtual Assurance user interface, server pane on the manager system and click  (Validate) in the upper-right corner.

   CA Virtual Assurance validates the server connection.

   If the connection to the server fails, verify the validity of the data you gathered according to the requirements for this scenario.

   Work with the administrator or support to fix the server connection problem.

Add the Discovered MSCS AIM Instance

After adding a Microsoft Cluster Service connection to the CA Virtual Assurance manager, add the AIM instance to manage the Microsoft Cluster Service environment.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.
2. Select Microsoft Cluster Service from the Provisioning section in the left pane.

3. Click (Add) on the Discovered Microsoft Cluster AIM Instances pane toolbar.

   The Add Cluster AIM Instance appears.

4. Select the AIM Host from the drop-down list.

   The list of discovered AIM Hosts appears.

5. Select the Registered Cluster from the drop-down list.

   CA Virtual Assurance populates the Registered Cluster drop-down list with the Cluster Names listed in the Registered Microsoft Clusters pane. You can only manage those clusters for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

   A new AIM instance for the selected cluster is added. If the instance is not in an error or in a stopped state, CA Virtual Assurance starts to discover the associated environment. When the discovery process is complete, you can start managing your Microsoft Cluster Service environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the number of managed objects that are related to virtual and physical resources in your environment. You can move the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Virtual Assurance adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance 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 AIM:

- Verify that the AIM Server is accessible.
- Verify that SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping servername`

2. Verify that the output of the commands has a valid DNS entry and IP address for the AIM server.

   If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 servername`

   Enter the correct IP address and AIM server name. For example:

   `192.168.50.51 myAIM`

4. Click (Validate) in the upper-right corner of the 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 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, AIM Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the AIM Server connection.

   If the error status remains unchanged, verify that the data you gathered is according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Virtual Assurance discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Virtual Assurance manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Microsoft Cluster Service in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand MSCS group.

   The MSCS resources appear.

CA Virtual Assurance is now ready to manage the configured MSCS environment. You can monitor the status and the properties of MSCS resources.

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 Cluster Services.

   The Microsoft Cluster Services 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.

Follow these steps:

1. Click Administration.

   The Administration page appears.

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

   The Microsoft Clusters Services page appears.

3. In the Registered Microsoft Clusters section, select the cluster that you want to remove.

4. Click - (Delete) on the Registered Microsoft Clusters toolbar.

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 Cluster Services.

   The Microsoft Cluster Services 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.

Microsoft Cluster Service Management

Microsoft Cluster Service Management lets you manage your Microsoft clusters, services and applications, and nodes. The Microsoft Cluster Service is the central location from which you can view all clusters and perform management operations.

This section describes the management operations that you can perform on Microsoft Cluster resources from the Resources page. The Resources page lets you view basic information and details about the following objects:

- Microsoft Clusters
- Services and Applications
- Nodes

Click Resources, open the Explore pane, and select one of the cluster resources; then click Summary for the resource.

The Summary page lets you view information associated with that object and events associated with the resource.

Monitor MS Cluster Services

You can monitor the status and the properties of MS cluster resources in detail.

To monitor cluster resources

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Expand the MS Cluster Service folder and click the cluster object.

   A list of cluster nodes and services object appears.

4. Click the Services and Applications object.

   A list of service appears.

5. Click the service object.

   The right pane displays general information, resources, and events.

   The General Information panel displays the service name, status, and the name of the cluster it belongs to.

   The Overview tab in the Resources panel displays resource details such as the resource name, type, and status. The Private Properties tab in the Resources panel displays private properties of each resource.

   The Events panel displays the current events.

Chapter 9: Agent-less Monitoring

CA Virtual Assurance provides agent-less monitoring of the supported virtual environments (except Hyper-V) and Windows systems (Remote Monitoring).

This section contains the following topics:

Remote Monitoring

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.

More information:

Interaction Between Remote Monitoring Components

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 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

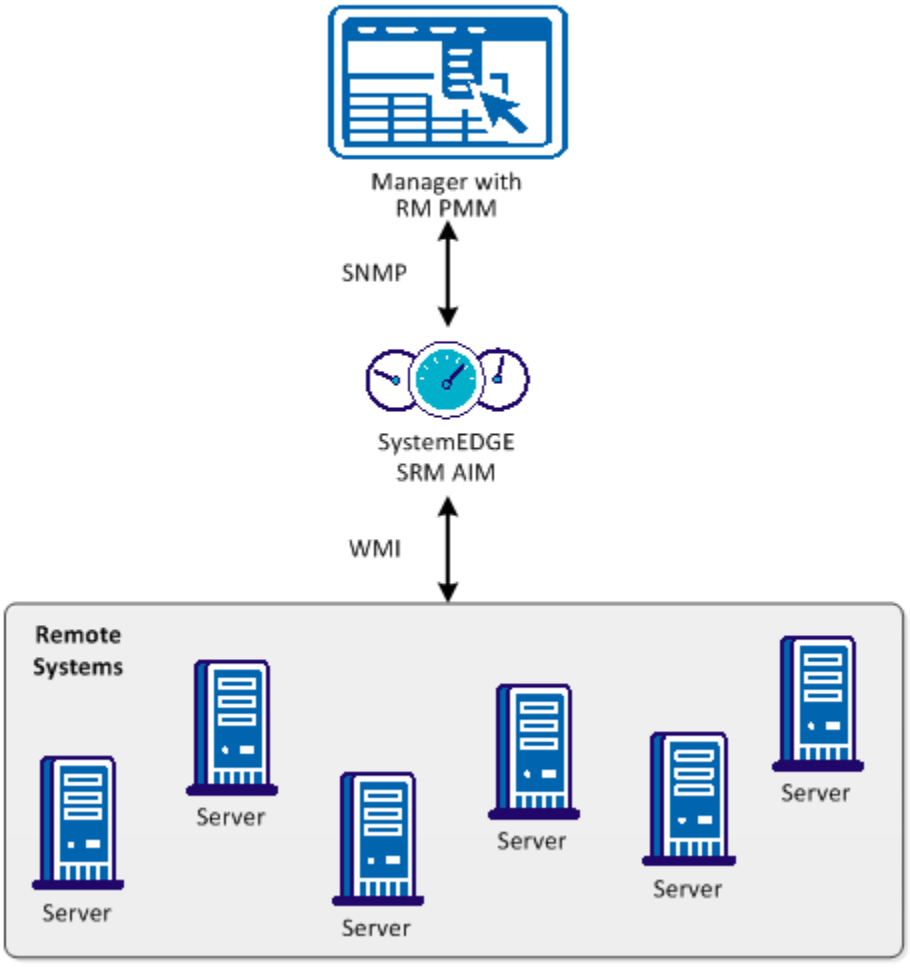

Advantages of Remote Monitoring

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 eHealth and Spectrum 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

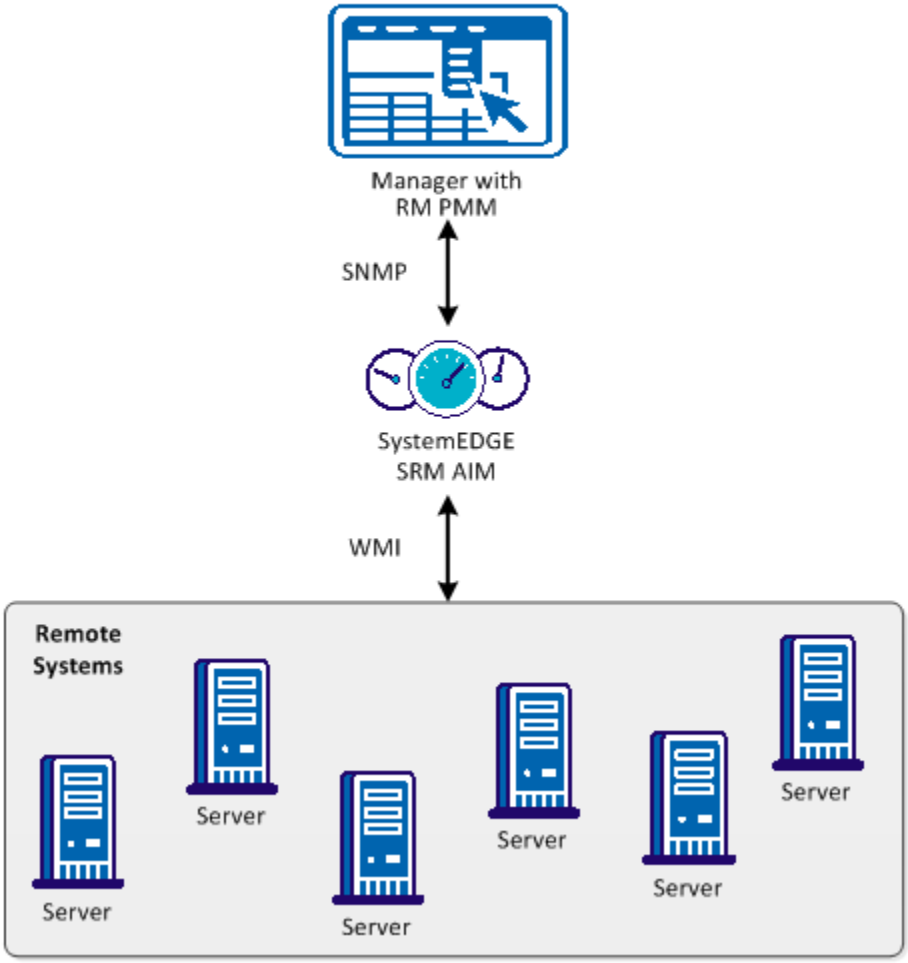

Use Case Scenario

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.

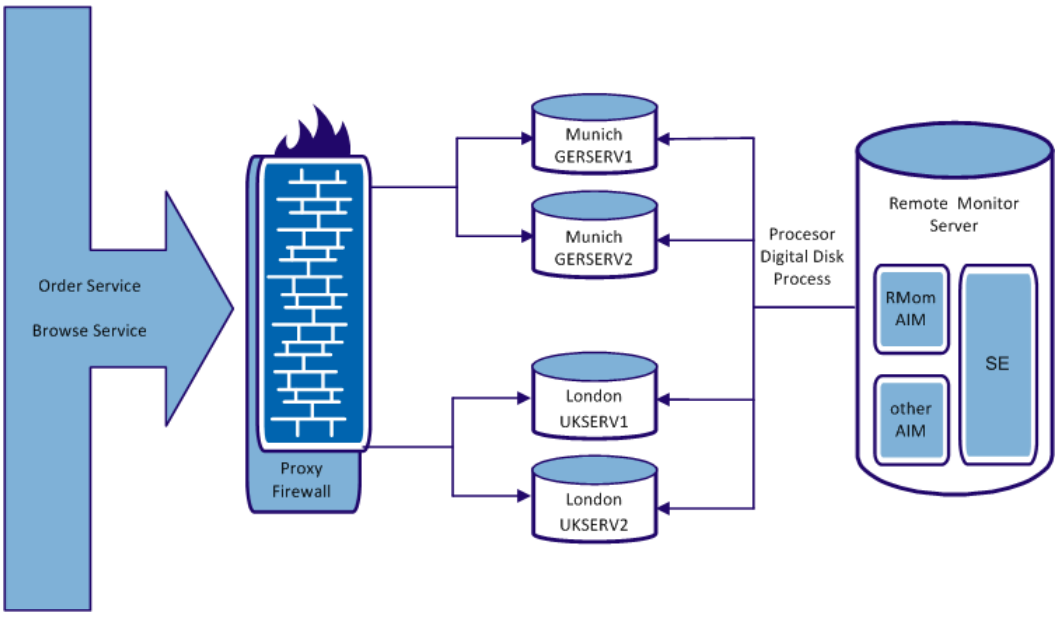

Monitoring of the services depends whether it is an order service or a browse service. In this example, two configuration sets (one for each service type) are defined. They comprise queries and monitors for the following information:

- 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.

Metric	Config Set
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 *Performance Metrics Reference*.

Create Configuration Sets

Remote Monitoring provides several out-of-the-box configuration sets that should not be changed. Use the Configuration Sets page to create custom configuration sets to suit your needs.

To create a configuration set

1. Click + (Create New).

   The Details of Individual Configuration Set pane appears.

2. Enter a name for the new configuration set, enter a description, and highlight the configuration sets to include in the new set (press the Ctrl key to highlight more than one entry).

3. Click Save.

   The new configuration set is added to the Config set list.

Note: You can also use the Actions drop-down list to clone and delete your custom configuration sets.

Support for Remote Monitoring Metrics

CA Virtual Assurance collects metrics and generates reports based on a fixed set of RM metrics in the default configuration set.

As a result, assign the default configuration set (or a configuration set or group of sets containing those metrics) to all systems for which you want to use reports.

The supported default configuration set metrics are as follows:

- CPU_PercentIdle
- Disk_PercentIdle
- Event_SystemErrors
- Mem_PercentUsed
- FSys_FreeMB
- Fsys_FreeMBDecrease
- Net_QueueLength
- Proc_PercentCPU
- Proc_PercentMemory
- Srvc_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

For the Dashboard, the following RM modules are available:

- CA SystemEDGE Machines Status
- CA SystemEDGE Objects Status

Add Remote Systems for Monitoring

Use the Systems page to enter system information for systems you would like to remotely monitor.

To add a system

1. Click + (Create New).

   The Create New pane appears.

2. Enter the name of the system you want to monitor remotely in the RM System name field and edit the following settings (if necessary):

   RM System Name

   Specifies the name of the RM System. Using user interface, you must enter RM System name in FQDN notation only, for example "vm1234.ca.com”. Using "rmonwatch" utility, you can also specify RM system name by Short name or IP Address.

   Status

   Specifies whether the system is active or in maintenance.

   Protocol

   Specifies whether the protocol is DCOM or SOAP.

   Max instances

   Specifies the maximum number of instances created in the instance table by any query to this system.

   Credentials

   Specifies the user credentials for the remote system.

   Config sets

   Specifies the config set (or group of metrics) that will be collected for the remote system.

3. Click Save.

   The system is added to the list of systems you are remotely monitoring.

Viewing Query Results

You can use the Queries page to view the query results associated with RM systems.

The Query page lets you perform the following actions:

- View detailed query results and settings (highlight a query in the Queries table and select Results or Settings).
- Filter query results based on system, status, configuration set, or specific query (use the binoculars to show or hide the query filters).
- Manage the information that displays in the query table (click a column header to sort columns in ascending or descending order and to add or remove columns).

Managing Credential Settings

You can use the Credentials page to manage the individual credential settings associated with RM systems.

The Credential page lets you perform the following actions:

- Add credentials (use the Create New (+) icon, enter settings in the Details of Individual Credentials pane, and click Save).
- Delete credentials (highlight existing credentials, and click (-) icon.
- Edit credentials (highlight existing credentials, update the settings in the Details of Individual Credentials pane, and click Save).

Managing Configuration Entries

Use the Configuration Entries page to view and manage the configuration settings associated with queries.

To view or manage configuration settings

1. Highlight a query in the Configuration Entries table. You can apply filters to the entries for configuration set, severity, query class, and escalation severity using the show and hide filter icon (binoculars).

   The Details of Individual Configuration Entry pane appears.

2. View or update the following values and click Save.

   General Settings

   Index

   Specifies a unique index for this configuration entry within the configuration set.

   Config set

   Specifies the name of the configuration set (do not use ',').

   Query name

   Specifies the name of the query (cannot include the '.' symbol).

   You can use the same query name in a different config set; however, when applying more than one config set to a system, ensure the uniqueness of all query names. If the Qualifier is set to fixed entry, you cannot rename the query.

   Description

   Specifies the description for the configuration entry.

Interval

Specifies the interval (in seconds) between successive executions of the query and evaluations of the monitor (the value must be a multiple of 30 seconds).

Query class

Specifies the query class for the configuration entry.

Query scope

Specifies the scope to apply to the query.

Query property

Specifies the property of the Query class.

Associated Monitors Definition

Obj. Class

Specifies the class name to use for the SysEDGE object state model (do not use '*').

Obj. Instance

Specifies the instance name to use for the SysEDGE object state model (do not use '*').

Obj. Attribute

Specifies the attribute name to use for the SysEDGE object state model (do not use '*').

Lag

Specifies the number of times the threshold (escalation) condition must be met to cause a state change in the SysEDGE object state model.

Result

Specifies the result attribute of the query in the query table or instance table to monitor with SysEDGE monitors.

Condition

Specifies the condition for comparing the result attribute value to the threshold and escalation threshold.

Threshold

Specifies the threshold that the result attribute value is compared against.

Severity

Specifies the severity to use for the SysEDGE object state model if the threshold condition is met.

Escal. delta

Specifies the difference to the threshold required to indicate an escalation condition.

Escal. severity

Specifies the severity to use for the SysEDGE object state model if the escalation condition is met.

Advanced Parameters

Query depends

Specifies that a query (Q2) depends on another query (Q1); such that Q2 is only created based on the result of Q1.

Query total

Specifies the property of the query class to apply as a total reference.

Query scale

Specifies the scale to apply for the property value (for example, *100, /1024 or /1024*100) This scale is used as default for the query scale in the query table. The value of the query property is multiplied or divided by the scale before storing the value in the result attributes.

Instance Key

Specifies the properties of the query class to use for instance naming in the instance table.

Qualifiers

Specifies additional information related to the configured queries and monitors. The possible values are as follows:

- Entry cannot be deleted and query name cannot be changed (fixed entry)
- Query is executed only once (at least one time successful)
- Query is no longer executed if it was unsuccessful
- Query is not shown in the query table
- Results are shown per instance
- Results show previous values instead of current ones
- Results show increasing delta values
- Results show decreasing delta values
- Aggregate Monitors with the same object data and severity as AND relation

The configuration settings are updated to reflect any changes.

Example

To monitor the free space on disk "C:" and generate the events when the space is less than 10GB or 5GB, set the following values:

General Settings

Index: 1

Config set: test

Query name: FreeSpace

Interval: 30

Query class: Win32_LogicalDisk

Query scope: DeviceID = "C:"

Query property: FreeSpace

Associated Monitors Definition

Obj. Class: Disk

Obj. Instance: C:

Lag: 0

Result: Minimum

Condition: <

Threshold: 10

Severity: Minor

Escal. delta: -5

Escal. severity: Major

Advanced Parameters

Query scale: /(1024*1024*1024)

Qualifiers: 0x0

Chapter 10: Install and Configure Active Directory and Exchange Server AIM

This section contains the following topics:

Introduction

The Active Directory and Exchange Server (ADES) AIM lets you monitor the health states and Key Performance Indicator (KPI) metrics of Active Directory and Exchange Server environments. The ADES AIM features include:

- Monitor the message records manager, logical disk usage, and logical disk read/write of the mailbox server.
- Monitor the network latencies, queue, mail delivery metrics, logical disk usage, and logical disk read/write of the hub transport server.
- Monitor the active directory performance, replication, logical disk usage, and logical disk read/write.

The ADES AIM collects the following data for monitoring:

- Configuration data from Active Directory and Exchange Server
- Performance data from Active Directory and Exchange Server

ADES AIM Scalability

When planning for the ADES AIM deployment, consider the following key factors that have an impact on the infrastructure sizing and system performance:

- Available memory for the ADES AIM, excluding the memory that operating system and other applications uses:
  - Host with 1-GB free memory can monitor 20 hosts (2 Active Directory hosts and 18 Exchange hosts).
  - Host with 2-GB free memory can monitor 40 hosts (6 Active Directory hosts and 34 Exchange hosts).
  - Host with 3-GB free memory can monitor 60 hosts (10 Active Directory hosts and 50 Exchange hosts).
- Geographic distribution of the environment:
  - When the ADES AIM is in geographical proximity, it reduces the time to discover and poll the environment.
  - High latency or packet loss can cause the AIM not to obtain all the data that is required.

Note: The sizing information is an approximate estimate of the deployment requirements and it is not definitive. The sizing information varies according to the monitoring environment.

Install the ADES AIM

Complete the following tasks to install ADES AIM:

1. Install the CA SystemEDGE Release 5.9 agent and CA Advanced Encryption Release 5.9.
2. Install the ADES AIM using one of the following methods:
   - Deploy through the CA Virtual Assurance Remote Deployment.
   - Install manually through command mode.
3. Configure the ADES AIM by specifying the domains to monitor:

Notes:

- When using CA Spectrum with ADES Manager, do not install the SpectroSERVER on the host that manages the ADES AIM host. Also, the ADES AIM must be the only AIM installed on the SystemEDGE host.
- Install the SystemEDGE and ADES AIM on a Windows host that is a member server in one of the domains, with a trust relationship to the other domains.
- The SystemEDGE agent and ADES AIM host must not have any Active Directory or Exchange Server roles.

Deploy the ADES AIM Using Remote Deployment

Create a software job to install the ADES AIM on the host using the CA Virtual Assurance Remote Deployment.

Follow these steps:

1. Log in to the CA Virtual Assurance application and go to the management view.
2. Find the host that you want to deploy the ADES AIM in the Resource tab.
3. Create a job and select the platform type as Windows and the available wrapper packages are displayed.

   Specify the following parameters in the wrapper package when creating the job:

   User

   Defines the name of a domain administrator without the Fully Qualified Domain Name (FQDN). For example, adminuser.

   Password

   Defines the password of the user.

   Domainname

   Defines the name of the domain that is monitored through the ADES AIM. Enter the FQDN.

Management Entity

Specifies which hosts to manage, based on the technology.

0

Monitor Active Directory hosts only.

1

Monitor Exchange Server hosts only.

2

Monitor both Active Directory and Exchange Server hosts.

Management Mode

Specifies which hosts to manage.

0

Discover and monitor all hosts in the domain automatically that the management entity defines (Domain-based management).

Note: Hosts of child domains are not monitored automatically.

1

Discover all the hosts in the domain but monitor only the hosts that are configured through the manager (Host-based management).

4. Select the necessary packages and deploy them on the host.

   Verify the job status from the jobs panel. If the job fails, redeploy the package again.

Note: For more information, see How to Deploy SystemEDGE and AIMs (see page 117).

Install the ADES AIM in Command Mode

Installing in command mode installs the ADES AIM on a host without using Remote Deployment.

Note: Verify that CA SystemEDGE Release 5.9 and CA Advanced Encryption Release 5.9 are installed on the host before you install the ADES AIM.

Follow these steps:

1. Go to *DVD1*\Installers\Windows\Data\SysMan and copy the following zip files to your local disk:
   - CA_SystemEDGE_ESAD-Windows.zip
   - CA_SystemEDGE_ESAD-Windows-metadata.zip
2. Extract the zip files that are copied to your local disk. The following files are available at the extracted location:
   - caesadaimx64.msi
   - ca-setup.exe
   - ca-setup.dat
3. Open the Command Prompt window and go to *Extracted_Path*\CA_SystemEDGE_ESAD\5.8.0\ENU\Windows_x64.
4. Run ca-setup.exe to install the ADES AIM. The command has the following format:

   ```
   ca-setup EULA_ACCEPTED="[yes|no]"
   CASE_ESAD_DOMAIN_NAME="domain_name"
   CASE_ESAD_DOMAIN_USER_NAME="username@fqdn"
   CASE_ESAD_DOMAIN_PWD="password"
   CASE_ESAD_MANAGEMENT_ENTITY="[0|1|2]"
   CASE_ESAD_MANAGEMENT_MODE="[0|1]"
   ```

   EULA_ACCEPTED="[*yes*|*no*]"

   Specifies whether the license is accepted or not.

   CASE_ESAD_DOMAIN_NAME="*fully_qualified_domain_name*"

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

   CASE_ESAD_DOMAIN_USER_NAME="*username@fqdn*"

   Specifies the name of a user with Domain Administrator and Exchange Organization Administrator or Organization Management privileges.

   CASE_ESAD_DOMAIN_PWD="*password*"

   Specifies the password of the user.

   CASE_ESAD_MANAGEMENT_ENTITY="[*0*|*1*|*2*]"

   Specifies which hosts to manage, based on the technology.

0

Monitor Active Directory hosts only.

1

Monitor Exchange Server hosts only.

2

Monitor both Active Directory and Exchange Server hosts.

CASE_ESAD_MANAGEMENT_MODE="[*0*|*1*]"

Specifies the hosts to manage.

0

Discover and monitor all hosts in the domain automatically that the management entity defines (Domain-based management).

Note: Hosts of child domains are not monitored automatically.

1

Discover all the hosts in the domain but monitor only the hosts that are configured through the manager (Host-based management).

5. Restart the SystemEDGE service to run the ADES AIM.

Example

The following example shows how to install the ADES AIM on a host and monitor the domain mydomain.com.

```
ca-setup EULA_ACCEPTED="yes"
CASE_ESAD_DOMAIN_NAME="mydomain.com"
CASE_ESAD_DOMAIN_USER_NAME="adminuser@mydomain.com"
CASE_ESAD_DOMAIN_PWD="domainpass123" CASE_ESAD_MANAGEMENT_ENTITY="2"
CASE_ESAD_MANAGEMENT_MODE="0"
```

How to Configure Active Directory and Exchange Server Monitoring

Follow these steps:

Requirements to Configure Active Directory and Exchange Server

The following prerequisites are necessary to install and configure the ADES AIM:

General requirements

- Knowledge to discover the server and deploy a package through CA Virtual Assurance.
- Required privileges:
  - User account with permissions for Remote Deployment.
  - User account with Local Administrator privileges on the host for manual installation.
  - Domain Administrator and Exchange Organization Administrator or Exchange Organization Management privileges for monitoring the domain.

    Note: Verify that Domain Administrator and Exchange Organization Administrator privileges are assigned to the same user.

Software Requirements

- Supported operating environments for ADES AIM host:
  - Windows 2008 Server SP2
  - Windows 2008 R2 SP2 x64
- Supported Domain Controller operating environments:
  - Windows 2008
  - Windows 2008 R2
- Supported Exchange Server versions:
  - Exchange 2007 SP3
  - Exchange 2010 SP2

  Note:
  - Monitoring Exchange 2003 hosts is not supported.
  - Monitoring Exchange 2007 hosts across the forests is not supported.
- Required applications:
  - .Net 3.5 or higher version
  - Windows PowerShell 2.0
  - Exchange 2007 Management Tools SP3 for monitoring Exchange 2007 hosts
  - CA SystemEDGE Release Release 5.9 and CA Advanced Encryption Release 5.9

How the Active Directory and Exchange Server AIM Works

The following diagram illustrates the ADES AIM architecture:

Interaction Between Active Directory and Exchange Server Management Components

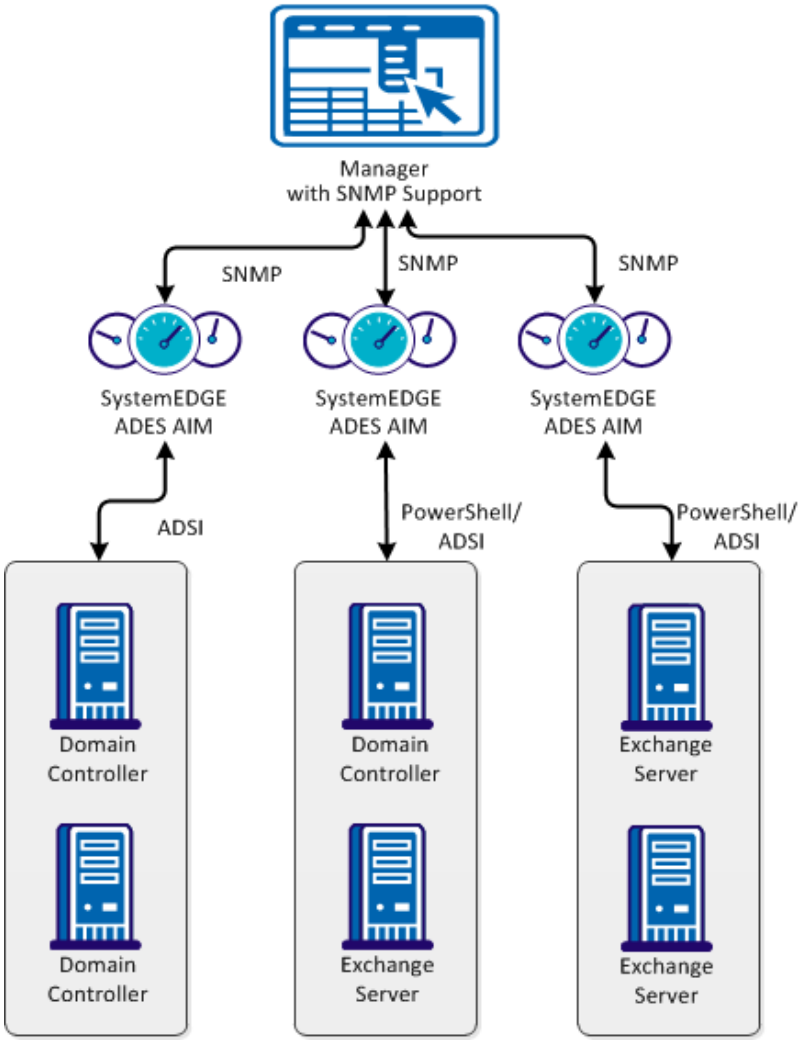

The following process explains how the ADES AIM works:

1. The ADES AIM discovers hosts by searching the Domain Controller. The ADES AIM collects the information about:
   - Active Directory server roles such as Domain Controller and Global Catalog.
   - Exchange Server roles such as Hub Transport, Mailbox, and Client Access Server.

   Note: Unified Messaging and Edge Transport roles are not supported for monitoring.
2. When the hosts are discovered, the AIM sends a message to collect the data from:
   - The Domain Controller using ADSI calls
   - The Exchange Server using PowerShell commands
3. The AIM receives the data and updates the MIB table for the SystemEDGE Agent.
4. The managers such as the CA eHealth and CA Spectrum, poll the SystemEDGE host and collect the data to display.
5. The AIM continually polls the managed hosts (Active Directory and Exchange Server hosts that are set for monitoring) and updates its MIB table.

Configure the Environment to Enable ADES AIM Monitoring

Apply the PowerShell configuration settings on the Exchange hosts to enable the ADES AIM to monitor a domain.

Note: Configure every Exchange Server before you start monitoring.

Follow these steps:

1. Select Start, Programs, Accessories, Windows PowerShell, Windows PowerShell (x86).

   The Windows PowerShell command prompt appears.

2. Run the following command to manage the host remotely through WinRM services:

   ```
   Enable-PSRemoting
   ```

   WinRM setup initiates remote management and creates a WinRM listener to accept WS-Man requests.

3. Run the following command to add hosts to the list of trusted hosts:

   ```
   Set-Item WSMan:Localhost\Client\TrustedHosts -Value * -Force
   ```

4. Run the following command to restart the WinRM service:

   ```
   Restart-Service WinRM
   ```

   The TrustedHosts settings are updated and the Exchange Server is available for monitoring.

Add a Domain Server or Exchange Server to the Manager

You can add a Microsoft Active Directory Domain Controller or Exchange Server connection to the manager using the 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 Microsoft Active Directory and Exchange Server from the Provisioning section in the left pane.

3. Click  (Add) on the Servers pane toolbar.

   The Add Server dialog appears.

4. Enter the required connection data (server name, user, password, mode, technology), specify the preferred AIM, enable Managed Status.

5. Click OK.

   CA Virtual Assurance validates the submitted connection data and tries to establish a connection to the server.

   When the network connection is established successfully, the Server is added to the top right pane with a green status icon.

   Note: If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Virtual Assurance adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Server Connection to the Manager Failed

Symptom:

After I have added a Server connection under Administration, Configuration, the validation of the connection to the Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Server connection data is still valid. If necessary, update the connection data.
- Verify, if the Server system is running and accessible.
- Verify, if the Management Service on the Server system is running properly.

To update the Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, 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 Server cannot be established, continue with the next procedure.

To verify if the Server system is running and accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `nslookup <Server Name>`
   `ping <IP Address of Server>`
2. Verify the output of the commands to find out whether the Server has a valid DNS entry and IP address.

   If the Server is not in the DNS, add the Server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 <Server Name>

   Enter the correct IP address and Server name. For example:

   `192.168.50.50 myServer`
4. Click  (Validate) in the upper-right corner.

   If the Server credentials and connection data are correct and you can ping the Server, the connection can still fail. In this case, it is possible that the Server causes the problem. If the connection to the Server cannot be established, continue with the next procedure.

To verify, if the Management Service on the Server system is running properly:

1. Contact the Administrator to access the Server system.
2. Log in to the Server system and open Administrative Tools, Services from the Start menu.

   The Services window opens.
3. Select the service and start or restart the service.
4. Change to the CA Virtual Assurance user interface, Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Virtual Assurance validates the Server connection.

   If the connection to the Server fails, verify the validity of the data you gathered according to the requirements for this scenario.

   Work with the administrator or support to fix the Server connection problem.

Add the ADES AIM Instance

After adding an Active Directory and Exchange Server connection to the CA Virtual Assurance manager, add the AIM instance to manage the environment.

Follow these steps:

1. Open the CA Virtual Assurance user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.
2. Select Microsoft Active Directory and Exchange Server from the Provisioning section in the left pane.

3. Click  (Add) on the AIM Servers pane toolbar.

   The Add AIM Server dialog appears.

4. Select the AIM Host from the drop-down list.

   The list of discovered AIM Hosts appears.

5. Select the server from the drop-down list.

   CA Virtual Assurance populates the server drop-down list with the server names listed in the Servers pane. You can only manage those servers for which your CA Virtual Assurance manager has a valid connection established.

   Note: If the AIM resides on a remote system, CA Virtual Assurance must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

   A new AIM instance for the selected server is added. If the instance is not in an error or in a stopped state, CA Virtual Assurance starts to discover the associated environment. When the discovery process is complete, you can start monitoring your Active Directory and Exchange Server environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the number of managed objects that are related to virtual and physical resources in your environment. You can move the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Virtual Assurance adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance 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 AIM:

- Verify that the AIM Server is accessible.
- Verify that SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Virtual Assurance manager system and run the following commands:

   `ping servername`

2. Verify that the output of the commands has a valid DNS entry and IP address for the AIM server.

   If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Virtual Assurance manager system. Continue with Step 3.

   If the 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 servername`

   Enter the correct IP address and AIM server name. For example:

   `192.168.50.51 myAIM`

4. Click  (Validate) in the upper-right corner of the 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 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, AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

   CA Virtual Assurance validates the AIM Server connection.

   If the error status remains unchanged, verify that the data you gathered is according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Virtual Assurance has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Virtual Assurance discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Virtual Assurance manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Virtual Assurance manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify Active Directory and Exchange Server Monitoring

After a successful configuration, CA Virtual Assurance starts monitoring the Active Directory and Exchange Server. Monitor the Active Directory and Exchange Server events in the user interface.

(Optional) Configure the ADES AIM using Node Configuration Utility

Using NodeCfgUtil instead of the user interface is an alternative configuration method for the ADES AIM. Configuring the ADES AIM lets you add, modify, or remove one or more domains that the ADES AIM manages. The NodeCfgUtil creates a configuration file for ADES AIM (esad.cfg), located in the *SystemEDGE_InstallPath*\plugins\AIPCommon directory.

Follow these steps:

1. Open Windows Explorer and navigate to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory.
2. Start NodeCfgUtil.exe.
3. Enter an option according to your choice. You can add, modify, or remove a domain. For example, enter 1 to add a new managed node.
4. Enter the number corresponding to the ADES AIM in the Choose Managed Node screen. For example, enter 1 to select the ADES AIM.
5. Follow the on-screen instructions to complete the configuration. Each domain requires a valid user name and password for authentication and appropriate management entity and management mode.
6. When the configuration is completed, enter 0 to save the configuration and exit the utility.
7. Restart the SystemEDGE service to apply the changes.

Example

The following example shows the Install Managed Node dialog for mydomain.net that has been successfully added to the configuration of the ADES AIM. Management Entity is set to Active Directory. Management Mode is set to domain-based.

```
**** Choose Managed Node ****

1. Microsoft Active Directory and Exchange Server

0. Go Back to Previous Menu

******************************

Enter choice: 1

Enter following information for the Microsoft Active

Directory and Exchange Server Node...

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

1. Domain Name(FQDN): mydomain.com

2. User Name(Example:adminuser@domain.com): administrator@mydomain.com

3. Password: ************

4. Management Entity(0-AD Only, 1-Exchange Only, 2-Both AD and Exchange): 0

5. Management Mode(0-domain based/automatic, 1-host based/manual): 0

CAAC1016 Authenticating, please wait...

CAAC1019 Authentication SUCCESSFUL.

CAAC1023 Added Node Successfully.

Press any key to continue . . .
```

Uninstall the ADES AIM

Uninstalling the agent removes the agent and its associated configuration data from the host.

Follow these steps:

1. Stop the SystemEDGE process using the SystemEDGE control panel.
2. Select Start, Control Panel, Programs, Programs and Features.

   The Uninstall or change a program window opens.
3. Right-click CA AIM for Exchange Server and Active Directory component and select Uninstall.

   A confirm message is displayed.
4. Click Yes.

   The ADES AIM component is removed. Verify that the ADES AIM component is not displayed in the Add/Remove control panel.

Troubleshoot Active Directory and Exchange Server

More information:

AIM is Inactive and not Collecting Data (see page 602)
One or More Domains are not Monitored (see page 602)
Some Counters are not Monitored (see page 603)
Some Hosts are not Monitored (see page 603)

AIM is Inactive and not Collecting Data

Symptom

The AIM is inactive and unable to collect data.

Solution

Verify the following:

- The caesadaim.exe process is running.
- The log file for the domain is created in the AIM directory for every configured domain.

If the process is not running or the log file is not created, restart the SystemEDGE service.

If the AIM is not running after restarting the SystemEDGE service, verify the following requirements and take appropriate action:

- .NET 3.5 SP1 Framework is installed on the AIM host.
- Exchange Management Tools 2007 SP3 is installed on the same host as the AIM (if the domain contains one or more Exchange 2007 servers).

One or More Domains are not Monitored

Symptom

The ADES AIM does not monitor one or more domains.

Solution

- Verify that a log file is created for each monitored domain in the ADES AIM folder with the name domain_AIM.log. If the log file is not created, verify that the domain is configured for monitoring using nodecfgutil.exe.
- If the log file is created for the domain, open the log file and look for the following error message:

  ```
  The specified domain does not exist or cannot be contacted.
  ```

  If this message exists in the log file, verify that communication between the ADES AIM host and the domain controller is not blocked. When the domain controller is accessible from the ADES AIM host, initiate Discovery for the AIM through CA Spectrum.

Some Counters are not Monitored

Symptom

Some of the performance counters are not monitored.

Solution

Reinitiate discovery in the ADES AIM to create the counter on the hosts where the counters do not exist.

Note: Performance counters that appear for specific configurations are monitored only when the required configuration or instance is available on the host.

Some Hosts are not Monitored

Symptom

Some Active Directory or Exchange Server hosts in the domain are not monitored.

Solution

Verify the following configurations:

- AIM is configured in domain mode or host mode.

  Note: In the host mode, change the management status using CA Spectrum or MIB browser for each of the hosts in the Universal Host Table.

- AIM is configured with a management entity to monitor only Active Directory hosts or only Exchange Server hosts. Change the value of the Management Entity option for the domain using NodeCfgUtil to 2 for the ADES AIM to monitor both the technologies.

Chapter 11: Using Rules and Actions

This section contains the following topics:

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 *Performance Metrics Reference*.

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 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.

List of Predefined Action Types

This section describes the following predefined action types that are available to create actions for policy rules.

Add Disk: VMware vCenter

The Add Disk action type lets you add a disk to a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to add a disk. Select one from the drop-down list.

Datastore

Specifies the name of the datastore associated with the ESX server for the selected VM. Select one from the drop-down list.

Drive Size

Specifies the size of the additional disk. Enter a value and select MB or GB from the drop-down list.

SCSI Controller

Specifies the SCSI controller to use to create the additional disk. Select one from the drop-down list.

Thin Provisioning check box

Specifies whether to enable thin provisioning.

Disk Mode

Specifies the disk mode. Select one of the following from the drop-down list:

- Persistent
- Independent Persistent
- Independent Non-persistent

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Add Network Interface: VMware vCenter

The Add Network Interface action type lets you add a virtual NIC to a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to add a virtual NIC. Select one from the drop-down list.

Device Type

Specifies the device type. Select one from the drop-down list.

Network

Specifies the network associated with the ESX server for the selected VM. Select one from the drop-down list.

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)

MAC Address

(Optional) Specifies a MAC address. Leave the field blank if you want the MAC address to be autogenerated.

Wake on LAN check box

Specifies whether to set the virtual NIC to wake on LAN.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Add Server to Service

The Add Server to Service action lets you add servers to an existing service.

The Details section of the action definition contains the following fields:

Service Name

Specifies the name of the service.

Server List (comma delimited)

Specifies the list of servers to add to the service.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Change Machine State: Microsoft Hyper-V

The Change Machine State action type controls the state changes of the virtual machines in your Hyper-V environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the server on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine for which to change the state. Select one from the drop-down list.

State

Specifies the desired state of the virtual machine. Select one of the following from the drop-down list:

- Turn off
- Shutdown
- Save
- Pause
- Start

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Clone Machine: Solaris Zones

The Clone Solaris Zone Machine action type configures and installs a new zone by copying the data from an existing zone. You cannot perform this operation for global zones or when a zone is in the installed state.

The Details section of the action definition contains the following fields:

Zone Host

Defines the Solaris Zones host that contains the zone to clone.

Zone

Defines the zone to clone. You can use text extracted from event messages.

Name

Defines the name of the new zone. You can use automatically generated text or text extracted from event messages.

Path

Defines the installation path of the new zone.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure CPU/Memory: IBM LPAR

The Configure CPU/Memory action type lets you set limits on CPU and memory resources allocated to a virtual machine in IBM LPAR environment.

The Details section of the action definition contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Displays the unique name for the partition.

Profile Name

Specifies the name of an existing profile for the selected LPAR.

Operations

Specifies the operation to perform. Select one of the following from the drop-down list:

- Add Memory Units
- Subtract Memory Units
- Add Processors
- Subtract Processors

Processors

Specifies the number of processors for addition or removal.

Adjustment Type

Specifies the Adjustment Type. Select one option:

- Dynamic Adjustment Only
- Dynamic Adjustment and Update Profile

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure CPU/Memory: Microsoft Hyper-V

The Configure CPU/Memory action type controls the number of CPU and memory shares allocated to a virtual machine in your Hyper-V environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the server on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine for which to change the state. Select one from the drop-down list.

CPU Allocation

Specifies the CPU Allocation of the virtual machine. Adjust one of the following from the drop-down list:

- Number of CPUs
- CPU Reserved %
- CPU Weight
- CPU Limit %
- Present CPUID

Memory Allocation

Specifies the memory share allocated to the virtual machine in megabytes.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure CPU/Memory: VMware vCenter

The Configure CPU/Memory action type lets you set limits on CPU and memory resources.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Target VM Machine

Specifies the name of the virtual machine for which to adjust resources. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Operations

Specifies the operation to perform. Select one of the following from the drop-down list:

- Set CPU Limit
- Set Memory Limit
- Set CPU Reservation
- Set Memory Reservation

MHz, MB

Enter a value appropriate to the operation you select.

Unlimited check box

Allows the operation you select unlimited use of the resource.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: Cisco UCS

This action type lets you configure a power management action for a UCS blade server.

The Details section of the action definition contains the following fields:

UCS Manager

Specifies the name of the UCS Manager

UCS Chassis

Specifies the name of the UCS Chassis

UCS Blade

Specifies the name of the UCS Blade

Power Operations

Select an operation from the drop-down list:

Cycle Immediate

Immediately power cycle the blade

Cycle wait

Power cycle the blade which notifies all applications about its shutdown

Hard Reset Immediate

Plug back to power on the blade similar to unplug the power of the blade

Hard Reset wait

Unplugs the power of the blade. Prior to unplugging, the blade notifies all applications about its shutdown

Soft Shut Down

Shuts down the blade. Prior to shutdown, the blade notifies all applications about its shutdown

Shut Down

Shuts down the blade immediately

Boot up

Boots up the blade

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: IBM LPAR

The LPAR Configure Power action type controls power settings on LPARs.

The Details section of the action definition contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Specifies the partition name to be controlled.

Operation

Specifies the power operation to perform. If you select Activate, complete the following fields in the Operation Options section:

Partition profile

Specifies the partition profile used to activate the power settings.

Key Lock

Specifies the key lock mode in the partition profile. Key Lock establishes the power-on and power-off modes allowed for the system and is either manual or normal. For security reasons, it is not recommended that you set the key lock position to manual.

Boot mode

Specifies the boot mode in the partition profile. The system uses this boot mode to start the operating system on the logical partition unless you specify otherwise when activating the partition profile. CA Virtual Assurance supports the following valid boot modes:

normal

Starts the logical partition as normal. (Use this option to perform most everyday tasks).

open_firmware

Boots the logical partition to the open firmware prompt. Service personnel use this option to obtain additional debug information.

If you select Shutdown, complete the following fields in the Operation Options section:

Delayed

Shuts down the logical partition using the delayed shutdown sequence. This sequence allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it ends abnormally. The next restart may be longer than typical.

Immediate

Shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

OS Shutdown

Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

OS Shutdown Immediate

Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

If you select Restart, select *one* option from the Operation Options section:

Partition profile

Specifies the partition profile used to restart the partition.

Immediate

Shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option causes undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

OS Shutdown

Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

OS Shutdown Immediate

Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: Microsoft Hyper-V

The Configure Power action type controls the startup and shutdown of a virtual machine in your Hyper-V environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the server on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine for which to change the state. Select one from the drop-down list.

Start Action

Specifies the action to perform when the Hyper-V Server starts. Select one of the following from the drop-down list:

- Always

  Always starts the VM when the Hyper-V Server starts.

- Auto

  Automatically starts the VM when the Hyper-V Server starts.

- None

  Does not start the VM when the Hyper-V Server starts.

Start Delay

Adjust the delay (in seconds) to start a VM after the Hyper-V Server starts. Select one from the drop-down list.

Shutdown Action

Specifies the action to perform when the virtual machine shuts down. Select one of the following from the drop-down list:

- Off

  Turns off the VM before Hyper-V Server shuts down.

- Save

  Saves (Suspends) the VM before Hyper-V Server shuts down.

- Shutdown

  Shuts down the VM before Hyper-V server shuts down.

Recovery Action

Specifies the action to regain the previous details of a virtual machine when the Hyper-V Server fails. Select one of the following from the drop-down list:

- None

  Does not take a specific action when the Hyper-V Server starts after the server fails.

- Restart

  Restarts the VM when Hyper-V Server starts after the server fails.

- Revert

  Reverts the VM with the latest snapshots when the Hyper-V Server starts after the server fails.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: VMware vCenter/Adjust vApp Power

The Configure Power action type controls power settings on the virtual machines and vApps in your VMware vCenter environment.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

VM/vAPP

Radio button to specify the target type, VM or vApp.

Target

Specifies the name of the virtual machine or vApp for which to adjust power. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Power Operation

Specifies the power operation to perform. Select one of the following from the drop-down list:

- VC Power On
- VC Power Off
- VC Power Reset
- VC Power Suspend
- VC Power Shutdown
- Power on vApp
- Power off vApp
- Suspend vApp

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Service Profile: Cisco UCS

The Configure Service Profile action type lets you associate, disassociate or failover a service profile to a UCS blade server.

The Details section of the action definition contains the following fields:

UCS Manager

Specifies the name of the UCS Manager

UCS Chassis

Specifies the name of the Cisco UCS chassis

UCS Blade

Specifies the name of the Cisco UCS blade

Service Profile

Specifies the name of the service profile

Profile Operations

Select a profile from the drop-down list:

Associate

Associates a service profile to a blade

Unassociate

Unassociates a service profile from a blade

Failover

Use this option, a check box appears against a service profile that is used to automatically failover the service profile to next available blade. By default, the check box is selected, and both chassis and blade drop-down are disabled.

Clear the check box to select the desired chassis and blade to failover a specific service profile.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Shares: VMware vCenter

The Configure Shares action type controls CPU and memory shares for virtual machines in your VMware vCenter environment.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Target VM Machine

Specifies the name of the virtual machine for which to adjust shares. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Operations

Specifies the operation to perform. Select one of the following from the drop-down list:

- Set CPU
- Add CPU
- Subtract CPU
- Set Memory
- Add Memory
- Subtract Memory

Values

Enter a value appropriate to the operation you select.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Convert Template to VM: VMware vCenter

The Convert Template to VM action type lets you convert a template to a virtual machine.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which the VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the data center where the VM is located. Select one from the drop-down list.

VC Compute Resource

Specifies the cluster or VMware ESX host where the VM is to be created. Select one from the drop-down list.

VC ESX Servers

Specifies the VMware ESX server where the VM will reside. Select one from the drop-down list.

VC Resource Pool

Specifies the name of the resource pool from which you want to select the VM for cloning. Select one from the drop-down list.

VC Template

Specifies the name of the template you want to convert. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Convert VM to Template: VMware vCenter

The Convert VM to Template action type lets you convert a powered off virtual machine to a template.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

VC Virtual Machine

Specifies the name of the virtual machine you want to convert. Select one from the drop-down list or use text extracted from event messages.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Create Event

The Create Event action type lets you create events like system discovery, system deletion, multiple system discovery, and system management status changes.

The Details section of the action definition contains the following fields:

Event Status

Specifies the status of the event.

Event Component

Specifies the component name involved in the event.

Event Message

Specifies the message that the event generated.

Event Source

Specifies the source of the event.

Event Target

Specifies the target of the event.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Create Report

The Create Report action type helps you automatically generate reports. You can schedule this action so that it creates reports regularly. This action type can also be created from the Reporting tab.

The Details section of the action definition contains the following fields:

Report Type

Specifies the type of the report that is created. For explanation of the available report types and the related creation options, see Reporting.

The generated report can be viewed on the Reporting tab in the Schedule Reports folder.

Create Service

The Create Service action type lets you organize the servers you monitor into a logical service that reflects the resources required by your business needs.

The Details section of the action definition contains the following fields:

Service Name

Specifies the name of the service.

Server List (comma delimited)

Specifies the list of available servers.

Lower Threshold

Specifies the lower threshold of the service as a whole.

Upper Threshold

Specifies the upper threshold of the service as a whole.

Lag

Defines how often the rule must evaluate as true before the action triggers. Some actions should trigger after a single occurrence, while others should trigger only after a number of occurrences signal a persistent problem.

Priority

Specifies the order in which to execute actions in a single polling cycle.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: IBM LPAR

The LPAR Delete Machine action type lets you delete a specified LPAR.

The Details section of the action definition contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Defines the partition name to be deleted.

Note: The partition being deleted must be powered off for this action. This action erases the logical partition and the logical partition configuration data stored in the partition profiles.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: Microsoft Hyper-V

The Delete Hyper-V VM action type deletes a virtual machine from your Hyper-V Server environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the host on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine that you want to delete. Select one from the drop-down list.

Attached Resources

Specifies the resources attached to the virtual machine that you want to delete. Select the resources that you want to delete:

- Hard Drive
- Floppy Drive
- DVD/ISO Image

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: Solaris Zones

The Delete Solaris Zone Machine action type deletes a zone from a Solaris Zones host.

The Details section of the action definition contains the following fields:

Zone Host

Defines the Solaris Zones host that contains the zone to delete.

Zone

Defines the zone to delete. You can use automatically generated text or text extracted from event messages.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: VMware vCenter

The Delete vCenter VM action type deletes a virtual machine from your VMware vCenter Server environment.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center on which the virtual machine resides. Select one from the drop-down list. Your selection populates the Target VM Machine drop-down list with the names of VMs associated with the data center.

Target VM Machine

Specifies the name of the virtual machine that you want to delete. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Discover Host by Name

The Discover Host by Name action type lets you discover a host using a specified host name.

The Details section of the action definition contains the following fields:

Host Name

Specifies the host name.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Discover Network

The Discover Network action type lets you discover networks that are available in your domain.

The Details section of the action definition contains the following fields:

Network ID

Specifies the network ID to be discovered.

Network Name

Specifies the network name to be discovered.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Distributed Switch: VMware vCenter

Use this action type to manage distributed virtual switches.

The Details section of the action definition contains the following fields:

Operation

Select one of the following operations:

- Add Port Group
- Remove Port Group
- Update Port Group

Virtual Center

Specifies the vCenter Server. Select one from the drop-down list.

Virtual Switch

Specifies the virtual switch you want to manage. Select one from the drop-down list.

Port Group

Specifies the port group name. Select one from the drop-down list.

Bind Type (Optional)

Select one of the following bind types:

earlyBinding

Assigns the ports when the VM binds to the portgroup. This type of binding ensures connectivity at all times, but permanently reserves the port. This binding type is the default.

lateBinding

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. LateBinding is configurable through vCenter.

ephemeral

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. Ephemeral binding is configurable through the ESX Host and vCenter.

VLAN ID (Optional)

Specifies an Integer value used for the virtual port group operations.

Number of Ports (Optional)

Specifies the number of ports of the port group.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Fault Tolerance: VMware vCenter

Use this action type to manage fault tolerance.

The Details section of the action definition contains the following fields:

Operation

Select one of the following operations for the specified VM:

- Turn On
- Turn Off
- Enable
- Disable
- Migrate Secondary VM

Virtual Center

Specifies the vCenter Server host name. Select one from the drop-down list.

Datacenter

Specifies the datacenter to which the VM belongs. Select one from the drop-down list.

Virtual Machine

Specifies the fault-tolerant VM. Select one from the drop-down list.

Secondary Host

Specifies the ESX server where the secondary VM resides. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage VM Snapshots: VMware vCenter

The Manage VM Snapshots action type lets you create, revert, or delete a virtual machine snapshot on a specified target system.

Note: If a Manage VM Snapshots action fails because the ESXi host was removed from vCenter and then added back, select the corresponding snapshot again and save the action.

The Details section of the action definition contains the following fields:

Operation

Specifies one of the following actions:

- Create Snapshot
- Revert Snapshot
- Delete Snapshot

If you select Create Snapshot, complete the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select a server from the drop-down list.

Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select a data center from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine on which to create the snapshot. Select a virtual machine from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Name

Defines a name for the virtual machine snapshot to create. You can use automatically generated text or text extracted from event messages.

Description

(Optional) Describes the virtual machine snapshot.

Capture Memory check box

(Optional) Specifies whether to create the snapshot with system running memory as part of the snapshot.

If you select Revert Snapshot, complete the following fields:

VC Server

Specifies the name of the server on which vCenter resides. Select a server from the drop-down list.

Data Center

Specifies the name of the data center in vCenter where the virtual machine resides. Select a data center from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine on which to revert the snapshot. Select a virtual machine from the drop-down list.

Name

Defines the name of the virtual machine snapshot to revert.

Enter the name or click the binocular icon and select the snapshot you want to revert to from the dialog.

ID

Defines the ID of the virtual machine snapshot to revert.

Note: You can use Name or ID to revert a snapshot, you do not need both. ID is required if you have multiple snapshots with the same name for a virtual machine.

If you select Delete Snapshot, complete the following:

VC Server

Specifies the name of the server on which vCenter resides. Select a server from the drop-down list.

Data Center

Specifies the name of the data center in vCenter where the virtual machine resides. Select a data center from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine on which to delete the snapshot. Select a virtual machine from the drop-down list.

Name

Defines the virtual machine snapshot name to delete.

Type the name or click the binocular icon and select the snapshot you want to delete from the dialog that opens.

ID

Defines the ID of the virtual machine snapshot to delete.

Note: You can use Name or ID to delete a snapshot, you do not need both. ID is required if you have multiple snapshots with the same name for a virtual machine.

Delete Children check box

(Optional) Specifies whether to delete all the children of the snapshot.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Virtual Switch: VMware vCenter

Use this action type to manage virtual switches.

The Details section of the action definition contains the following fields:

Operation

Select one of the following operations:

- Add Port Group
- Remove Port Group
- Update Port Group

Virtual Center

Specifies the vCenter Server. Select one from the drop-down list.

Datacenter

Specifies the Datacenter. Select one from the drop-down list.

ESX Server

Specifies the ESX Server to which the virtual switch belongs. Select one from the drop-down list.

Virtual Switch

Specifies the virtual switch you want to manage. Select one from the drop-down list.

Port Group

Specifies the port group name. Select one from the drop-down list.

VLAN ID (Optional)

Specifies an Integer value used for the virtual port group operations.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Windows Service

The Manage Windows Service action type controls Windows services using the AutoShell command line and scripting environment.

The Details section of the action definition contains the following fields:

Operation Options

Specifies the operation performed with the service.

Note: The Windows service status that the query service operation returns can be obtained only from the %STDOUT% parameter; it is not available from the Events table. This parameter is only valid for actions running in an action sequence.

Note: The following behavior differs from service management performed directly in Windows:

- Even if a service is in the "Stopped" status, it is possible to perform the Restart Service operation. The service status will change to "Started".
- If the service is in the "Started" status, and the Disable Service operation is performed, the service will be disabled and its status will change to "Stopped".

Host Name

Defines the name of the computer on which the service is running.

User Name

Defines the user name.

Password

Defines the password. Reenter the password for confirmation.

Service Name

Defines the name of the service for which the operation is performed. Type the name or use a text extracted from event messages.

Note: Check the Service name in the Properties dialog of the service in Windows. Do not confuse it with its Display name visible in the Computer Management window.

If you select Change Service Startup Type, complete the following field:

Startup Type

Specifies the startup type that is set for the service. The options include Automatic, Manual, Disabled. The Boot option means that a device driver is loaded by the boot loader. The System option means that a device driver is started during kernel initialization.

If you select Change Service Dependencies, complete the following field:

Dependencies

Defines the dependencies (other services, system drivers, or load order groups) that must be running before the service can be started. If you define multiple dependencies, separate them by a forward slash.

If you select Change Service Account, complete the following field:

Local System Account / This Account

Specifies the account under which the service logs in. You can use the LocalSystem account or define an account here.

Migrate Machine: VMware vCenter

The vCenter VMotion Migration action type uses VMware VMotion to migrate a virtual machine. The VMware ESX servers must be configured correctly for this operation and the VMotion license must be present on the target computer.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Source Data Center

Specifies the name of the data center where the source virtual machine resides. Select one from the drop-down list.

Source Virtual Machine

Specifies the name of the server to use as the source VM. Select one from the drop-down list.

Destination ESX Server

Specifies the name of the ESX server that is to be the target of the migration. Select one from the drop-down list.

Note: VM migration between ESX hosts is only supported when the VMs datastore/disk is shared between the two ESX hosts.

Destination Resource Pool

Specifies the name of the resource pool to use. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Modify CPU: VMware vCenter

The Modify CPU action type lets you modify the number of CPUs allocated to a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to modify memory. Select one from the drop-down list.

CPU

Specifies the number of CPUs to allocate to the VM.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Modify Memory: VMware vCenter

The Modify Memory action type lets you modify memory allocation for a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to modify memory. Select one from the drop-down list.

Memory

Specifies the amount of memory to allocate to the VM.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Open HelpDesk Ticket

The Open HelpDesk Ticket action type lets you define the properties with which to open help desk tickets.

The Details section of the action definition contains the following fields:

Summary

Summarizes the ticket details.

Description

Describes the ticket.

Entity

(Optional) Defines the name of the server or service that is used to match the ticket with a known configuration item in the help desk system. If the configuration item host name is the same as the entity name, the ticket is associated with that configuration item.

Type

Specifies the type of the ticket.

Template

Specifies the template for the ticket.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Provision Machine: IBM LPAR

This action type provisions an LPAR.

The Build Partition section contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Defines the name of the partition for image creation.

Profile Name

Defines the name of an existing profile for the selected LPAR.

The Memory Settings section contains the following fields:

Installed Memory

Identifies installed memory.

Available Memory

Identifies installed memory.

Minimum

Specifies the minimum memory.

Desired

Specifies the desired memory.

Maximum

Specifies the maximum memory.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

The Processor section contains the following fields:

Processing Mode

Specifies the processing mode.

Select between the following options:

- Partial Processor Units (Shared)
- Entire Processors (Dedicated)

Units Available

Identifies available processor units.

Minimum

Specifies the minimum processor units.

Desired

Specifies the desired processor units.

Maximum

Specifies the maximum processor units.

I/O Components

Specifies the I/O components that you want to associate with the LPAR.

I/O Pools

Lets you add, delete and modify the I/O pools.

Maximum Virtual Adapters

Defines the maximum number of virtual adapters.

Virtual Adapters

Identifies the number of virtual adapters

Virtual Serial Adapters

Lets you add, delete and modify the virtual serial adapters.

Virtual Ethernet Adapters

Lets you add, delete and modify the virtual ethernet adapters.

Virtual SCSI Adapters

Lets you add, delete and modify the virtual SCSI adapters.

The provisioning process starts on the client computer and a confirmation message notifies you when the job has completed successfully.

Provision Machine: Microsoft Hyper-V

The Provision Hyper-V VM action type creates and installs a VM. Specify the following parameters.

The Details section of the action definition contains the following fields on the first page:

SCVMM Server

Specifies the Microsoft System Center Virtual Machine Manager (SCVMM) library server. Select one from the drop-down list.

Hyper-V Server

Specifies the Hyper-V Server. Select one from the drop-down list.

Template

Specifies the template. Select one from the drop-down list.

Destination Path

Specifies the destination path of the VM that you want to create (template is stored). Select one from the drop-down list.

VM Name

Specifies the name of the VM.

Start VM

Starts the VM automatically after it is created. By default, the new VM is in powered-off state.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

The Details section contains the following fields on the second page after you click Next:

Hardware Profile

Specifies the name of the hardware profile defined by the Microsoft System Center Virtual Machine Manager (SCVMM) library server.

Virtual Processors

Specifies the number of virtual processors that you want to assign to the VM.

Default: 1

Memory

Specifies the RAM memory in megabytes (MB) for the VM that you want to create.

Default: 1024

The Details section contains the following fields on the third page after you click Next:

Guest OS Profile

(Optional) Specifies the name of the guest operating system profile defined by the Microsoft System Center Virtual Machine Manager (SCVMM) library server. This parameter overwrites the operating system configuration settings stored in the SCVMM library server. This parameter is valid when you use SCVMM integration to provision VMs.

Product Key

(Optional) Specifies the Windows product activation key for the VM. Support for this parameter requires a Windows image created using Sysprep tool. This option is invalid for asynchronous execution of command.

Full Name

Specifies the user name of the Windows image (created using sysprep tool) which is installed on the new VM.

Organization

(Optional) Specifies the organization name of the Windows image (created using sysprep tool) which is installed on the new VM. Support for this parameter requires a Windows image created using Sysprep tool. This option is invalid for asynchronous execution of command.

Admin Password

(Optional) This option is used to set the default administrator account password for the VM. Support for this parameter requires a Windows image created using Sysprep tool. This parameter is ignored in asynchronous execution.

Note: To set this option successfully, set the Windows Server administrator password which is created using Sysprep tool as empty.

Join Workgroup

Specifies the workgroup that you want to create for the VM. Domain and workgroup specifications are mutually exclusive.

Join Domain

Specifies the domain name for the VM. Domain and workgroup specifications are mutually exclusive.

Domain User

Specifies the domain user name that you want to create as a part of the default Administrators group.

Domain User Password

Specifies the password of the domain user account that you want to create as a part of the default Administrators group.

The Details section contains the following fields on the fourth page after you click Next:

Use DHCP

Specifies an option to enable DHCP for a network interface of the VM. If the template image has more than one network adapter, DHCP is turned on for the first interface. If enabled, the other network parameters are not accessible.

IP Address

Specifies the static IPv4 address that you want to assign to the VM.

Network Mask

Specifies the subnet mask that you want to assign for the VM.

Default Gateway

Specifies the default gateway for VM.

DNS Server

Specifies the DNS server that you want to set for the VM.

IP Metric

(Optional) Specifies the interface metric that you want to set for the VM. This option is used with -ip4ddr option. If an interface name is specified in the -ip4addr option, same interface name must be used in this option. Support for this parameter requires a Windows image created using Sysprep tool. This option is invalid for asynchronous execution of command.

Default: 1

Provision Machine: Solaris Zones

The Provision Solaris Zone Machine action type creates and installs a zone. Specify a Solaris Zones host, a zone name, the zone type, and other zone properties. The zone installs automatically after creation.

The Details section of the action definition contains the following fields on the first page:

Host

Defines the Solaris Zones host on which to create the zone.

Name

Defines the zone name. You can use automatically generated text or text extracted from event messages.

Description

(Optional) Defines a description of the zone.

Type

Defines whether the Zone is Native, Whole Root, or Branded. A Branded Zone is based on an existing Zone template.

Template

(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.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

The Details section contains the following fields on the second page after you click Next:

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

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

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.

Provision Machine: VMware vCenter

The Provision vCenter machine action type provisions a Virtual Machine (VM). A template and a target vCenter specification that works with the template is required. If a service rule exists for provisioning a VM, that new VM is placed in the service for which the rule was created.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center on which to provision the machine. Select one from the drop-down list.

VC Compute Resource

Specifies the name of the server on which the compute resource resides. Select one from the drop-down list.

VC ESX Server

Specifies the name of the VMware ESX server that is to be the target of the provisioned VM. Select one from the drop-down list.

VC Datastore

Specifies the name of the data store to use. Select one from the drop-down list.

VC Target Location

Specifies the VC target location. Select one from the drop-down list.

Hostname/VM Name

Specifies the name or VC name to use from the specification. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

User Name

Specifies the user name credentials to access the specification.

Password

Specifies the password to access the specification.

VC Virtual Machine

Specifies which available VC virtual machine to use. If selected, click one from the drop-down list.

VC Template

Specifies which available VC template to use. If clicked, select one of the software package groups that has previously been created from the drop-down list.

NICs (VC Template)

Specifies the number of network interface cards used by the VC template.

VC Specification

Specifies the name of the VC specification to use. Select one from the drop-down list.

NICs (VC Specification)

Specifies the number of network interface cards used by the VC specification.

OS System Type

Displays the type of operating system for the provisioned VM.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Memory

Specifies the amount of memory to allocate to the VM, in megabytes.

Virtual Processors

Specifies the number of virtual processors to allocate to the VM.

Datastore

(Optional) Specifies the storage datastore under which to create additional hard disks.

Drive Size

(Optional) Specifies the size of an additional hard drive.

SCSI Controller

(Optional) Specifies the SCSI controller to use to create the additional hard drive.

Network Management

Lets you change network connection settings.

Global NIC Settings

Lets you add DNS search suffixes.

Remove Disk: VMware vCenter

The Remove Disk action type lets you remove a disk from a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to add a disk. Select one from the drop-down list.

Hard Drive

Specifies the disk to be removed. Select one from the drop-down list.

Delete Disk File(s) check box

Specifies whether to delete the disk data or not.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Remove Network Interface: VMware vCenter

The Remove Network Interface action type lets you remove a virtual NIC from a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to remove the virtual NIC. Select one from the drop-down list.

Network Interface

Specifies the virtual NIC to be removed. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Remove Server From Service

The Remove Server From Service action type lets you remove a server from an existing service.

The Details section of the action definition contains the following fields:

Service

Specifies the name of the service.

Server List (comma delimited)

Specifies the list of servers to remove from the service.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Run Action

The Run Action action type lets you run actions.

The Details section of the action definition contains the following fields:

Action Name

Specifies the action.

Event Source

Specifies the source of the action.

Event Message

Specifies the event message.

Rule Name

Specifies the rule for the action.

Server Name

Specifies the server for the action.

Service Name

Specifies the service for the action.

Propagate

Specifies that the action runs against all servers in the service specified in the -service_name option.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Run Action Sequence

The Run Action Sequence action type lets you select multiple actions for a rule and run them in a defined sequence.

The Details section of the action definition contains the following fields:

Restart if interrupted

Restarts the action sequence if it is interrupted. The action does not resume from the point that it was interrupted, but it starts the sequence over from the beginning.

Action Sequence

Provides single or multiple row selections of actions and conditions. When you select Action Sequence, a drop-down list is enabled. If it is not selected, text is displayed in the table cell.

Sequence

Specifies the sequence of actions.

Note: If an action sequence exits without a condition being met, the default return code is -1.

Action

Specifies the action name. You can select an action from the available actions in the drop-down list.

Condition Name

Specifies the condition that determines the next action to run. You can create your own custom conditions or use one of the following predefined conditions:

- On Failure
- On Success

Note: Conditions are evaluated in the order that they are created.

Next Step

Specifies the next action to run based on the results of the condition.

Continue

Continues to the next action when the condition evaluates to true.

Exit (RC=0)

Exits the sequence and returns a code of 0 to the log when the condition evaluates to true.

Exit w/RC (RC=action RC)

Exits the action sequence and the action's return code when the condition evaluates to true.

Abort (RC=-1)

Stops the action sequence and returns a code of -1 to the log when the condition evaluates to true.

Go to #

Continues to the action sequence number specified when the condition evaluates to true.

Add Action

Adds a new action to the table and automatically generates a new sequence number.

Add Condition

Adds a new condition to the action.

Delete

Deletes the selected row and updates the sequence numbers. A row can contain an action or a condition. This function permits you to delete a condition of an action without deleting the entire action.

Save

Saves the action sequence.

Note: If the Next Step in the last action in the sequence is set to *Continue*, the setting is automatically changed to “Exit w/RC (RC=action RC)”, and an informational message notifies you that the Next Step for the last action in the sequence has been changed.

Run Command Script

The Run Command Script action type lets you use a script to run an external command from the server that processed the command. For example, if the command is run from the Initiation page, the target command must be on the same server as the Windows scheduler that runs the command. When the command runs as a result of a rule evaluation, the action runs on the computer hosting the Windows scheduler server, as a result of running a job.

The Details section of the action definition contains the following fields:

Command Line

Specifies the command or substitution string to run. Alternatively, you can use automatically generated text or text extracted from event messages.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Set Health State

The Set Rule Health State action type lets you set the health state of a system to any one of the following: Warning, Minor, Major, or Critical.

The Details section of the action definition contains the following fields:

Health State

Specifies one of the following actions:

- Warning
- Minor
- Major
- Critical

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

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 dpmpolicy 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.

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.

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 697)
Use Case: Adding a New Rule to a Service (see page 698)
Use Case: Defining an Action (see page 698)

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 data is collected in your data center, including:

- Time intervals for metrics collection
- Systems from which to collect metrics (filtering)
- Metrics to collect for each server
- Data aging and data expiration (how long to retain data).

More information:

Configure Data Collection for a Server (see page 703)
Configure Data Collection for a Data Center (see page 702)
Configure Performance Thresholds (see page 707)
Key Points About Metrics Collection (see page 699)
Configure the Metric Filter (see page 707)
Configure Data Collection for a Virtual Resource (see page 705)

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 Systems Performance LiteAgent or with the SystemEDGE agent on the remote computer to collect the specified system metrics.

  Install CA Systems Performance LiteAgent or the SystemEDGE agent on any server from which you want to collect the base system metrics. If SystemEDGE agents are present, then the CA Systems Performance LiteAgent 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 that are 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.

  Note: Select Include for Overall Calculation in the Policy, Metrics, Thresholds section of the user interface to include new metric in the overall utilization calculation. If you include the metric, CA Virtual Assurance provides up-to-date results when evaluating the state of the servers.

- How metric evaluations affect the overall utilization? The metric details provided in the tables help you understand how CA Virtual Assurance evaluates the different metrics. Each metric has a method property that is 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 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. The 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, sometimes a more immediate means of determining if collected data is available is needed, or performance data is critical. Configure CA Virtual Assurance 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
```

CA Virtual Assurance modifies 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. These changes reflect 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 *Performance Metrics Reference*.

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

   Note: CA Technologies recommends that for every 1000 machines in your monitored environment, increase the data recording interval by 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: 10 days

   Daily roll-up data retention (days)

   Specifies how long to store the average of the daily data in the Performance DB.

   Maximum: 365

   Default: 365
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

Note: CA Technologies recommends that for every 1000 machines in your monitored environment, increase the data recording interval by 300 seconds.

Daily roll-up data retention (days)

Specifies how long to store the average of the daily data in the Performance DB.

Maximum: 365

Default: 365

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: 10 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 Virtual Assurance.

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 PowerVM 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

   Note: CA Technologies recommends that for every 1000 machines in your monitored environment, increase the data recording interval by 300 seconds.

   Daily roll-up data retention (days)

   Specifies how long to store the average of the daily data in the Performance DB.

   Maximum: 365

   Default: 365

   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: 10 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.

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^{Th} 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 CA Virtual Assurance collects the metrics specified by the filter by default. Unless a metrics filter is set as default, CA Virtual Assurance does 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.

Appendix A: FIPS 140-2 Encryption

This section contains the following topics:

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: Tools

This section contains the following topics:

Configure AIMs with NodeCfgUtil

The Node Config Utility lets you configure SystemEDGE AIMs without using the CA Virtual Assurance user interface.

This section describes the Dialog Mode and the Command Mode of the utility.

More information:

NodeCfgUtil Overview

To configure the AIMs and discover virtual environments, do *one* of the following actions:

- 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 that are specified for accessing virtual environments or 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.

- Citrix XenServer AIM can only connect to pool masters or standalone Citrix XenServers. Otherwise, the AIM does not work.
- To manage VMware vSphere, enter the credentials for the corresponding vCenter Servers.
- To optimize the virtualization of your VMs, install the corresponding system tools on your VMs. Many features are available only if these tools are installed. Use the following system tools, depending on your environment:
  - (Valid for VMware) VMware Tools
  - (Valid for XenServer) XenTools
  - (Valid for RHEV) RHEV Guest Tools

  Note: For further information about the corresponding system tools, see the Third-party documentation.

To discover supported environments from an AIM Server:

1. 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.
2. 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 for each AIM in a file (for example, zone.cfg, vc.cfg, ...).
3. Open the user interface on the Manager Server and click Resources, Data Center in the navigation pane.
4. Right-click and select Management, Discover.

   The discovery options appear.
5. Select one of the following actions:
   - Discover a system
   - Discover a network

   The corresponding dialog opens.
6. Enter a system name of 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 473)

Configure AIMs with NodeCfgUtil in Dialog Mode

NodeCfgUtil.exe lets you modify the AIM configurations for IBM PowerVM, IBM PowerHA, Solaris Zones, VMware vCenter, VMware vCloud, Microsoft Clusters, Cisco UCS, Citrix XenServer, Citrix XenDesktop, RHEV, Active Directory and Exchange Server (ADES), or Huawei GalaX. The utility writes a configuration file for the corresponding AIM to the s*ysedge_InstallLpath**plugins**AIPCommon* directory. You can also use the NodeCfgUtil utility to edit or remove existing entries.

Use the utility in dialog mode to configure which nodes the appropriate AIMs manage.

Note: Run NodeCfgUtil.exe as Windows Administrator.

Follow these steps:

1. Log in as Administrator and open Windows Explorer on the computer on which the AIM is installed.
2. Change to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory, and start NodeCfgUtil.exe.

   NodeCfgUtil discovers and lists the installed AIMs in subsequent dialogs.
3. Enter *1* to add a new managed node.
4. Follow the on-screen instructions to complete the configuration. Each node requires a valid user name and password for authentication.

   After the 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), vCloud Director (vcloud.cfg), Microsoft Clusters (mscs.cfg), Citrix XenServer (cxen.cfg), UCS (ucs.cfg), PowerVM (lpar.cfg), PowerHA (hacmp.cfg), RHEV (kvm.cfg), Huawei GalaX (galaxa.cfg), Citrix XenDesktop (xendesktop.cfg), or ADES (esad.cfg) to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory.

Note: You can also use the NodeCfgUtil utility to edit or remove existing entries. The corresponding dialogs are self-explaining.

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 can add more vCenter Servers that you want to manage with this AIM.

```
*****  Main MENU  *****
1. Install Managed Node
2. Modify Managed Node

3. Remove Managed Node

0. Exit

************************

Enter choice:

****  Choose Managed Node  ****
1. IBM PowerVM
2. Oracle Solaris Zones
3. VMware vCenter
4. Cisco UCS
5. Microsoft Cluster Service
6. Microsoft Active Directory and Exchange Server
7. IBM PowerHA
8. VMware vCloud Director
9. Citrix XenDesktop
10.Go Back to Previous Menu
*****************************
Enter choice: 4
Enter following information for the VMware vCenter Node...

(At any point to go back to the 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 can add more entities that you want to manage with this AIM.

```
****  Choose Managed Node  ****
1. Microsoft Cluster Service
2. Microsoft Active Directory and Exchange Server
0. Go Back to Previous Menu ******************************
Enter choice: 2
Enter following information for the Microsoft Active Directory and Exchange Server
Node...

(At any point to go back to the 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...
```

The following example shows the Managed System dialog for the HMC1 server that has been successfully added to the configuration of the LPAR AIM. After the AIM discovers all Virtual I/O Servers that are related to the HMC server, they are visible in NodeCfgUtil and each one can be modified to specify its credentials. The AIM uses the credentials of the first fully configured VIO Server as default credentials for all VIO Servers not yet configured. Thus, it is sufficient to specify the credentials for only one VIO Server if all of them share credentials. Otherwise, it is necessary to configure each VIO Server with different credentials. The AIM is now ready to manage the HMC Server.

```
****  Choose Managed Node  ****
1. IBM PowerVM
0. Go Back to Previous Menu
******************************
Enter choice: 1
List of existing entries...
1. hmc: HMC1.company.com
2. vio: ibm101.company.com
```

```
Select the entry to be modified (0 to go back to the previous menu): 2
Enter following information for the IBM LPAR Node...
(At any point to go back to the previous menu, Enter 'CTRL Q').
1. Server Name: ibm101
2. User Name: admin
3. Password: *********

CAAC1016 Authenticating, please wait...
CAAC1019 Authentication SUCCESSFUL.
CAAC1024 Modified Node Successfully.

Press any key to continue...
```

The following example shows the Install Managed Node dialog for *myserver* that has been successfully added to the configuration of the GalaX AIM. For details, see the NodeCfgUtil command mode. The GalaX AIM is a multi-instance AIM. So you can repeat this procedure and can add more entities that you want to manage with this AIM.

Note: To configure Huawei GalaX component you need to specify the certificate filename.

```
****  Choose Managed Node  ****
1. Huawei GalaX
0. Go Back to Previous Menu
*****************************
Enter choice: 1
Enter following information for the Huawei GalaX Node...

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

1. Server Name: myserver
2. Certificate file name: certificatename123.p12
3. Password: ************
4. Port [default =8773]:
5. Protocol [default =http]:

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 PowerVM, IBM PowerHA, Solaris Zones, VMware vCenter, VMware vCloud, Microsoft Clusters, Cisco UCS, Citrix XenServer, Citrix XenDesktop, RHEV, Active Directory and Exchange Server (ADES), or Huawei GalaX. The utility writes a configuration file for the corresponding AIM to the s*ysedge_InstallLpath**plugins**AIPCommon* directory. You can also use the NodeCfgUtil utility to edit or remove existing entries.

When you use the utility in command mode, you can only add managed nodes to an AIM configuration.

Note: Run NodeCfgUtil.exe 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 {vc|ucs} -u user -p password -h hostname -t port -c protocol
(4) nodecfgutil ades -u user -p passwd -d domainname -e entity -o option
(5) nodecfgutil {xen|vcloud|xenserver} -u user -p passwd -h hostname
(6) nodecfgutil {powerha|kvm} -u user -p password -h {cluster_name|hostname} [-t port]
(7) nodecfgutil galax -u usercertificate -p password -h hostname [-t port] [-c
protocol]
```

-help

Displays usage information about the console.

lpar|ucs|vc|zone|mscs|ades|xen|vcloud|powerha|kvm|galax|xendesktop

Specifies the virtual or physical environment.

-u *user|usercertificate*

Specifies the name of an administrative user or the user certificate, accordingly.

-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 cluster.

-t *port*

(Optional) 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.

Follow these steps:

1. Open a command prompt on the system on which the AIM is installed.

   The command prompt appears.

2. Enter *one* of the following commands:

   ```
   (1) nodecfgutil -help
   (2) nodecfgutil {lpar|zone|mscs} -u user -p password -h
   {pvmname|hostname|cluster_name}
   (3) nodecfgutil {vc|ucs} -u user -p password -h hostname -t port -c protocol
   (4) nodecfgutil ades -u user -p passwd -d domainname -e entity -o option
   (5) nodecfgutil galax -u usercertificate -p password -h hostname [-t port] [-c
   protocol]
   ```

(1) Displays the usage information about the console.

(2) Authenticates and stores the passed credentials for Solaris Zones, IBM PowerVM, or MSCS.

(3) Authenticates and stores the passed credentials for vCenter or Cisco UCS.

(4) Authenticates and stores the passed credentials for Active Directory and Exchange Server (ADES).

(5) Authenticates and stores the passed credentials and user certificate for HUAWEI Galax.

Support Agent

The Support Agent collects diagnostics information. To access the Support Agent, use the following address:

http://<Manager Server>:8556

The user interface is self-explanatory and provides the following information:

- The performance metrics of important parts of the system
- Detailed Web Service usage statistics
- Log files monitoring
- Long run SQL queries

Appendix C: Troubleshooting

This section contains the following topics:

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.

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 nodecfgtutil.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 about troubleshooting a full transaction log.

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

Domain Server is not available

Symptom:

Domain server is unavailable, stopped, nonfunctional or not servicing request and Service Controller (SC) shows that the component is up and running.

Solution:

This behavior is due to a database connection failure or an AIM password expiry and can potentially impact the behavior of Policy Configuration and Remote Deployment components. Monitoring Support Service Web Service (ISM) monitors the functionality of the Domain Server periodically. When ISM identifies an unexpected behavior, the user is notified about the changes in the status with the message: *The CA SM Domain Service is down or not responding.*

You can monitor the status with the following command:

```
Caaipscutil /status /id=ISM /user=<user> /password=<password>
```

Administration panel indicates the Domain Server status to make sure the infrastructure state and functionality.

eHealth does not discover LPAR Physical Disks

Symptom:

The eHealth does not discover any LPAR Physical Disks with Release 12.9 LPAR AIM.

Solution:

If you use Release 12.9 LPAR AIM, upgrade your eHealth to the following version:

- 6.2.2 D11 if you have eHealth 6.2.2.10 or below
- 6.3.0.06 or later if you have eHealth 6.3.0.05 or below

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.

Solution:

For seamless local and remote monitoring, identical monitored object names can be chosen. However, different APIs 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. Monitored metrics are highly volatile, and samples are likely to differ.

Naming Limitations of IBM Logical Partitions

Symptom:

When I specify the name of a logical partition, CA Virtual Assurance does not support it.

Solution:

The name of an LPAR provided to the IBM Create LPAR request is case-sensitive. This product however does not support management of two logical partitions in the same PowerVM server where the name differs only by case. For example, the following LPAR names are not supported in PowerServer1:

LPAR1 and lpar1

LPAR names must not contain the `/` character, because it is used as a separator of entities in object instances of monitors. A `/` character produces ambiguous monitor object instances. For example, an LPAR name ‘lpar/blue’ is not supported.

Navigation Problem in SystemEDGE Installer on AIX Systems

Symptom:

When I install SystemEDGE on AIX 6.1 and 7.1, navigation does not work in the lsm (UNIX Installer) text user interface. This problem also occurs with Advanced Encryption and the SRM AIM.

Solution:

Unlike on other UNIX operating systems and older AIX versions, navigation in lsm text user interface does not work on AIX 6.1 and 7.1 using the keyboard arrow keys when TERM is set to the (common) value of xterm. The problem does not occur when using the Java-based graphical lsm UI.

Workaround is to either set TERM to a different value (for example, vt100) before starting the installation, use + and – keys to navigate, or (PuTTY specific) set “Disable application cursor keys mode”.

NodeCfgUtil Fails to Validate the Connection to XenDesktop Controller

Symptom:

NodeCfgUtil fails to validate the connection to XenDesktop controller.

Solution:

Verify that the following components are installed on the machine where XenDesktop AIM is installed:

- Microsoft .NET Framework 4.0
- Windows Management Framework Core (Windows PowerShell 2.0, Windows Remote Management (WinRM) 2.0)

Performance Chart Shows Zero Memory Usage on LPAR Level

Symptom:

When I monitor memory usage on LPAR level, the performance chart shows zero.

Solution:

The platform only provides memory utilization metrics, if memory is used in shared mode, that is, if memory is virtualized.

For dedicated memory, no utilization metrics can be collected and hence a utilization value of 0 is reported in the utilization chart.

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 PowerVM, 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.

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.

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.

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.

Status Icon of Component Shows Not Configured

Symptom:

After CA Virtual Assurance installs a component, the status icon of this component shows (Not configured). This status appears if CA Virtual Assurance registered a component that is connected to an unconfigured server.

Solution:

To change the status of the component to ready, add the missing Server connection settings and validation.

Upgrading SystemEDGE

Symptom:

When I upgrade SystemEDGE to Release 5.9, the AIMs of the previous CA Virtual Assurance release do not run.

Solution:

Upgrade Advanced Encryption and all AIMs to CA Virtual Assurance Release 12.9. SystemEDGE Release 5.9 does not load AIMs of previous CA Virtual Assurance releases.

Unable to Connect to Microsoft SQL Server

Symptom:

Attempts to authenticate credentials to a Microsoft SQL Server 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.

User Interface Does Not Reflect Product Upgrade

Symptom:

After I upgrade to the new version of CA Virtual Assurance, the user interface does not reflect the new version.

Solution:

If you used the same browser instance before and after upgrade, the user interface may not reflect the new version. Close the browser session, open a new one, clean the browser cache, and log in to the user interface.

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.

User Interface is not Working

Symptom:

When I use a remote SQL Server with Windows Authentication, the user interface does not work properly.

Solution:

During the installation, you are prompted to add an appropriate domain user and grant "Logon as a Service" permission. Verify that the CAAIPTOMCAT, CAAIAPACHE, and CA SM Domain Server services are configured for this domain user account. If you fail to reconfigure the services, the CA Virtual Assurance user interface is not functional (empty dashboard or features not working).

These conditions are not required for SQL Server authentication.

Follow these steps:

1. Open the Services dialog from the Control Panel, Administrative Tools.

   The list of available services appears.

2. Open the Properties dialogs for the CA SM Domain Server, CAAIAPACHE, and CAAIPTOMCAT services.

3. In each dialog change to the "Log In" tab, select "This account", and enter valid credentials that can be browsed (domain user account).

4. Add this domain user account to the local administrator groups on both systems (manager server and database server).

5. Add this domain user account to the sysadmin (or at least dbcreator) server role of the SQL Server.

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 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, port, web client protocol, web client port (optional), web client user(optional), web client user password(optional). 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 discovers 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 add 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 number of managed objects that are 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 finishes, 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 add 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 servername`

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 host 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 servername`

   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 No Polling

Symptom:

After I add a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (no polling).

Solution:

No specific actions are required for the associated instance. This icon informs you that the CA Virtual Assurance manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular vCenter Server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

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.

Glossary

access control list

The access control list or ACLs specify 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.

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.

capped logical partition (LPAR)

A *capped logical partition* is a logical partition that cannot use more processor power than its assigned processing units. The capped partition is assigned a maximum capacity and guarantees a capacity that cannot be exceeded and cannot affect the overall behavior of the physical system.

catalog (VMware)

Organizations provide *catalogs* to store vApp templates and media files. The members of an organization can use the vApp templates and media files in the catalog to create their own vApps.

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 cap

A *cpu cap* limits the amount of CPU resources for a zone.

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. Depending on their virtual configuration, 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 infrastructure.

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.

dual HMC (LPAR)

A *dual HMC* is a redundant Hardware Management Console (HMC) management system that provides high availability.

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) 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.

Elastic Service Controller (ESC)

An *Elastic Service Controller (ESC)* is a Huawei controller that provides centralized management of virtual resources, computing, storage, and other services.

entitled pool capacity (LPAR)

The *entitled pool capacity* of a shared processor pool defines the guaranteed processor capacity that is available to the group of partitions in the processor pool.

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 that is used to perform management tasks on IBM PowerVM Systems. 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 Systems, 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)*.

Huawei SingleCLOUD

Huawei SingleCLOUD is a cloud service solution for cloud computing data centers.

Hyper-V

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.

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.

IBM High Availability Cluster Multiprocessing (HACMP)

IBM High Availability Cluster Multiprocessing (HACMP) is a solution for building high-availability clusters on the AIX UNIX and Linux for IBM system p platforms.

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 System. 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.

kernel-based virtual machine (KVM)

The k*ernel-based virtual machine (KVM)* is a hardware-assisted virtualization infrastructure for the Linux kernel.

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).

Multiple Shared-Processor Pools (MSPPs)

The *multiple shared-processor pools (MSPPs)* is a capability that is supported on Power6 and later servers. This capability enables the creation of multiple processor pools to make allocation of the CPU resource more flexible.

Multiple Virtual I/O Servers

Multiple Virtual I/O servers offer capability that increase application availability by enabling Virtual I/O server maintenance without a downtime for the client partitions.

NetApp filer

A *NetApp filer* is a disk storage device that 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.

network object (XenServer)

Each XenServer host has one or more *network objects*, which are virtual Ethernet switches. Network objects have a name and description, a globally unique UUID, and a collection of virtual and physical network interfaces (VIFs and PIFs) connected to them. VM and host objects that are attached to a particular network object can send network packets to each other.
Networks without an association to a PIF are considered *internal*, and provide connectivity only between VMs on a XenServer host, with no connection to the outside world. Networks with a PIF association are considered *external*, and provide a bridge between VIFs and the PIF connected to the network.

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.

Open Virtualization Format (OVF)

The *Open Virtualization Format (OVF)* is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels that are associated with it.

P12 file

A *P12 file* is an archive file that stores a private key together with its certificate. A P12 file is used in Huawei GalaX environments.

physical block device, PBD (XenServer)

A *physical block device (PBD)* object represents an attachment between a host and a storage repository object. PBDs store the device configuration fields that are used to connect to and interact with a given storage target.

physical network interface, PIF (XenServer)

A *physical network interface (PIF)* object represents a physical network interface on a XenServer host. PIF objects have a name and description, a globally unique UUID, the parameters of the NIC they represent, and specify the network and server they are connected to. PIF objects abstract both physical interfaces and VLANs.

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 pools (LPAR)

A *processor pool* is a set of physical processors that can be shared across different logical partitions.

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.

Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization (RHEV) is an enterprise virtualization product that is based on the KVM hypervisor.

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.

resource pool master (XenServer)

A resource pool consists at least of one physical node, the *resource pool master*. Other physical nodes that join existing pools are described as members. Only the master node exposes an administration interface used by XenCenter and the CLI. The master forwards commands or requests from outside the pool to individual members as necessary.

resource pool, overcommitted (XenServer)

A *resource pool* comprises multiple XenServer host installations, bound together into a single managed entity that can host VMs. When combined with shared storage, a resource pool enables VMs to start on any XenServer host with sufficient memory and then move dynamically between XenServer hosts (XenMotion).
A resource pool is overcommitted if the VMs that currently run in the resource pool cannot be restarted elsewhere following a user-defined number of failures. XenServer dynamically maintains a failover plan for what to do if a set of hosts in a resource pool fail at any given time. The host failures to tolerate value defined as part of the high availability (HA) configuration determines the number of failures that is allowed without any loss of service. If a plan is not available, the pool is considered to be overcommitted. The plan is dynamically recalculated based on VM lifecycle operations and movement.

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.

storage repository, SR (XenServer)

A *storage repository* (SR) describes a particular storage target, in which virtual disk images (VDIs) are stored. The interface to storage hardware allows VDIs to be supported on many SR types.

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.

template (XenServer)

Templates are VMs with the *is_a_template* parameter set to true. A template contains all the various configuration settings to instantiate a specific VM. XenServer ships with a base set of templates, which range from generic raw VMs that can boot an OS vendor installation CD or run an installation from a network repository to complete preconfigured OS instances.
With XenServer you can create VMs, configure them in standard forms for your particular needs, and save a copy of it as a template for future use in deployment.

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).

universally unique identifier, UUID

A *universally unique identifier, UUID* is an identifier standard that is used in distributed systems to identify information uniquely. Labeling information with UUID limits identifiers conflicts when information is stored in a single database.

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. The *Open Virtualization Format (OVF)* is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels that are 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.

vCloud Director (VMware)

VMware vCloud Director lets you build secure, multitenant clouds by pooling virtual infrastructure resources into virtual datacenters and exposing them to users. vCloud Director resources depend on underlying vSphere resources such as CPU, memory, storage, or vNetwork Distributed Switches to run virtual machines. You can use these underlying vSphere resources to create virtual machines and vApps in vCloud.

vCloud Organization (VMware)

A *vCloud Organization* is a unit of administration that represents a collection of users, groups, and computing resources. Associated virtual datacenters provide the required computing resources. After users authenticate at the organization level, they can create, use, and manage virtual machines or vApps.

virtual block device, VBD (XenServer)

A *virtual block device (VBD)* object represents an attachment between a virtual machine (VM) and a virtual disk image (VDI). When a VM is booted, its VBD objects are queried to determine which disk images should be attached.

virtual datacenter, vDC (VMware)

A *virtual datacenter (vDC)* provides virtual computing resources to a vCloud organization. You can provision, run, and store virtual systems in a virtual datacenter. A vCloud organization can have multiple virtual datacenters.

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 disk image, VDI (XenServer)

A *virtual disk image (VDI)* is an on-disk representation of a virtual disk provided to a VM. VDIs are the fundamental units of virtualized storage in XenServer.

Virtual I/O Server, VIOS (LPAR)

A *Virtual I/O Server (VIOS)* is a special logical partition that is 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 input output devices has one or more Virtual I/O Servers.

virtual local area networks, VLAN (XenServer)

Virtual local area networks (VLANs) allow a single physical network to support multiple logical networks. To use VLANs with XenServer, the host's NIC must be connected to a VLAN trunk port.

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 machine, VM (XenServer)

A *virtual machine (VM)* specifies virtualized x86 environments in which guest operating systems and applications can run. VMs are created from templates. A template contains all the various configuration settings to instantiate a specific VM.

XenServer provides a base set of templates, which range from generic raw VMs that can boot an OS vendor installation CD or run an installation from a network repository to complete preconfigured OS instances. XenServer supports Linux and Windows guest operating systems.

virtual network interface, VIF (XenServer)

A *virtual network interface (VIF)* object represents an attachment between a VM and a network object. VIF objects have a name and description, a globally unique UUID, and specify the network and VM they are connected to. When a VM is booted, its VIF objects are queried to determine which network devices it must create.

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 Private Cloud (VPC)

A Virtual Private Cloud (VPC) is a private local network for a Huawei SingleCLOUD user with several virtual machines and associated virtual disks.

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.

XenCenter (XenServer)

XenCenter is a Windows client application to manage a XenServer environment. It must be installed on a remote Windows computer that can connect to the XenServer hosts through the network, but it cannot run on the same system as the XenServer host.

XenMotion (XenServer)

XenMotion provides the ability of live migration of VMs within a resource pool.

XenServer host (XenServer)

A *XenServer host* object represents a physical host on which XenServer and its VMs run. A XenServer host can be a stand-alone host or associated with a XenServer pool.

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

(

A

B

C

D

E

F

G

H

I

N

O

P

R

S

T

U

V

X

Z