CA Server Automation

Administration Guide

Release 12.6

CA Technologies Product References

This document may reference the following CA Technologies products and components or third-party components:

- CA AppLogic®
- CA Configuration Automation, formerly CA Application Configuration Manager (CA ACM)
- CA eHealth®
- CA Embedded Entitlements Manager (CA EEM)
- CA IT Asset Manager
- CA IT Client Manager (CA ITCM)
- CA Network and Systems Management (CA NSM)
- CA Network Automation
- CA Patch Manager
- CA Process Automation, formerly CA IT Process Automation Manager (CA IT PAM)
- CA Server Automation
- CA Service Desk Manager (CA SDM)
- CA Spectrum® Infrastructure Manager (CA Spectrum)
- CA SystemEDGE
- Racemi® DynaCenter®
- Software Delivery, a component of CA IT Client Manager

Contact CA Technologies

Contact CA Support

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

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

Providing Feedback About Product Documentation

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

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

Contents

Chapter 1: Introduction

This section contains the following topics:

Related Publications

The following publications provide information about CA Server Automation and are located on your installation media, the CA Bookshelf, and the CA Support Online website:

Administration Guide (DCA_Admin_ENU.pdf)

Describes product architecture, troubleshooting, concepts, and configuration tasks for administrators.

Installation Guide (DCA_Install_ENU.pdf)

Describes installation prerequisites, best practices, and procedures for CA Server Automation.

Reference Guide (DCA_Ref_ENU.pdf)

Provides detailed information about AutoShell, CLI scripting commands, log files, and performance metrics.

Online Help (HTML)

Provides information to help you complete tasks using the CA Server Automation user interface.

Reservation Manager Help (HTML)

Provides information to help users and administrators complete tasks using the Reservation Manager user interface.

Release Notes (DCA_Release_ENU.pdf)

Provides information about new and changed features and product implementation information including operating system support, system requirements, and how to contact Technical Support.

Service Response Monitoring User Guide (VAS_SRMUser_ENU.pdf)

Provides installation and configuration details of SRM.

SystemEDGE Release Notes (SE_Release_ENU.pdf)

Provides information about new and changed features and agent implementation information including operating system support, system requirements, and how to contact Technical Support.

SystemEDGE User Guide (SE_User_ENU.pdf)

Provides end-user information about the SystemEDGE agent.

CA Bookshelf (HTML)

Provides HTML and PDF versions of all guides and search capabilities.

To view PDF files, download and install the Adobe Reader from the Adobe website if it is not already installed on your computer.

Conventions

This guide uses the following conventions to communicate ideas consistently:

Case-Sensitivity

All class, command, directive, environment parameter, function, and property names mentioned in this guide are case-sensitive. The names must be spelled exactly as shown. System command and environment variable names can be case-sensitive, depending on the requirements of your operating system.

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 a chapter name in this guide or another guide.

Chapter 2: Overview

This section contains the following topics:

Architecture

CA Server Automation is a policy-based product that automatically monitors, reconfigures, and provisions physical and virtual resources to meet the load demands of complex service-oriented data centers. CA Server Automation is built on a service-oriented architecture (SOA) and continuously analyzes your data center to help ensure that your servers are optimally provisioned to perform required tasks. Use the web-based CA Server Automation user interface to manage your data center and obtain detailed information about each managed computer in your data center.

CA Server Automation leverages the following existing CA technologies:

- CA Service Desk Manager for request escalation and resolution
- CA IT Client Manager (CA ITCM) for application and operating system deployment
- CA EEM for security
- CA Configuration Automation for discovery and change management
- CA Network Discovery Gateway for lightweight standalone discovery capabilities
- CA Patch Manager for patch maintenance and delivery to relevant data center resources
- CA Process Automation to automate data center workflow processes
- CA AppLogic for provisioning complete applications within an AppLogic environment.

CA Server Automation integrates with the following external technologies:

- Cisco Unified Computing System (UCS) for heterogeneous hardware provisioning and virtualization
- VMware vCenter Server for virtual machine (VM) operating environment management
- Microsoft Hyper-V Server for virtual machine (VM) operating environment management and optional integration with Microsoft System Center Virtual Machine Manager (SCVMM) for provisioning
- Amazon Elastic Compute Cloud (EC2) for remote virtualized Windows and Linux provisioning in the Amazon Cloud
- IBM Logical Partitioning (LPAR) for AIX virtualization operating environment management
- IBM Network Installation Manager (NIM) for AIX provisioning
- Rapid Server Imaging (RSI) powered by Racemi for cross-platform and dissimilar hardware provisioning

The following diagram illustrates the product architecture:

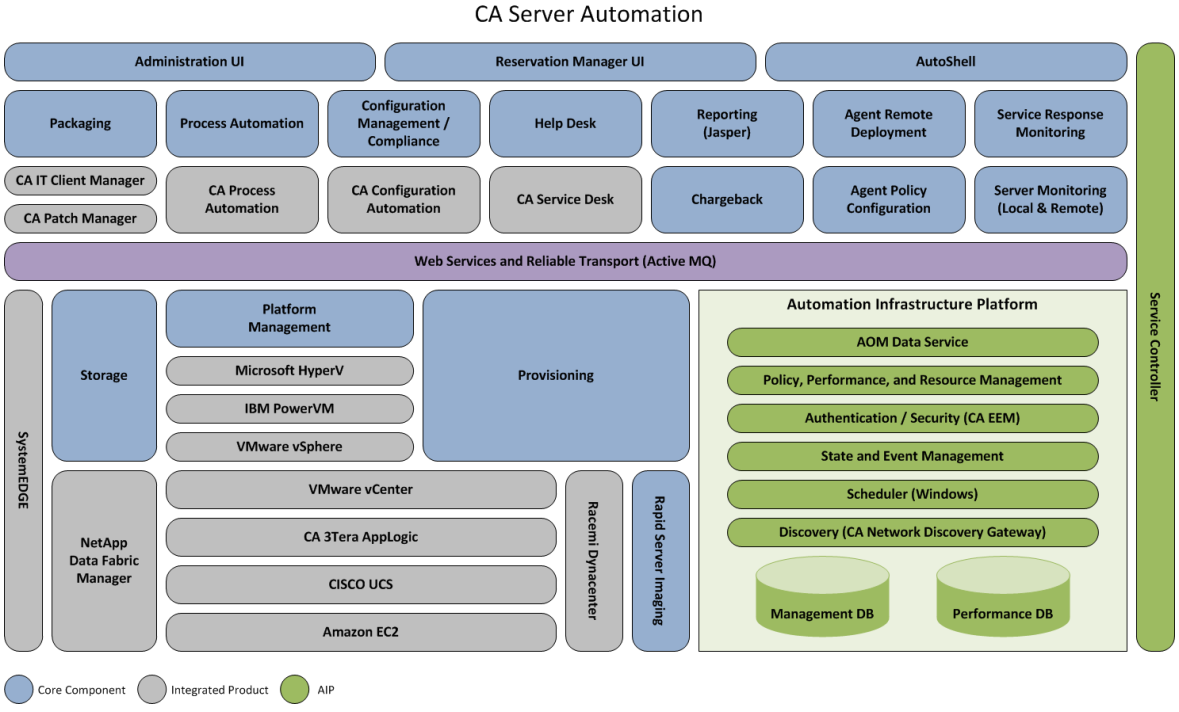

Core components:

Service Controller

Provides a centralized location for identifying the location and status of all components. This centralized location allows for the distribution of CA Server Automation components across multiple management servers, if necessary. In some circumstances, the components must be installed on the same server as the integrating technology. For example, you must install the associated components on the CA Software Delivery server. These components must also be registered back to one main service controller for proper functioning.

Provisioning

Provides the product with a uniform integration point for CA IT Client Manager OS Installation Management technology (OSIM), VMware vCenter, Hyper-V, Amazon EC2, AppLogic, and Rapid Server Imaging.

Performance Monitor

Gathers information and performance metrics using agents that collect system metrics. The supported agents are CA Performance Agent, CA NSM Performance Agent and SystemEDGE agent through SNMP. The managed servers are the core servers in the data center infrastructure. After the Management DB is populated with the managed servers, the Performance Monitor begins collecting data.

Resource Manager

Provides the ability to create and modify resources such as services and static IP pools and also sets the management status of each system.

Verification

Integrates with CA Configuration Automation to provide configuration and change management functionality.

Packaging

Integrates with CA IT Client Manager and CA Patch Manager to provide deployment of software packages and patches to managed servers.

Initiation

Integrates with Microsoft Task Scheduler for job scheduling. You can schedule long running or repeated maintenance tasks and actions as jobs.

CA Process Automation

Integrates with CA Process Automation for scheduling, setting, monitoring, and automation of IT processes. Provides visualization of your processes so you can see exactly where you are in a process at any point in time.

Management DB

Provides a common data repository that stores information about all managed objects. For example, server information, service relationships, service thresholds, rules and actions, events, credentials for other components such as CA Configuration Automation or CA Software Delivery, data center-level polling and recording interval, and data center-level thresholds and lag.

Performance DB

Provides a repository that stores all the metrics collected by the Performance Monitor. Also stores which metrics are collected from which servers, values of those metrics (aggregated over time), server-level polling and recording interval, and server-level thresholds lag (for overall server utilization).

Policy

Analyzes the collected performance 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 take to resolve a particular issue in advance and the policy component uses your parameters to make intelligent decisions. After the server or service is identified, you can take various actions to resolve the situation. For example, you can submit the job to CA Software Delivery to deliver software packages to a remote target server, run custom scripts, provision new systems, and complete many more corrective actions.

Help Desk

Provides integration with CA Service Desk Manager to support opening, updating, and monitoring the status of help desk tickets.

Reporting

Provides reporting capability using the data stored in the Management and Performance Databases.

Event Management

Captures all events generated by CA Server Automation components and provides SNMP forwarding, which can be used to forward events to any CA or third-party product capable of receiving SNMP traps.

Authentication

Integrates with CA EEM to manage all authentication requests and authorization, and provides common access policies.

Reservation Manager

Provides a self-service resource reservation system for end users through a web-based interface. Users can quickly and securely reserve, configure, and provision physical and virtual servers without administrator intervention.

SNMP Trap Receiver

Listens and converts incoming SNMP traps to events (indications) and delivers them to various components.

State Engine

Propagates health state information gathered by agents up a hierarchy of CA Server Automation entities.

Storage

Provisions new or additional storage to virtual and physical systems through integration with NetApp Provisioning Manager.

Platform Management Modules (PMMs)

Provide monitoring and management interfaces for the following virtualization platforms that CA Server Automation integrates with:

- Hyper-V
- IBM PowerVM Logical Partitions
- VMware vCenter
- Cisco UCS
- Solaris Zones

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 Server Automation uses the Management DB to store the following information:

- Server information
- Service relationships
- Service thresholds

- Rules and actions
- Events
- Credentials for other components

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

Performance DB

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

CA Server Automation 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 Server Automation also uses the data in this database and user-created rules to make logical business decisions.

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

User Interface

You can use the CA Server Automation web-based user interface to manage your data center from a central location. The web-based interface lets you use the functions of the components embedded in CA Server Automation without opening the component interfaces separately.

For example, you can use CA Service Desk Manager for issue management from the CA Server Automation 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 opening CA EEM.

You have the option to access user interfaces or integrated products directly to perform more advanced functions. For example, you can open CA EEM to use native security. You can log in to the server where the integrated product is installed to access its user interface, or you can go to the Administration, Configuration page in the CA Server Automation user interface, select a component, and launch the component Home page.

Access the User Interface

Access the user interface to discover and provision systems, to deploy software applications, to create policy, to schedule jobs, and so on. The Start menu shortcuts are only available on the CA Server Automation server. Access the server directly and use the Start menu to access product features such as the user interface and CLI command window. Users that access 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 Server Automation, Launch CA Server Automation on the CA Server Automation server.

   The CA Server Automation login page appears at the following URL:

   `https://servername:port`

   servername

   Identifies the server where the CA Server Automation User Interface is installed.

   port

   Identifies the Apache Tomcat Server port on which the server is listening.

   Default: 8443

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

2. Enter your login credentials and click Log In.

   The user interface appears with the Dashboard displayed by default.

Chapter 3: Managing Users

This section contains the following topics:

User Access Control

CA EEM secures all communication between CA Server Automation components. You can select one of the following configurations:

- Active Directory
- Native Security

Note: For more information about configuring external directories, see the *CA EEM Getting Started* and *Online Help* installed with CA EEM. 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 Server Automation or CA EEM.

CA Server Automation 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.

Native Security

Native Security lets the CA EEM administrator create users, user groups, and policies specifically for CA Server Automation 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 Server Automation

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 Server Automation, 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 Server Automation.

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 Embedded Entitlements Manager Log In window appears.

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

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

   The User Name field is populated with EiamAdmin.

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

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

Create CA EEM Users

To give a user access to CA Server Automation, 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 Server Automation. You can also add users by referencing an external directory.

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

To create CA EEM users

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

   The Users page is selected by default.

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

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

   All CA Server Automation users are listed in a hierarchical tree in the Users pane.

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

   The New User pane appears on the right.

5. Enter the user ID for this user in the User Name field and click Add Application User Details in the User Details pane.

6. Select the application group from the Available User Groups box in the Application Group Membership pane, and click the right arrow  .

   The application group is added to the Selected User Groups.

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

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

   A confirmation message appears below the Users pane.

Create Default User Groups

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

To create user groups

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

   The Users page is selected by default.

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

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

3. Click New Application Group in the left pane.

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

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

   The new Application User Group is created.

Password Management

User credentials are essential for the communication between CA Server Automation components. CA Server Automation stores user and password information internally. When you change passwords of external components or applications CA Server Automation integrates with, change these passwords in CA Server Automation for consistency. Otherwise, CA Server Automation 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 Server Automation.

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 Server Automation

1. Navigate to Start, Programs, CA, CA Server Automation, CA Server Automation 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 Server Automation 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 Server Automation 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 Server Automation

1. Navigate to Start, Programs, CA, CA Server Automation, CA Server Automation 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 Server Automation works as expected.

Change the System User Password for Native Security

CA Server Automation 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 Server Automation. If you change the password for sys_service in CA EEM later, also change it in CA Server Automation to help ensure that all CA Server Automation 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 Server Automation

1. Navigate to Start, Programs, CA, CA Server Automation, CA Server Automation 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 Server Automation works as expected.

Change the System User Password for Active Directory Security

If your CA Server Automation installation is configured to connect to Active Directory, the user who installs CA Server Automation is automatically registered with CA EEM. This registration allows CA Server Automation to authenticate users from the Active Directory domain. If the user password changes, users cannot log in to the CA Server Automation 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 Server Automation, CA Server Automation Command Prompt.

   The command prompt appears.

8. Enter the following command:

   ```
   dpmutil –set –sysuser
   ```

   Sysuser is the same user who installs CA Server Automation. 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 Server Automation 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.

Tasks for External Directory Users

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

Create User Groups

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

To create user groups

1. Click Administration.

   The Administration page appears.

2. Click User Groups.

   The User Groups page appears.

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

4. (Optional) Type a Description.

5. Click Save.

   The new user group appears in the left pane.

Assign Users to Groups

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

To assign users to groups

1. Click Administration.

   The Administration page appears.

2. Click User Group.

   The User Groups page appears.

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

   A submenu appears.

4. Select the Membership submenu.

   A series of membership panes appears.

5. Enter the user name to add in the Value text box, and click Search.

   The search results appear in the Available User/User Group pane or a message notifies you that no match was found. If you are unsure of the user name, you can search for users or user groups. (see page 38)

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

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

7. Click Save to finish adding users.

   Users are granted the access privileges of their user group.

Remove Users or User Groups from a User Group

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

To remove users or user groups from a user group

1. Click Administration, then Configuration.

   The Configuration page appears.

2. Select User Groups.

   The User Groups menu appears on the left pane.

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

   A submenu appears on right pane.

4. Select the Membership submenu.

   A series of membership panes appears.

5. Select the user or user group to remove from the Selected User/User Group pane and click the left arrow  .

   The user or user group is moved to the Available User/User Group pane.

6. Click Save when you finish removing users and user groups.

Assign User Groups Access Rights to Services

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

Set User Group Permissions

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

To set user group permissions

1. Click Administration.

   The Administration page appears.

2. Click User Group.

   The User Groups page appears.

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

   The Privileges page appears.

4. Select the functional areas or specific functions for which you want to grant or restrict access, and click Save.

   The user permissions are updated.

Delete User Groups

You can delete user groups that you no longer need.

To delete user groups

1. Click Administration.

   The Administration page appears.

2. Click User Group.

   The User Groups page appears.

3. Right-click a User Group, and select Delete User Group.

   The user group is removed.

Search for Users or User Groups

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

To search for users or user groups

1. Click Administration.

   The Administration page appears.

2. Click User Group.

   The User Groups page appears.

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

   The user group page appears in the right pane.

4. Click Membership.

   The User/User Group page appears.

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

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

Chapter 4: Configuring Resources

Depending on your implementation, perform the tasks described in one or more of these sections after installing CA Server Automation.

This section contains the following topics:

Discovery

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

Note: Add the SNMP credentials before starting discovery.

If discovery was already done, discover the system again.

Add Proxy Servers

During installation, you can set up a proxy server for use by integrated components. Proxies are required to access AppLogic grids that reside on public IP networks and EC2 AMIs (Amazon Machine Images). You can add multiple proxy servers using the Administration tab.

Follow these steps:

1. View the Administration tab, and select Proxy Servers from the Configuration menu.

   The Proxy Servers page opens showing the currently configured proxies.

2. Click + (Add) in the Proxy Servers toolbar.

   The Proxy Server Configuration panel opens.

3. Enter the server connection parameters and credentials and click OK.

   The server is added to the list of available proxies for use by integrated components.

More information:

Add AppLogic Servers (see page 304)

Configuration Requirements for Managing Environments

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

Discover Environments

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

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

- Open the Administration tab from the user interface, change to Configuration, Provisioning and select the appropriate server type to add credentials and configure the AIM. CA Server Automation 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 Server Automation manager. This option lets you manually perform the required steps.

Consider the following guidelines

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

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

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

1. Run the NodeCfgUtil.exe utility on a Windows AIM Server to update the configuration data for the corresponding AIMs.

   The NodeCfgUtil.exe utility stores the data in zone.cfg, vc.cfg, ucs.cfg, lpar.cfg, mscs.cfg, or esad.cfg.

2. Open the user interface on the Manager Server and click Resources, Data Center in the navigation pane.

3. Right-click and select Management, Discover.

   The discovery options appear.

4. Select one of the following:

   - Discover a system
   - Discover a network

   The corresponding dialog opens.

5. Enter a system name of a vCenter Server, Solaris Zones Server, Microsoft Cluster, Cisco UCS Server, or HMC/IVM Server that you want to manage. Alternatively, you can enter network properties for the discovery process. Click OK.

   CA Server Automation starts the discovery process.

   The discovered resources appear in the Explore pane.

More Information

How to Configure the vCenter Server Management Components (see page 616)

Configure AIMs with NodeCfgUtil in Dialog Mode

NodeCfgUtil lets you modify the AIM configurations for vCenter Server, IBM LPAR, Solaris Zones, Microsoft Clusters, Cisco UCS Server. When you use the utility in dialog mode, you can configure which nodes the appropriate AIMs manage.

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

Follow these steps:

1. Log in as Windows Administrator, open Windows Explorer on the system on which the AIM resides, change to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory, and start NodeCfgUtil.exe. NodeCfgUtil discovers the installed AIMs and lists only those AIMs in subsequent dialogs.

2. Enter **1** to add a new managed node.

   NodeCfgUtil can only discover installed AIMs.

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

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

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

Examples

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

```
****  Choose Managed Node  ****
1. IBM LPAR
2. Solaris Zones
3. VMware vCenter
4. Cisco UCS
5. Microsoft Cluster
0. Go Back to Previous Menu
******************************
Enter choice: 3

Enter following information for the VMware vCenter Node...

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

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

CAAC1016 Authenticating, please wait...

CAAC1019 Authentication SUCCESSFUL.

CAAC1023 Added Node Successfully.

Press any key to continue . . .
```

Configure AIMs with NodeCfgUtil in Command Mode

NodeCfgUtil.exe lets you modify the AIM configurations for IBM LPAR, Solaris Zones, VMware vCenter, Microsoft Clusters, or Cisco UCS. When you use the utility in command mode, you can only add managed nodes to an AIM configuration.

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

This command has the following format:

```
(1) nodecfgutil -help
(2) nodecfgutil {lpar|zone|mscs} -u user -p password -h
{pvmname|hostname|cluster_name}
(3) nodecfgutil lpar -u user -p password -h pvmname -s serial
(4) nodecfgutil {vc|ucs} -u user -p password -h hostname -t port -c protocol
```

-help

Displays the usage information about the console.

lpar|ucs|vc|zone|mscs

Specifies the virtual or physical environment.

-u *user*

Specifies the name of an administrative user.

-p *password*

Specifies the password of that user.

-h *hostname*

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

-h *pvmname*

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

- h cluster_name

Specifies the name of the MSCS cluster.

-s *serial*

Specifies the HMC/IVM serial number.

-t *port*

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

-c *protocol*

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

Return codes: 0 success, -1 failure

To configure AIMs with NodeCfgUtil in command mode

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

   The command prompt appears.

2. Enter one of the following commands:

   ```
   (1) nodecfgutil -help
   (2) nodecfgutil {lpar|zone|mscs} -u user -p password -h {pvmname|hostname|mscs}
   (3) nodecfgutil lpar -u user -p password -h pvmname -s serial
   (4) nodecfgutil {vc|ucs} -u user -p password -h hostname -t port -c protocol
   ```

   (1) Displays the usage information about the console.

   (2) Authenticates and stores the passed credentials for Solaris Zones, LPAR, or MSCS.

   (3) Authenticates and stores the passed credentials for LPAR.

   (4) Authenticates and stores the passed credentials for vCenter or Cisco UCS.

   The utility writes a configuration file for vCenter Server (vc.cfg), Solaris Zones (zone.cfg), Cisco UCS (ucs.cfg), or LPAR (lpar.cfg) to s*ysedge_InstallLpath**plugins**AIPCommon* directory.

Manage Multiple vCenter AIMs

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

Adding vCenter Server

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

Removing vCenter Server

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

Switching vCenter Server

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

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

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

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

To manage multiple vCenter AIMs

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

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

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

Manage a Renamed vCenter Server

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

To manage a renamed vCenter Server

1. Discover the renamed vCenter Server.
2. Change to Administration, Configuration, vCenter Server.

   The vCenter Server pane and the vCenter AIM Servers pane appear.
3. Click the plus (+) button to add the renamed vCenter Server to the vCenter Server pane. Use valid credentials of your vCenter Server environment.

   The renamed vCenter Server appears in the list in normal state.
4. Change to the vCenter AIM Servers pane and select the 1:1 relationship between the vCenter AIM Server and the vCenter Server (previous name). Click the Configure button.

   A dialog appears displaying the server names of that relationship.
5. Select the renamed vCenter Server from the drop down menu, and click OK.

   The new relationship is active and CA Server Automation starts to manage the vCenter Server with the new name.

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

Requirements for Solaris Zones Management

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

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

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

Requirements for LPAR Management

Verify that the following components are installed:

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

How to Configure AIX NIM Imaging

To provision AIX NIM images, use the following process:

1. Install the NIM Adapter on the AIX NIM server.
2. Edit the ca_post_install.sh script file.
3. Configure the NIM master server. Configuration can be done during CA Server Automation installation or after installation using the Administration, Configuration page.

Install NIM Adapter on AIX NIM Server

You can install the NIM adapter using the graphical interface or a command line text console.

To install the NIM adapter

1. Insert the installation media into the computer, navigate to the DVD2\Installers\AIX_aix\NIM directory and copy ca-nim-adapter.AIX to the AIX NIM server.
2. Enter the following command on the AIX NIM server:

   ```
   ./ca-nim-adapter.AIX
   ```

   If X-windows and DISPLAY are configured on the computer where the AIX UNIX terminal is open, a graphical interface installer is launched. Otherwise, a command line interface installer is launched.
3. Press Next.

   The License Agreement appears.
4. Read the License Agreement and select I agree.

   The Installation directory option appears.
5. Specify the directory where you want to install, and click Next.
6. Review the installation path, and click Install Product.

   Post-installation instructions appear.
7. Click OK to exit the installer.

Note: The script file *install_path*/imaging/etc/ca_post_install.sh must be updated to configure CA Server Automation to work with NIM. The hash password must match the hash password on the target AIX computers. Before you start the CA Server Automation NIM adapter daemon, read and update the ca_post_install.sh script, which contains instructions on setting the appropriate options.

Edit the ca_post_install.sh script File

You can edit the ca_post_install.sh script file to set the hashed password and to increase the size of the /tmp and /opt filesystems.

Update the Hashed Password Variable

The IBM AIX installer leaves an empty root password after a NIM client is imaged. To prevent empty root passwords, update the ca_post_install.sh script file before using it. Set the hashed password (DES format) and update the HASH_PASSWORD variable for NIM clients to use.

To update the HASH_PASSWORD variable

1. Access a system that is configured with the password with which to configure your NIM clients.
2. Change to the /etc/security directory and open the passwd file.

   The hashed root password entry in the passwd file resembles the following:

   ```
   root:
       password = YmB7AkapuLf8/s
   ```
3. Copy the hashed root password, change to the *install_path*/imaging/etc directory, and open the ca_post_install.sh script file.
4. Paste the password into the HASH_PASSWORD variable in the ca_post_install.sh script file.
5. Save the file and exit.

Increase the size of the /tmp and /opt filesystems

If there is insufficient space in the /tmp or /opt filesystems the deployment of agents to the NIM clients can fail. The defaults selected are too small for the agents to install, so increase them to at least 400 MB for /tmp and 700 MB for /opt. If you do not have NIM scripts that increase the filesystem sizes, uncomment the chfs commands in the ca_post_install.sh script. Uncommenting these lines enables the NIM adapter to increase the file systems size by using this script as a NIM script resource after the NIM client has been imaged.

Note: Do not enable these lines unless you are sure that your scripts do not already enable these lines.

To increase the size of the /tmp and /opt filesystems

1. Change to the *install_path*/imaging/etc directory and open the ca_post_install.sh script file.
2. Uncomment both chfs commands in the script by removing the leading # character.

   The commented line looks resembles the following example:

   ```
   #chfs -a size=$OPTFSSIZE /opt
   ```

   The uncommented line resembles the following example:

   ```
   chfs -a size=$OPTFSSIZE /opt
   ```

   Note: Optionally change the OPTFSSIZE and TMPFSSIZE variables to higher values than the default but do not set them any lower.
3. Save the file and exit.

Start the NIM Adapter Daemon

After you configure the ca_post_install.sh script file, start the NIM adapter daemon by running the following command:

```
install-path/imaging/bin/canimstart.sh start
```

Configure NIM Master Server

You can configure NIM master servers after CA Server Automation installation.

To configure NIM master servers

1. Configure the NIM adapter on the NIM master server.
2. Log in to the CA Server Automation user interface.
3. Click Administration.

   The Administration page appears.
4. Click Configuration.

   The Configuration page appears.
5. In the left pane, click NIM Master Server.

   The NIM Master Server page appears.
6. Click + (Add).

   The Add NIM Master Access Credentials dialog appears.
7. Add credentials, and click OK.
8. Click Validate to verify the connection status.

Synchronize NIM Master Servers

NIM master server resource and property synchronization is based on the CONFIG_KEY_IMG_IMAGELIST_SYNC_INTERVAL key in *install_path*\conf\caimgconf.cfg. The default setting is 12 hours. You also can synchronize on demand.

Note: caimgconf.cfg changes require a restart of the CAAIPApache service.

To synchronize NIM master servers on demand

1. Log in to the CA Server Automation user interface.
2. Click Administration.

   The Administration page appears.
3. Click Configuration.

   The Configuration page appears.
4. In the left pane, click NIM Master Server.

   The NIM Master Server page appears.
5. Select one or more NIM servers.
6. Click >> (Refresh NIM server properties).

Dynamic NIM Machine Resource Support

CA Server Automation provides a web service method to create NIM machine resources dynamically in a NIM environment. An IP address is provided to the method to determine which NIM master has the NIM network resource to support it. The method only creates the NIM machine resource if it does not exist for that address.

The resources are created using the following convention:

ca_*UUID*

UUID is a randomly generated UUID. The web service method that deletes a NIM machine resource only deletes NIM machines created by CA Server Automation. If the NIM master is removed from CA Server Automation configuration and then added back later, all NIM machines created by the web service before the configuration change are no longer considered for deletion. All CA Server Automation creation records are removed when the NIM master is removed from configuration.

The Imaging Service exposes the web services but only external consumers (such as Reservation Manager) decide how to use them.

Specific required and optional properties are used to create a NIM machine. CA Server Automation uses default properties defined in the NIM Machine group in the caimgconf.cfg file. Properties are global and apply to all NIM masters. You can update these values to suit your environment.

Environment requirements for the web services include:

- NIM network resources must already be defined in the NIM environment.
- NIM network resources cannot have overlapping network address ranges, including network resources that are defined on the same NIM master and across multiple NIM master servers.
- IP addresses must be DNS-resolvable.
- NIM machines are created without a CPUID.

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>

More information:

Configure the Cisco UCS AIM from the Command Line (see page 55)

Configure the Cisco UCS AIM from the Command Line

You can use the general AIM Configuration Utility to specify the connection parameters that the AIM requires to communicate with UCS Manager. You can run the utility on all servers that have SystemEDGE and the Cisco UCS AIM installed. The utility is required on remote servers that have the AIM without a CA Server Automation manager installed.

Note: You also can register the UCS Manager AIM using the CA Server Automation user interface. To register the UCS AIM, click Administration, click Cisco UCS Servers, and click + (Add) in the UCS AIM Servers toolbar.

The following application scenarios are possible using the AIM Configuration utility:

- Change the Cisco UCS user credentials on the Cisco UCS Manager; then update the Cisco UCS Manager connection parameters.
- During Cisco UCS AIM installation, skip the configuration step; then specify connection parameters.

The general AIM Configuration Utility (nodecfgutil.exe) is a command line utility located in the Install_Path\SystemEDGE\plugins\AIPCommon directory.

To configure the UCS AIM from the command line

1. Open a command prompt, change to the *Install_Path*\SystemEDGE\plugins\AIMCommon directory, and enter the following command:

   ```
   nodecfgutil.exe
   ```

   The utility displays the usage information.
2. Follow the directions in the menu to add, update, or remove Cisco UCS Manager access information. The access information is encrypted and stored in the *Install_Path*\SystemEDGE\plugins\AIPCommon\ucs.cfg file.
3. The Cisco UCS AIM automatically acquires the changes to monitor a new Cisco UCS Manager instance or to remove monitoring of an existing Cisco UCS Manager instance without recycling the SystemEDGE agent.

JumpStart

The JumpStart adapter integrates with CA Server Automation so that you can provision Solaris systems with images. The JumpStart adapter searches the system for information to determine what images are available on the server. You can select one of the available images and submit an imaging job against a specific target computer.

The CA Server Automation documentation assumes that you are familiar with the JumpStart solution. All requirements and restrictions imposed by Oracle on JumpStart server technology are valid for CA Server Automation.

Overview

This section describes the steps to create an installation image that can be used by JumpStart and the steps required to configure the system after the operating system is installed.

JumpStart server refers to the server that stores the operating system images that are available in your environment. *JumpStart client* refers to the systems that can be provisioned with one of the operating system images stored on the JumpStart server.

Multiple installation and boot servers can be configured.

Solaris 10 x86 clients with service processors that are IPMI compliant can be provisioned as long as you also have a properly configured Solaris DHCP server.

One JumpStart adapter must be installed and configured on each JumpStart boot server to provide support for multiple JumpStart adapter environments.

JumpStart Prerequisites

The prerequisites for using the JumpStart solution include:

- The server must run on Solaris 10 SPARC or Solaris 10 x86.
- JumpStart boot servers must exist on the same subnet of each SPARC client to be imaged.
- All host names must be configured as static with the DNS server regardless of client architecture.
- Protocols used for configuration cannot be blocked by firewalls.
- The JumpStart solution requires RARP/BOOTP/TFTP protocols for SPARC systems.
- The JumpStart solution requires PXE/DHCP/TFTP for Solaris 10 x86 based systems.
- SPARC and Solaris 10 x86 systems use Network File System (NFS) to access remote JumpStart installation images.

- The full Solaris 10 media is required for server installation.
- Initial installation only is supported, but upgrades of Solaris images are not supported.
- Only one JumpStart configuration directory and one profile/rules directory per supported Solaris version is allowed.
- SPARC client computers must already be running Solaris to be reimaged for dynamic use.
- SPARC client computers must be preconfigured for Secure Socket Shell (SSH) root access to allow rebooting.
- Solaris 10 X86 client computers must have service processors that are Intelligent Platform Management Interface (IPMI) 1.5 or 2.0 compatible.
- Each service processor must have a static IP address configured.
- The IPMI feature for the service processor must be enabled and configured in the BIOS.
- Each service processor must be preconfigured so that it can be reached from the public network.
- The DHCP server must be configured for SSH access to allow the CA Server Automation JumpStart adapter to communicate with it during JumpStart x86 provisioning requests.

JumpStart Adapter Installation

Use one of the following methods to install the JumpStart adapter:

- From the command line
- From a response file

Install the JumpStart Adapter from the Command Line

To install the JumpStart adapter, run ca-jumpstart-adapter.Solaris (or ca-jumpstart-adapter.SolarisIntel for x86 installs). Follow the installation prompts, and choose either the default location (/opt/CA/ServerAutomation) or another location.

Install the JumpStart Adapter using a Response File

To create a response file, run the interactive installation process to produce the response file.

To create a response file

```
./ca-jumpstart-adapter.Solaris -a ca-jumpstart-adapter.Solaris.@pif -r resp.txt
```

To use a response file to do a silent installation

```
./ca-jumpstart-adapter.Solaris -r resp.txt
```

Install Imaging on a JumpStart Server Using the Text Terminal Console

You can install the adapter interactively using the text terminal console.

To install the adaptor using the text terminal console

1. Insert the installation media into the computer, navigate to the DVD2\Installers\Solaris_sparc\JumpStart or DVD2\Installers\Solaris_x86\JumpStart directory, and copy ca-jumpstart-adapter.Solaris or ca-jumpstart-adapter.SolarisIntel, respectively, to the JumpStart server. If you use an ftp client, copy the file in binary format, and also set execute permissions on the file.
2. Enter the following command on the JumpStart server:

   ```
   ca-jumpstart-adapter.Solaris or ca-jumpstart-adapter.SolarisIntel
   ```

   The console appears and prepares for installation.
3. Press Enter.

   The License Agreement appears.
4. Scroll to the bottom of the License Agreement.
5. Tab to "I Agree", and press Enter.

   The installation folder option appears.
6. If the default location is acceptable, tab to Next. If not, specify the installation location, tab to Next, and press Enter.
7. To accept the selections and start the installation, tab to "Install Product". To visit a previous screen, tab to Previous. To cancel the installation, tab to Cancel and press Enter.

Uninstall a JumpStart Adapter

The JumpStart adapter uninstaller is located in the *install_path*/Uninstall directory. If the default installation location is selected, the uninstaller is located in /opt/CA/ServerAutomation/Uninstall.

To perform a silent uninstall (no response file needed)

```
./uninstall.ca-jumpstart-adapter -s
```

To perform an interactive uninstall

```
./uninstall.ca-jumpstart-adapter
```

JumpStart for Solaris

To deploy Solaris images using JumpStart, first install the CA Server Automation JumpStart adapter on the Solaris JumpStart server. The installation procedures are described in the JumpStart Adapter Installation section. After you have installed the adapter, manually edit the cajmpst.cf file.

Configure Solaris DHCP Servers Using the dpmutil Utility

After you install the Solaris JumpStart server, run the dpmutil command-line utility to configure the Solaris Dynamic Host Configuration Protocol (DHCP) servers and enable provisioning on Solaris x86 computers. You must have a user name with administrator privileges to run dpmutil.

To add a DHCP server with the dpmutil command

1. Log in to the CA Server Automation server using your administrator user name and password.
2. Type the following command at a command prompt and press Enter:

   ```
   dpmutil -set -—dhcp-u
   ```

   The utility prompts you for your CA Server Automation user name and password.
3. Type your user name and password and press Enter.

   The utility prompts for the name of the DHCP host server.

   Note: The DHCP server must be configured for SSH access to allow the CA Server Automation JumpStart adapter to communicate with it during JumpStart x86 provisioning requests.

4. Type the host name and press Enter.

   The utility prompts you for a user name and password for the DHCP server.

5. Type your user name and password and press Enter.

   The DHCP server is configured to allow CA Server Automation to provision x86 computers.

Edit the cajmpst.cf File

You can edit the cajmpst.cf file to specify the location of the JumpStart Configuration server and the JumpStart Profile server.

To edit the cajmpst.cf file (default installation location)

1. Change to the /opt/CA/ServerAutomation/imaging/etc directory and open the cajmpst.cf file in a text editor.

   The file opens.

   Note: /opt/CA/CAM/imaging/etc is the default path. If you selected to install to a different path, navigate accordingly.

2. Navigate to the following lines in the file:

```
# Rules file path (JumpStart profile server) for Solaris 10.
#Solaris_10_Profile_Server = /jumpstart/ca/profile/Solaris_10

# sysidcfg file path (JumpStart configuration server) for Solaris 10
#Solaris_10_Config_Server = /jumpstart/ca/profile/Solaris_10
```

   Note: You must include a sysidcfg file in the top-level file path. You also can create subdirectories in which to place additional target-specific sysidcfg files. Subdirectory names can use the MAC address (lowercase with no colons) or user-defined host name to identify the target servers.

```
# Rules file path (JumpStart profile server) for Solaris 9.
#Solaris_9_Profile_Server = /qa/jumpstart/Solaris_9

# sysidcfg file path (JumpStart configuration server) for Solaris 9
#Solaris_9_Config_Server = /qa/jumpstart/Solaris_9

# Rules file path (JumpStart profile server) for Solaris 8.
#Solaris_8_Profile_Server = /qa/jumpstart/Solaris_8

# sysidcfg file path (JumpStart configuration server) for Solaris 8
#Solaris_8_Config_Server = /qa/jumpstart/Solaris_8
```

3. Remove the # characters before the two variables containing path information that relate to your version of Solaris, and update the location of the JumpStart servers. For example, if you are running Solaris 9, edit the following variables:

   ```
   Solaris_9_Profile_Server = <path>
   Solaris_9_Config_Server = <path>
   ```

   The updated section of the file resembles the following:

   ```
   # Location of JS profile server for Solaris 10.
   #Solaris_10_Profile_Server = /qa/jumpstart/Solaris_10

   # Location of JS configuration server for Solaris 10
   #Solaris_10_Config_Server = /qa/jumpstart/Solaris_10
   # Location of JS profile server for Solaris 9.
   Solaris_9_Profile_Server = <path>

   # Location of JS configuration server for Solaris 9
   Solaris_9_Config_Server = <path>
   ```

   Note: If you are running both versions of Solaris, edit all four of the variables that contain path information. Solaris 8 variables can be ignored unless you plan to provision using Solaris 8.

4. Save the file and exit.

   The edit is complete.

Copy the post_install.sh File

CA Technologies provides a post_install.sh finish script that is required. Do not remove any original content. You can add content and modify specific parameters that are identified.

To use this file, change to the /opt/CA/ServerAutomation/imaging/etc directory (or the user-selected installation path), and copy the post_install.sh file to the directory that is specified in your JumpStart rules file for Solaris 8, 9, or 10.

How To Create a Solaris 8 Image

Creating the Solaris 8 image involves extracting an installable image from the CDs, adding the packages that are required for integrating the client with CA Server Automation, adding the patches that these packages require, and modifying the configuration files.

More information:

Prepare Your Directories (see page 62)
Extract an Installable Image from the Media (see page 63)
Add Packages to the Image (see page 64)
Add Patches to the Image (see page 65)
How To Modify Configuration Files (see page 66)
Configure SSH for JumpStart (see page 71)

Prepare Your Directories

Configure directories for JumpStart imaging. The image parent directory will contain the operating system images that will be installed on the JumpStart client computers. The config parent directory will contain the configuration files that JumpStart uses to install the operating system and configure a JumpStart client. Edit the directory values specific to your site, create your parent directories and then share them.

To prepare your directories

1. Edit the following code and command examples which contain values that are specific to your site. This is not an all inclusive list.

   image_hostname

   Specifies the host name of the JumpStart server.

   client_hostname

   Specifies the host name of the JumpStart client.

   image_parent

   Specifies the path of the directory on the JumpStart server that contains the operating system image directories.

   Example: /jsimages

   config_parent

   Specifies the path of the directory on the JumpStart server that contains the JumpStart configuration files.

   Example: /jumpstart

   sol_8 is

   (Optional) Specifies a subdirectory that can be replaced or removed.

   Your site-specific settings are set.

2. Create the installation directory and, if necessary, the configuration directories as follows:

   ```
   mkdir –m 755 /image_parent/sol_8
   ```

   ```
   mkdir –m 755 /config_parent
   ```

   ```
   mkdir –m 755 /config_parent/bin
   ```

3. Change to the /etc/dfs/dfstab file and add these lines:

   ```
   share –F nfs –o ro,anon=0 /image_parent/sol_8
   ```

   ```
   share –F nfs –o ro,anon=0 /config_parent
   ```

   The directories are shared.

4. Enter this command:

   ```
   shareall
   ```

   The shared directories are activated.

Extract an Installable Image from the Media

Extract an installable image from the Solaris media.

To extract an installable image from the media

1. Insert the Solaris 8 Software 1 CD into the CD-ROM drive, mount the CD if it is not automatically mounted, and enter these lines in the Command Prompt window:

   ```
   cd /cd_drive/cdrom0/s0/Solaris_8/Tools
   ```

   ```
   ./add_to_install_server /image_parent/sol8
   ```

   The files are extracted from software CD 1.

2. Insert the Solaris 8 Software 2 CD into the CD-ROM drive, mount the CD if it is not automatically mounted, and enter these lines in the Command Prompt window:

   ```
   cd /cd_drive/cdrom0/s0/Solaris_8/Tools
   ```

   ```
   ./add_to_install_server /image_parent/sol_8
   ```

   The files are extracted from software CD 2. After the image is created, the packages must be added to the image so the JumpStart client can be integrated with CA Server Automation.

Add Packages to the Image

Packages are required to integrate the JumpStart client with CA Server Automation. Download the required and optional packages from the www.sunfreeware.com website and add them to the image.

To add packages to the image

1. Download libgcc-3.4.6 into a working directory and unzip the package:

   ```
   cd /working_directory
   gunzip libgcc-3.4.6-sol8-sparc-local.gz
   ```

2. Enter the following commands:

   ```
   pkgtrans libgcc-3.4.6-sol8-sparc-local . all
   cp –r SMClgcc /image_parent/sol_8/Solaris_8/Product
   ```

   The package is converted and the libgcc package is added to the image.

3. Download openssh-5.0p1 into a working directory and unzip the package:

   ```
   cd /working_directory
   gunzip openssh-5.0p1-sol8-sparc-local.gz
   ```

4. Enter the following commands:

   ```
   pkgtrans openssh-5.0p1-sol8-sparc-local . all
   cp –r SMCosh501 /image_parent/sol_8/Solaris_8/Product
   ```

   The package is converted and the SSH package is added to the image.

5. Download openssl-0.9.8h into a working directory and unzip the package:

   ```
   cd /working_directory
   gunzip openssl-0.9.8h-sol8-sparc-local.gz
   ```

6. Enter the following commands:

   ```
   pkgtrans openssl-0.9.8h-sol8-sparc-local . all
   cp –r SMCossl /image_parent/sol_8/Solaris_8/Product
   ```

   The package is converted and the SSL package is added to the image.

7. Download zlib-1.2.3 into a working directory and unzip the package:

   ```
   cd /working_directory
   ```

   ```
   gunzip zlib-1.2.3-sol8-sparc-local.gz
   ```

8. Enter the following commands:

   ```
   pkgtrans zlib-1.2.3-sol8-sparc-local . all
   ```

   ```
   cp –r SMCzlib /image_parent/sol_8/Solaris_8/Product
   ```

   The package is converted and the zlib package is added to the image. After you finish adding packages to the image, add the patches required by these packages to the image.

Add Patches to the Image

Add the patches required by the packages you added to the image and any optional patches. There are also two patches required to integrate the JumpStart client with CA Server Automation. Download the patches from the www.sun.com website and add them to the image.

Note: You must have a Sun Online account to download patches. Log in or register on the self-registration page of the Sun website.

To add patches to the image

1. Download patch 108434-22 from the libc patch page of the Sun website into a working directory:

   ```
   cd /working_directory
   ```

2. Enter the following command:

   ```
   jarunzip -x–xf 108434-22.jar.zip
   ```

   The file is uncompressed.

3. Enter the following command:

   ```
   cp –r 1084340-22 /image_parent/sol_8/Solaris_8/Patches
   ```

   The libc patch is copied to the image.

4. Download patch 112438-03 from the random patch page of the Sun website into a working directory:

   ```
   cd /working_directory
   ```

5. Enter the following command:

   ```
   unzip jar –xxf 112438-03.jarzip
   ```

   The file is uncompressed.

6. Enter the following command:

   ```
   cp –r 112438-03 /image_parent/sol_8/Solaris_8/Patches
   ```

   The random patch is copied to the image.

How To Modify Configuration Files

After you have created the image, create or modify the configuration files to help ensure that the packages are installed, the patches are applied, and any other required configuration is performed when a JumpStart client is provisioned.

More information:

Edit the Order File (see page 66)
Edit the Package Table of Contents File (see page 67)
Edit the Profile File (see page 68)
Edit the Rules File (see page 68)
Edit the Configuration File (see page 69)
Edit the Finish File (see page 70)

Edit the Order File

The Order file indicates which packages are installed with the image and the sequence in which they are installed.

To edit the order file

1. Change to the image parent directory and edit the Order file.

   ```
   /image_parent/sol_8/Solaris_8/Product/.order
   ```

2. Add the new packages to the end of the package list as follows:

   SMClgcc

   Defines the libgcc package and can be listed in any order.

   SMCossl

   Defines the SSL package and can be listed in any order.

   SMCzlib

   Defines the zlib package and can be listed in any order.

   SMCosh501

   Defines the SSH package and must be listed after the other packages—SSL, libgcc and zlib.

   If you are installing optional packages, you must add them in order.

3. Save the file.

Edit the Package Table of Contents File

The Package Table of Contents file contains information about the packages being installed. Each package has an information file that contains information that is required in the Package Table of Contents file.

To edit the package table of contents file

1. Change to the image parent directory and edit the Package Table of Contents file.

   ```
   /image_parent/sol_8/Solaris_8/.packagetoc
   ```

2. Locate the information file for each package, open with a text editor of your choice, and extract the information required for the Package Table of Contents File. For example, the Open SSL information is located in the following directory:

   ```
   /image_parent/sol_8/Solaris_8/Product/SMCossl/pkginfo
   ```

3. Edit the Package Table of Contents File with information for all required and optional packages. Some information is required, but some is optional, such as size values.

4. Save the file.

Edit the Profile File

Update the Profile file with the names of the packages that have been added to the image. The file can be copied from an existing JumpStart image or from the newly created Solaris 8 image. Advanced users may need to create multiple profile files with unique information. A specific profile is associated with an image and one or more computers in the rules file.

To edit the profile file

1. Change to the configuration parent directory and copy the Profile file:

   ```
   cd /config_parent/ca/profile/Solaris_8

   cp /image_parent/sol_8/Solaris_8/Misc/jumpstart_sample/ any_machine profile

   cp –r /image_parent/sol_8/Solaris_8/Misc/jumpstart_sample/* .
   ```

2. Open the Profile file with a text editor and add the new packages to the end of the package list and before the filesys entries:

   ```
   package SMClgcc add

   package SMCossl add

   package SMCzlib add

   package SMCosh501 add
   ```

3. Save the file.

Edit the Rules File

The Rules file includes the names of the Profile and Finish files for the image. The Rules file can be copied from an existing JumpStart image or from the newly created Solaris 8 image. Advanced users may need to write rules to control not only which images can be installed but also which rules start or finish scripts run on specific JumpStart clients. They also can require site-specific finish scripts, which are acceptable if they include a call to bin/post_install.sh.

To edit the rules file

1. Navigate to the configuration parent directory and open the Rules file with a text editor.

   ```
   /config_parent/config_dir/Solaris_8/rules
   ```

2. Write your rule with the following format. Rules are free form.

   ```
   rule value  begin_file  profile_file finish_file
   ```

3. Rename or delete the existing rules file and create one with the following rule for the minimum Rules file:

   ```
   any  profile  bin/post_install.sh
   ```

4. Save the file.
5. Run the check shell script to validate your changes to the Rules file:

   ```
   ./check
   ```

   JumpStart uses the rules.ok file which is created by the check shell script.

Edit the Configuration File

The Configuration file contains the information that is required to allow JumpStart to perform a silent (unattended) installation. The syntax and keywords for the Configuration file are detailed in the Solaris 8 sysidcfg document on the Sun website.

To edit the configuration file

1. Change to the configuration parent directory and open the Configuration file with a text editor.

   ```
   /config_parent/ca/profile/Solaris_8/sysidcfg
   ```

2. Edit and save the file.

Example of Configuration file

This is an example of a Configuration file:

```
system_locale=en_US

timezone=US/Pacific

timeserver=localhost

terminal=sun-cmd

name_service=DNS {domain_name=domain.com name_server=ip_address}

security_policy=none

network_interface=PRIMARY {default_route=ip_address netmask=255.255.255.0
protocol_ipv6=no}
```

Edit the Finish File

The Finish file is executed after the operating system image is installed on the JumpStart client. This script contains setup steps that complete the installation and make the JumpStart client fully operational.

To edit the finish file

1. Navigate to the configuration parent directory and copy the post_install.sh file:

   ```
   /config_parent/config_dir/Solaris_8/bin/post_install.sh
   ```

   ```
   cp $CA_DCA_MANAGER/imaging/etc/post_install.sh
   config_parent/config_dir/Solaris_8/bin/
   ```

   The post_install.sh file provided with CA Server Automation JumpStart Adapter is copied.

2. Open the Finish file with a text editor of your choice and replace the values of the following variables:

   PASSWD

   Sets the password.

   Example: PASSWD=pZWXCvV5eAkJU.

   PATCH_LOCATION

   Specifies the path to the patches.

   Example:
   PATCH_LOCATION=server_hostname:/image_parent/sol_8/Solaris_8/Patches

   IPC Tunables

   Specifies the tunable Solaris parameters. Requires a system reboot for settings to take effect.

   Example:

   SHMMAXv8_HEX=0x400000

   SHMSEGv8_HEX=0x100

   SEMMNIv8_HEX=0x100

   SEMMNSv8_HEX=0x12c

   SEMUMEv8_HEX=0x20

   SEMMNUv8_HEX=0x100

Optional Patches

Specifies any additional patches added to the installation image. Add patchadd statements in this section to add these patches during provisioning:

```
if [ "$OSVER" = "5.8" ] ; then
        echo "${ID}Install SUN patches"
        echo "${ID}mount -f nfs ${PATCH_LOCATION} ${A_ROOT}/mnt"
        mount -f nfs ${PATCH_LOCATION} ${A_ROOT}/mnt

        echo "${ID}patchadd -R ${A_ROOT} ${A_ROOT}/mnt/112438-03"
        patchadd -R ${A_ROOT} ${A_ROOT}/mnt/112438-03

        echo "${ID}patchadd -R ${A_ROOT} ${A_ROOT}/mnt/108434-22"
        patchadd -R ${A_ROOT} ${A_ROOT}/mnt/108434-22

        echo "${ID}umount ${A_ROOT}/mnt"
        umount ${A_ROOT}/mnt
```

The password, patch location, CA IPC tunables and optional patches are set in the Finish file.

Configure SSH for JumpStart

CA Server Automation JumpStart uses the SSH service to communicate with JumpStart clients. CA Server Automation cannot monitor or control JumpStart clients unless this service is functional. If this service is not functional on the JumpStart client, you can install it by either running JumpStart on the JumpStart client manually or manually installing and configuring the service on the JumpStart client. To install and configure SSH using JumpStart, the JumpStart server must know about the client and then the JumpStart process must be started on the JumpStart client.

To configure a client for manual JumpStart provisioning

1. Replace sol_10/Solaris_10 with the path to the image of the highest release operating system in the following command:

   ```
   /image_parent/sol_10/Solaris_10/Tools/add_install_client \
   -s server_hostname:/image_parent/sol_8 \
   -p server_hostname:/config_parent/config_dir/Solaris_8 \
   -c server_hostname:/config_parent/config_dir/Solaris_8 \
   -e ethernet_address client_hostname client_class
   ```

   The add_install_client shell script updates /etc/bootparams with information about the JumpStart client including path names, Ethernet or MAC address, host name, and hardware class.

2. Enter the following command to obtain the Ethernet (MAC) address:

   ```
   ifconfig –a
   ```

   Note: The output of this command displays the MAC address without leading 0s. The add_install_client script requires a MAC address with leading 0s. This MAC address must also match the entry in the /etc/ethers file on the JumpStart server.

3. Enter the following command to obtain the node name. The host name is node name with domain information.

   ```
   uname –n
   ```

4. Enter the following command to obtain the computer hardware name (class):

   ```
   uname –m
   ```

   The JumpStart client is added.

5. Log in as root and enter the following command:

   ```
   reboot “net – install”
   ```

   The JumpStart process is started on the client.

Note: For information about configuring SSH manually, see the Solaris 8 section of www.sunfreeware.com.

Rapid Server Imaging

Rapid Server Imaging (RSI) provides cross-platform and dissimilar hardware provisioning. Functionality includes capturing and deploying images between Physical to Physical, Physical to Virtual, Virtual to Physical, and Virtual to Virtual machines.

Configure the RSI Environment

Before you start provisioning, configure the RSI environment (RSI servers, hypervisors, and networks), and register them in the CA Server Automation user interface.

Follow these steps:

1. Install the RSI server.

   Note: For more information about installing RSI server, see the *RSI Server Installation Guide* on DVD3.

2. Configure the RSI server.

   Note: For more information about configuring RSI server, see the *RSI Server Administration Guide* on DVD3.

3. Configure hypervisors.

   Note: For more information about configuring hypervisors, see the *RSI Server Administration Guide* on DVD3.

4. Configure external networks.

   Note: For more information about configuring external networks, see the *RSI Server Administration Guide* on DVD3.

   Note: You can make any external network available to the RSI environment.

   a. Open a port on the RSI server for RSI agent remote connections.

      Example: port 4105.

   b. If necessary, update your firewall rules to allow traffic on the new port, and update NAT rules to allow redirection to the RSI server.

      Example: https://RSIserver:443.

   c. Install the RSI agent on the target server.

   d. Edit the dpad.ini file on the target server with the appropriate RSI URL.

      Example: https://RSIserver:4105.

   e. Restart the RSI agent.

5. Log in to the CA Server Automation user interface.

6. Select the Administration tab.

7. In the left Configuration panel, scroll down to Provisioning, and click Rapid Server Imaging.

8. In the RSI Servers panel, click + (Add), add RSI servers, and validate that the connection status is green.

9. In the RSI Registered Hypervisors panel, click + (Add) and add hypervisors.

   Hypervisor credentials are validated.

10. In the RSI Registered Networks panel, click + (Add) and add boot and external networks.

    Networks are not validated.

11. In the RSI Depot panel, click click + (Add) and add any additional depots for captured image storage.

How to Deploy Rapid Server Imaging Using CA ITCM

As a System Administrator, you use CA Server Automation to manage the optimal policy-driven provisioning of software images across your physical and virtual server resources.

Rapid Server Imaging (RSI) provides cross-platform and heterogeneous hardware provisioning, physical and virtual server migration, disaster recovery, and image capture and deployment. Images can be deployed across dissimilar hardware with multiple operating environments provided the hardware belongs to the same processor family.

Use the optional integration products, CA ITCM and Software Delivery, to distribute preconfigured packages for RSI Server and RSI agents.

Prerequisites:

1. Install and configure CA ITCM to manage the environment where you want to deploy RSI.

   Note: For more information, see the *CA ITCM Implementation Guide*.

2. During CA Server Automation installation, set the ITCM Domain and Software Delivery Scalability Server settings to match your CA ITCM deployment.

3. Verify that the following CA Server Automation modules are installed and configured to support RSI deployment using CA ITCM:

   - Provisioning Manager
   - Deployment and Configuration Distribution (Scalability) Server
   - Software Delivery Adapter
   - Rapid Server Imaging (RSI)

   Note: For more information, see the *CA Server Automation Installation Guide*.

Important! The supplied RSI Server MWS package can only be deployed to a Red Hat Linux server.

To perform this process, you require the following:

- Access to DVD3 of the CA Server Automation installation media that contains the RSI packages and documentation.
- Administrator privileges for CA Server Automation and CA ITCM
- Administrator credentials for all target servers for RSI Server and Agent deployment
- (Optional) Credentials for hypervisors controlling virtual machines

The following process provides a high-level overview of how to install and deploy RSI using CA Server Automation and CA ITCM:

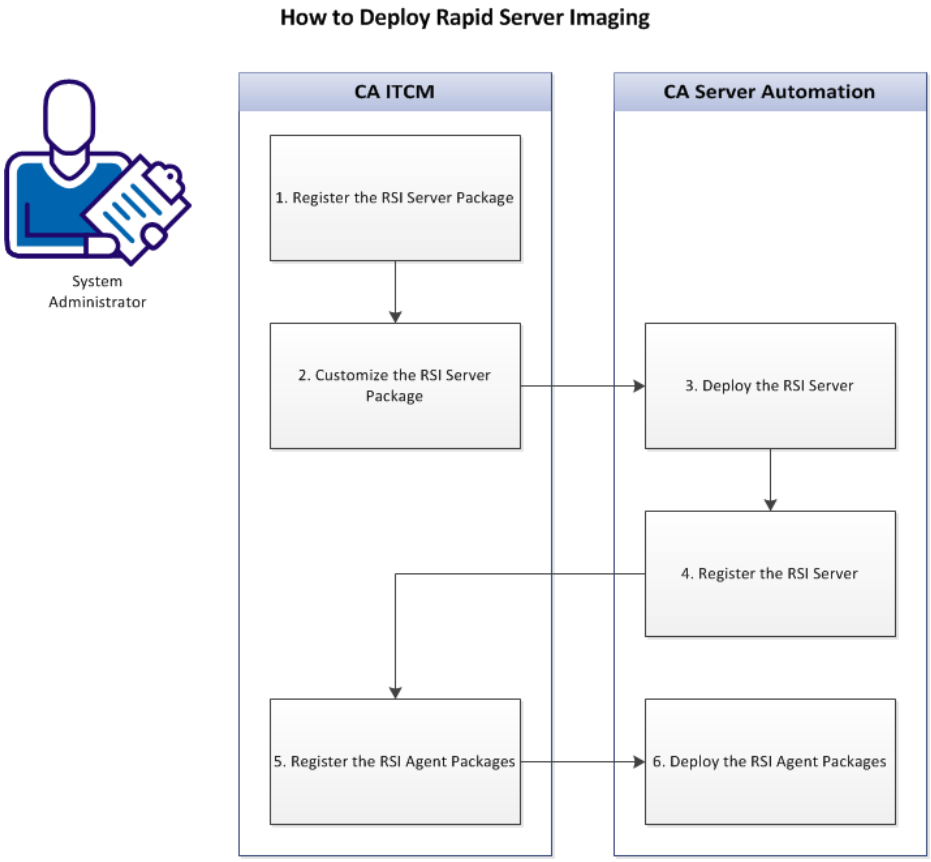

1. Register the RSI Server Package (see page 76)
2. Customize the RSI Server Package (see page 77)
3. Deploy the RSI Server (see page 79)
4. Register the RSI Server (see page 79)
5. Register the RSI Agent Packages (see page 80)
6. Deploy the RSI Agent Packages (see page 81)

Register the RSI Server Package

To use the imaging and software delivery features of CA Server Automation to deploy RSI Server, register the RSI Server package with CA ITCM.

Follow these steps:

1. Copy the 'linux' folder from the following location on the CA Server Automation installation media to a temporary folder on the CA ITCM server:

   DVD3\Racemi\SoftDeliveryPackages\MWS

   Note: The folder requires at least 3.5GB of free space.

2. Copy the entire contents (five folders and one dc-install file) from the following location:

   DVD3\Racemi\DynaCenter_*version*.

3. Place the contents in the following location:

   local_drive:*temporary_folder*\linux\1.vol\DynaCenter_*version*

   Note: Use the zero length file 'COPY CONTENTS HERE' to identify the correct location.

4. Click Start, All Programs, CA, IT Client Manager, DSM Explorer on the CA ITCM server.

5. Expand Software, Software Package Library, All Software in the DSM Explorer Tree.

6. Right-click All Software and select Import, Software Package from the shortcut menu.

   The Register Software Package dialog opens.

7. Enter the following path and click OK:

   local_drive:*temporary_folder*\linux

   The software package is imported into CA ITCM and is listed in DSM Explorer as 'Racemi DynaCenter MWS *version*'.

Customize the RSI Server Package

The RSI Server package supplied in the installation media must be customized for your RSI deployment environment.

The configuration file specifies the following environment parameters:

- Network settings for the RSI Server
- Networks that the RSI Server can access
- Supported operating systems for imaging target servers
- The location of the image storage depot

Follow these steps:

1. In CA ITCM, open the DSM Explorer Tree and navigate to Software, Software Package Library, All Software.

   Right-click the Racemi DynaCenter MWS package and select Unseal.

2. Expand Racemi DynaCenter MWS, Source in the left pane and click Package.

   The right pane displays the folder structure of the package.

3. Edit the config.ini file.

   Uncomment all lines beginning with #, and enter parameter values as appropriate for your environment.

   For example, the parameters in the following configuration file define:

   - A single RSI Server using the loop back address 127.0.0.1
   - Support for imaging target servers running all supported operating systems (Windows and Red Hat Linux are provided by default)
   - Access to target servers using network addresses 10.130.64.0-255
   - Primary network interface eth0 accessible at 10.130.64.162
   - Storage depot accessible at 10.130.64.162 with the following storage paths:
     - Default storage at path /repo/R for components [default]
     - Path /repo/I for agent images and image metadata [instance]
     - Path /repo/images for captured images [image]

```
[general]
oem_configuration = True
agent_addressing = dhcp
os_support = Solaris-sun4u, Solaris-i86pc
mws_address = 127.0.0.1
database_address = 127.0.0.1
[client_networks]
[[10.130.64.0/24]]
gateway = 10.130.64.1
addresses = 10.130.64.0-10.130.64.255
mws_interface = eth0
mws_ip = 10.130.64.162
[storage]
[[default]]
path = /repo/R
server_address = 10.130.64.162
type = component
[[instance]]
path = /repo/I
server_address = 10.130.64.162
type = image_metadata
[[image]]
path = /repo/images
server_address = 10.130.64.162
type = captured_image
```

   Note: For config.ini file parameter details, see "Silent Installation" in the *RSI Server Installation Guide* on DVD3 of the CA Server Automation installation media.

4. In the left pane, right-click Racemi DynaCenter MWS, select seal, and click OK.

   The RSI server package is ready for deployment.

Deploy the RSI Server

Use CA Server Automation to deploy the customized RSI Server package registered in CA ITCM.

1. Click Resources, right-click Data Center in the Explore tree, and select Packaging, Manage Packages.

   The Packages, Packages tab opens.

2. Select the Domain Server that the RSI Server package is registered with.

3. Select the Racemi DynaCenter MWS package from the Available Packages section, and click the down-arrow to move it to the Selected Packages section.

4. Click Save.

   A confirmation message notifies you that the Managed Packages list is updated. The RSI Server package is ready to deploy.

5. Right-click Data Center in the Explore tree, and select Packaging, Deploy Software.

   The Managed Resource section appears.

6. Select the managed server for RSI Server package deployment.

7. Select the Racemi DynaCenter MWS package, select procedure Install, and the domain server to use.

8. Click OK.

   Note: If there is no CA Software Delivery agent installed on the target server, a dialog opens prompting you to install the agent.

9. Enter valid credentials for the target server, the scalability server name, the operating system type, and click OK.

   The panel submits a request to install the RSI Server package to the selected server.

Register the RSI Server

To enable provisioning to the RSI environment, register the RSI Server and optional environment settings in CA Server Automation.

Follow these steps:

1. Log in to CA Server Automation and select the Administration tab.

2. In the left Configuration panel, locate Provisioning, and click Rapid Server Imaging.

3. In the RSI Servers panel, click the + (Add) icon, add the RSI Server, and validate that the connection status is green.

4. (Optional) If you use virtual machines, in the RSI Registered Hypervisors panel, click + (Add) to add hypervisors.

   Note: Hypervisor credentials are validated.

5. (Optional) If you are not using the default RSI network, in the RSI Registered Networks panel, click + (Add) to add boot and external networks.

   Note: Networks are not validated.

6. (Optional) If you are not using the default RSI depot, in the RSI Depots panel, click + (Add) to add any additional depots for captured image storage.

Register the RSI Agent Packages

To enable the provisioning features of CA Server Automation to deploy RSI Agent Software Delivery packages, register the packages in CA ITCM. Each agent package supports a specific operating system and enables the RSI Server to capture and deploy images on servers using that operating system.

Note: Import the package for each operating system individually.

Follow these steps:

1. Click Start, All Programs, CA, IT Client Manager, DSM Explorer on the CA ITCM server.
2. Expand Software, Software Package Library, All Software in the DSM Explorer Tree.
3. Right-click All Software and select Import, Software Package from the shortcut menu.

   The Register Software Package dialog opens.

4. Perform *one* of the following tasks:

   - Enter the path.

     For example, enter the following path for Windows:

     DVD_drive:\DVD3\Racemi\SoftDeliveryPackages\DPADAgent\win

   - Navigate to the appropriate folder for the target operating system in the SoftDeliveryPackages folder on DVD3 of the CA Server Automation installation media. Click Choose to select the folder path, and click OK.

   The RSI agent software package is imported into CA ITCM and listed in the DSM Explorer.

Deploy the RSI Agent Packages

Use CA Server Automation to deploy the RSI Agent packages registered in CA ITCM to the target servers.

Follow these steps:

1. Click Resources, right-click Data Center in the Explore tree, and select Packaging, Manage Packages.

   The Packages, Packages tab opens.

2. Select the Domain Server that the RSI Agent packages are registered with.

3. Select the RSI Agent packages from the Available Packages section, and click the down-arrow to move them to the Selected Packages section.

4. Click Save.

   A confirmation message notifies you that the Managed Packages list is updated. The RSI Agent packages are ready to deploy.

5. Right-click Data Center in the Explore tree, and select Packaging, Deploy Software.

   The Software Deployment panel opens.

6. Select the managed servers for RSI Agent package deployment.

7. For each server, select the RSI Agent package to deploy, select procedure Install, and the domain server to use.

   Note: Select the appropriate RSI Agent for the operating system on each managed server.

8. Click OK.

   Note: If there is no CA Software Delivery agent installed on the target server, a dialog opens prompting you to install the agent.

9. Enter valid credentials for the target servers, the scalability server name, the operating system type, and click OK.

   The panel submits a request to install the selected RSI agent package to each selected server.

   When agent package deployment is complete, you can use the RSI image capture and deploy features of CA Server Automation on the target servers.

How to Setup Rapid Server Imaging for AppLogic

As a System Administrator, you use CA Server Automation to manage the optimal policy-driven provisioning of software images across your physical and virtual server resources.

You use Rapid Server Imaging (RSI) to provide cross-platform and heterogeneous hardware provisioning, physical and virtual server migration, disaster recovery, and image capture and deployment. Images can be deployed across dissimilar hardware with multiple operating environments provided the hardware belongs to the same processor family.

Rapid Server Imaging (RSI) is deployed to AppLogic which acts as a DynaCenter domain for images.

You want to use CA Server Automation to simplify the deployment of RSI images to the DynaCenter domain inside an AppLogic grid.

Prerequisites:

1. Install and configure CA AppLogic to manage the grid where you want to deploy RSI.

   Note: For more information, see the *CA AppLogic Backbone Fabric Controller User Guide*.

2. Deploy Rapid Server Imaging to the AppLogic grid.

   Note: For more information, see the *RSI for AppLogic Installation and User Guide* on DVD3 of the installation media.

The following process provides a high-level overview of how to setup and use CA Server Automation to simplify RSI image deployment to AppLogic:

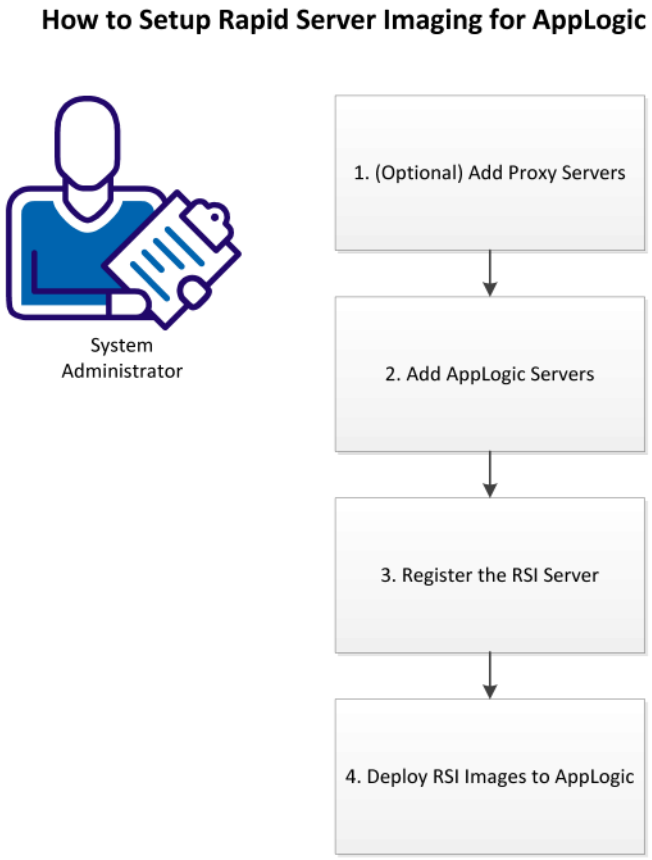

1. (Optional) If the AppLogic server is outside your internal network, Add Proxy Servers (see page 40).
2. Add AppLogic Servers (see page 304).
3. Register the RSI Server (see page 79) deployed to AppLogic to CA Server Automation.

   Note: To deploy images captured by other RSI servers, the RSI server in the AppLogic grid must share their depots.
4. Deploy RSI Images to AppLogic (see page 425).

Storage Provisioning Manager for NetApp

The CA Server Automation Storage Provisioning Manager lets you provision new or additional NetApp storage for virtual and physical servers. Install the CA Server Automation Storage Provisioning Manager component and establish the storage provisioning environment for NetApp before provisioning new or additional datastores and raw storage.

How to Configure the Storage Provisioning Manager

Use the following process to establish the CA Server Automation Storage Provisioning Manager environment for NetApp:

1. Verify that the data center uses NetApp Storage Systems that include NetApp Operations, Provisioning, and Protection Managers.
2. Confirm that the Storage Architect (or equivalent) has created a storage service catalog using the NetApp Provisioning Manager.

   A storage service is a combination of provisioning policies, protection policies, resource pools, and vFiler templates. These objects are preconfigured and then applied as a package to different data sets with different storage needs.

   Note: For more information about NetApp Storage Systems or NetApp Operations, Provisioning, and Protection Managers, see the corresponding NetApp documentation.
3. Define credentials for CA Server Automation that match existing credentials for the NetApp DataFabric Manager server.
4. Configure DataFabric Managers for use with CA Server Automation.

   Note: For more information about configuring DataFabric Managers, see "Configure DataFabric Managers" in this chapter.
5. (Optional) Perform the prerequisites for VMware vCenter ESX server provisioning.

   Note: For more information about these prerequisites, see "Prerequisites for VMware vCenter ESX Servers" in this chapter.
6. (Optional) Perform the prerequisites for supporting Windows physical server attachment during storage provisioning.

   Note: For more information about these prerequisites, see "Prerequisites for Windows Physical Server Attachment" in this chapter.
7. (Optional) Confirm that the prerequisites for using CA Server Automation with CA Process Automation have been met. You can use CA Process Automation connectors to provision, discover, resize, move, and deprovision storage.

   Note: For more information about these prerequisites, see the section "IT PAM Prerequisites" in this Guide.

Configure DataFabric Managers

You can use the CA Server Automation user interface to configure DataFabric Managers.

To configure a DataFabric Manager

1. Click Administration, Configuration.
2. In the Network/Storage section, click NetApp DataFabric Manager .
3. Click + (Add).

   The Network Automation Server dialog appears.
4. Provide the following information:

   Server Name

   Specifies the name of the NetApp DataFabric Manager.

   User Name

   Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

   Note: The credentials specified during configuration are only used to verify whether the NetApp DataFabric Manager exists. The CA Server Automation login credentials are the credentials used to connect to the NetApp DataFabric Manager for storage provisioning. This user must have privileges to schedule and provision storage on the storage system (NetApp Power Users and Administrators have these privileges).

   Password

   Specifies the password of the user authorized to connect to the NetApp DataFabric Manager.

   Port

   Specifies the port number used by the NetApp DataFabric Manager.

   Default: 8088

   Protocol

   Specifies the protocol used to access the NetApp DataFabric Manager.

   Default: HTTP

   To verify the connection status, click Validate. To complete the configuration, click OK.

   The DataFabric Manager is configured.

Note: You can also provide this information during installation or using the dpmutil -set -netapp CLI command. For more information about using this command, see the *Reference Guide*.

Prerequisites for VMware vCenter ESX Servers

Verify that the following prerequisites are met before provisioning storage for VMware vCenter ESX servers.

- Install the required adapters

  Depending on the type of storage available, your VMware vCenter ESX Server requires the following adapters for connectivity to specific storage devices or networks:

  - ISCSI Adapter for an ESX

    The ESX/ESXi software iSCSI initiator built into VMkernel communicates with the network adapter through the network stack to facilitate the iSCSI connection. Use the vSphere Client Configuration tab and click Storage Adapters in the Hardware panel to configure the iSCSI adapter.

  - Fiber Channel HBA

    Use the default configuration settings for Fiber Channel HBA.

- Configure Discovery Addresses for iSCSI Initiators

  Set up target discovery addresses so that the iSCSI initiator can determine which storage resource on the network is available for access.

Note: For more information about installing and configuring the required adapters, see the VMware website.

Prerequisites for Windows Physical Server Attachment

If the following prerequisites are met, the Storage Provisioning Manager supports attachment on Windows physical servers.

- Install the NFS Client on the remote Windows host for NFS storage attachment.

  For Windows 2008, the NFS client is likely already installed; navigate to Control Panel, Server Manager, Roles, File Services, Add Role Services, Services for Network File System for verification.

- Install Microsoft iSCSI initiator on the remote Windows host for iSCSI storage attachment. For Windows 2008, the iSCSI initiator is likely already installed. For verification, look for iSCSI initiator in the Control Panel under Administrative Tools.

  For Windows 2003, you can download the iSCSI initiator from the Microsoft website under downloads.

- Configure the primary DNS suffix for the fully qualified domain name to provision and attach new storage on the server where CA Server Automation is installed.

  Note: For more information about configuring the primary DNS suffix, see the Microsoft website.

- Verify that the login credentials you use to connect to the Windows host for attachment have access to WMI.

  For Windows 2008, navigate to Control Panel, Server Manager, Configuration, WMI Control, Properties, Security tab, Security to verify that the login is listed:

- (64-bit Windows Server 2003 only) Install Microsoft hotfix 942589 if CA Server Automation is installed on a 32-bit host.

Note: For NFS and CIFS, logout and log back in to see the new drive after storage provision and attachment succeeds. For CIFS attachment, an error occurs if the storage is already mapped and you attach to the same storage with the same CIFS credentials.

Configure CA Network Automation

CA Network Automation lets you discover and manage Cisco routers and switches running the Cisco Internetwork Operating System (Cisco IOS) within an enterprise location. You can also use it to add, delete, and update discovered network devices. CA Network Automation provides scripts to perform additional changes to the network devices it manages.

In the folder *install-directory*\nma_scripts, CA Server Automation provides the following additional scripts for use with CA Network Automation:

- Create VLANs
- Create VLANs and subnets
- Delete VLANs

To import these scripts into the NetMRI system, navigate to the Configuration Management, Job Management page, and click Import.

Configure the CA Network Automation Server

To run Network Automation scripts from CA Server Automation, configure one CA Network Automation server for use with CA Server Automation.

Follow these steps:

1. Open the CA Server Automation user interface.
2. Click Administration, Configuration.
3. In the Network/Storage section, click Network Automation Server.
4. Click + (Add) and provide the following information:

   Server Name

   Specifies the name or IP address of the CA Network Automation server.

   User Name

   Specifies the user name for the CA Network Automation server.

   Password

   Specifies the password for the CA Network Automation server user.

   Protocol

   Specifies the protocol used to access the CA Network Automation server.

   Default: HTTP

   Port

   Specifies the port number used by the CA Network Automation server.
5. Click OK.

   The CA Network Automation server is configured.

Automating Processes with CA Process Automation

CA Process Automation automates routine administration tasks, improves operations efficiency and incident response handling, and helps ensure best practice and regulatory controls compliance. CA Process Automation can automate and manage many processes, including the following:

- Applications monitoring and restart
- Disaster recovery
- Virtual Infrastructure Management
- IT Infrastructure Library (ITIL) compliance
- Security
- Discovery
- Change detection
- Provisioning
- Performance monitoring
- Storage provisioning

The CA Process Automation integration with CA Server Automation enhances rules and actions handling by providing a graphical user interface that lets systems administrators configure and manage processes that CA Server Automation activates. CA Process Automation uses these processes to run operational processes automatically. CA Process Automation also supports client applications that let operators and other users schedule, start, and monitor automated processes.

A typical usage scenario follows:

- The administrator configures a rule that requires provisioning a single or multiple virtual machines on a specific date or when a particular metric is reached on a server.
- This rule activates the CA Server Automation interface to CA Process Automation and that triggers a process that has already been configured on that server.
- When the process terminates, an event is sent to CA Server Automation and a Service Desk ticket is created, if available.

CA Process Automation Prerequisites

Verify that the following requirements have been met before using CA Process Automation:

- The latest releases of CA Server Automation, CA Process Automation, CA EEM, and the public version of JRE are installed.
- The CA Process Automation Server is configured.
- CA Process Automation is configured for single sign-on (see page 90).
- Access to CA Process Automation web services is enabled.

Note: If you skipped installation of a component, you can configure it later using the dpmutil command-line utility or using the Administration, Configuration page. For more information about dpmutil, see the *Reference Guide*.

Configure CA Process Automation for Single Sign-On

Run the CA EEM safex utility to generate a certification file and specify a CA Server Automation server for single sign-on for that CA Process Automation server. To use single sign-on, CA Process Automation must be configured with CA EEM and the safex tool must be run against the ITPAM_eem.xml file.

To configure CA Process Automation and CA EEM for single sign-on

1. Locate the ITPAM_eem.xml file on DVD2 of the CA Process Automation installation media and copy the file to a directory on the C: drive.
2. (Optional) Edit the certificate name or password or change other properties such as application name, group name, and user name. The name and password entries follow:

   ```
   Register certfile="itpamcert.p12" password="itpamcertpass"
   ```

3. Locate the safex.exe tool on the CA EEM installation media and issue the following command:

   ```
   CA_EEM_Installation_Path/safex.exe –h hostname -u EiamAdmin –p eiamadminpass -f ITPAM_EEM.xml
   ```

   Note: The default directory for the safex tool is Program Files/CA/SharedComponents/iTechnology/ and the tool is automatically copied when you install CA EEM.

   The safex.exe tool runs and creates the certificate file, itpamcert.p12 in the default directory. The tool also creates the groups, ITPAMAdmins and ITPAMUsers and the users itpamadmin and itpamuser.

4. Log in to the CA EEM UI (see page 28) on the CA Process Automation server.

5. Verify that the groups and users created by the safex tool are available in Manage Identities, Groups in the CA EEM UI.
6. Reset the passwords of the users to the desired passwords.
7. Select EEM from the drop-down list in the CA Process Automation installation and complete the following fields:

   EEM Server

   Identifies the host name of the CA EEM server.

   Example: itpamserver.itpam.ca.local

   EEM Application Name

   Identifies the name of the CA EEM application instance.

   Example: ITPAM

   EEM Certificate File

   Identifies the full path to the certification file.

   Example: Program Files/CA/SharedComponents/iTechnology/itpamcert.p12

   EEM Certificate Password

   Identifies the password for the certificate file.

   Example: itpamcertpass

   The CA EEM security settings are configured.

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 Embedded Entitlements Manager Log In window appears.

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

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

   The User Name field is populated with EiamAdmin.

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

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

Access the CA Process Automation User Interface

The CA Process Automation Client is the administrative tool for CA Process Automation. Access the CA Process Automation Client user interface to design, deploy, monitor, control, and audit your IT processes. The CA Process Automation documentation is also located on the Home Page after you log in. The Start menu shortcut is only available on the CA Process Automation server. Users accessing the interface from a separate server must enter the URL in a web browser. You can also access the CA Process Automation Client user interface from the CA Server Automation UI.

To access the CA Process Automation user interface

1. Select Start, Programs, CA, CA Process Automation Domain, Start CA Process Automation Client from the CA Process Automation server.

   The CA Process Automation page opens at the following URL:

   `https://servername:port/itpam`

   servername

   Identifies the name of the server where the CA Process Automation user Interface is installed.

   port

   Specifies the server listening port.

   Default: 8080

2. Enter your administrator login credentials and click Log In.

   The CA Process Automation home page appears.

3. Click CA Process Automation Client in the upper right corner of the page.

   The CA Process Automation Domain Browser opens. Use this interface to configure your processes.

Configure a CA Process Automation Process

CA Process Automation processes provide you with a visual representation of your actions. You can see exactly where you are in the process and you can view multiple instances of actions. Before you configure a CA Process Automation process, create an action (see page 338) and a rule first. The rule is mapped to and triggers the process. Create rules and actions and configure the processes from the CA Server Automation UI.

To configure a CA Process Automation process

1. Click Resources.

   The Resources page appears.

2. Open the Explore pane.

   Available groups, services, and systems appear.

3. Select the Data Center node.

4. Click Policy, then click Actions.

   The Actions page appears.

5. Click + (New) on the toolbar.

   The Action Definition: New page appears.

6. Enter a meaningful name for the action in the Action Name text box and select the Run CA Process Automation Process action type from the drop-down list.

   The Details section appears.

7. Select *one* of the following settings in the Action Start drop-down list:

   No Delay

   Specifies that the same action can be rerun immediately when a rule using that action is triggered again.

   Delay For

   Specifies the amount of time that must elapse before the same action can be rerun when a rule using that action is triggered again.

   Limits: seconds

   Note: The Action Start setting has no effect when a scheduled job runs the action.

8. Select one of the following settings in the Action Completion drop-down list:

   No Wait

   Specifies no waiting period 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 for the action to complete before running succeeding actions in an action sequence.

Wait Indefinitely

Specifies to wait indefinitely for the action to complete before running succeeding actions in an action sequence.

Note: The Action Completion drop-down list appears only for long running actions.

9. Select a form from the drop-down list. Forms let you create an interface so that users can launch a process and make appropriate inputs to that process at startup.

   Required input fields depend on the form you selected.

10. Complete all input fields.

    Note: Connect Parameters are configured in CA Process Automation. The CA Server Automation URL identifies the Service Controller, which is displayed at the bottom of the Administration page, Configuration list.

11. (Optional) Click Open process in CA Process Automation Client to log in to CA Process Automation and view the process definitions.

12. Select the Help Desk Approval check box if the ticket requires approval by a third party.

    Note: Configure CA Service Desk Manager properly to use this option.

    The Ticket Types and Templates fields become enabled.

13. Select the Auto close ticket on approval check box if you want to close the ticket automatically after it is approved.

14. Select a ticket type from the Ticket Types drop-down list. The following types are valid options, but are dependent on your configuration:

    - Incident
    - Problem
    - Request

    The Templates drop-down list is updated with the templates associated with the ticket type you selected.

15. Select a template from the Templates drop-down list.

    The fields are populated with predetermined values depending on the ticket model you are using.

16. Select Save from the Actions drop-down list.

    A confirmation message notifies you that your save was successful.

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

More information:

Action Types (see page 340)

CA Process Automation Use Cases

Use cases are internal names assigned to the processes that perform automation steps based on CA Process Automation operators, methods, and custom operators (CA Server Automation connectors).

The CA Server Automation installation media provides use cases for the CA Process Automation integration. The names of the use case processes, the CA Server Automation connectors called by the use cases, and additional comments are described.

LoginInfoProc Process

CA Server Automation Connectors Called

- Query Service Controller
- Create Event
- Create Ticket

Comments

This use case creates the indications and/or help desk tickets required by the other use cases that call this use case.

This process takes the following connect parameters:

T_SVC_URL__

T_VC_EVENT_COMPONENT__

T_VC_EVENT_TIMESTAMP__

T_VC_EVENT_MESSAGE__

T_TICKET_TYPE_OF_REQUEST__

T_TICKET_ENTITY__

T_TICKET_DESCRIPTION__

T_TICKET_SUMMARY__

T_TICKET_USER__

ConfigurationAudit Process

CA Server Automation Connectors Called

- Query Service Controller
- Get All Cohesion Snapshots
- Create Event
- Create Ticket
- Run Discovery Profile
- Get Current Activity
- Performance Change Detection

Comments

This use case detects whether a particular server is in compliance with its baseline standard.

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_SYSTEM_NAME__

T_SNAPSHOT_TYPE__

CyberMonday Process

CA Server Automation Connectors Called

- Query Service Controller
- Get Service List
- Run Process
- Get Service Machine List
- Power On VC System
- Get Machine Status VC
- Remove Machine From Service
- Add Machine to Service

Comments

This use case brings online a VM system from a source CA Server Automation service to a target CA Server Automation service.

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_SVC_STANDBY_SERVICE__

T_SVC_SERVICE_TO__

T_SVC_VC_SERVER__

T_SVC_DATACENTER_NAME__

T_SVC_VM_MACHINE__

CyberMondayPowerDown Process

CA Server Automation Connectors Called

- Query Service Controller
- Get Service List
- Run Process
- Get Service Machine List
- Remove Machine From Service
- Add Machine to Service
- Get Machine Status
- Create Event

Comments

This use case moves a VM system from a CA Server Automation service back to the service that lent the system during the CyberMonday use case. After the VM system has been returned, the system is powered down.

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_SVC_STANDBY_SERVICE__

T_SVC_DATACENTER_NAME__

T_SVC_VM_MACHINE__

T_SVC_SERVICE_TO__

T_SVC_VC_SERVER__

AlmostGoldByService Process

CA Server Automation Connectors Called

- Query Service Controller
- Get Service Machine List
- Almost Gold Black And White -- In detached mode

Comments

This use case lists the systems in a CA Server Automation service. After this list is created, the AlmostGoldBlackAndWhite use case is called to determine if all the required components are present.

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_SVC_SERVICE__

T_SVC_COREFILE__

T_SVC_UNACCEP_FILE__

AlmostGoldBlackAndWhite Process

CA Server Automation Connectors Called

- Query Service Controller
- Get All Components
- Start Script
- Create Event
- Create Ticket

Comments

This use case determines if a particular server contains all required components and has no invalid applications.

For this use case, you must create the following files:

- newcore2.lst
- unacceptablecomponents.lst

newcore2.lst

```
#-------------------------------------------------------------------------
#            Black-list Components (BL)
#
#      These components MUST exist on the target server
#       and are version specific
#  CC -- represents the core components
#  component name
#  version number
#-------------------------------------------------------------------------
CC:Windows:5.2
CC:Java Web Application:2.4
CC:Java Web Application:2.4
CC:Java Web Application:2.3
```

unacceptablecomponents.list

```
#-------------------------------------------------------------------------
#            Black-list Components (BL)
#
#      These components MUST exist on the target server
#       and are version specific
#  BL -- represents the black list components
#  component name
#  version number
#-------------------------------------------------------------------------
BL:VMware Server:2.*
#BL:CA Threat Manager:*.*
BL:winamp:*
BL:JRE 1.2 or 1.3 (Windows):1.3.1_02
#BL:Log4J:1.2.8
```

Server Automation Disk Usage Process

CA Server Automation Connectors Called

- Query Service Controller
- Start Script (check disk space)
- Create Event
- Create Ticket

Comments

This use case determines the amount of available free disk space on the CA Server Automation server.

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_SVC_TOUCHPOINT__

T_SVC_LAG__

T_SVC_INTERVAL__

T_MIN_DISK_SIZE__

T_MIN_PERCENT_DISK_SIZE__

T_SEC_DELAY__

Server Automation SQL Pct Free Process

CA Server Automation Connectors Called

- Query Service Controller
- Start Script (Check DB Space Free)
- Create Event

Comments

This use case verifies that the following databases have enough free space:

- AOM2
- DPM
- tempdb

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_SEC_DELAY__

T_DB_SERVER__

T_DB_USER__

T_DB_PASSWORD__

T_DB_PORT__

T_DB_MIN_PCT_FREE__

SSRM VM Reservation Process

CA Server Automation Connectors Called

- Query Service Controller
- SSRM Get Resrc Pool
- SSRM Get Data Software
- SSRM Get Data Templates
- SSRM Get System Requirements
- SSRM Check System Availability
- SSRM Create Reservation
- SSRM Get Res Status

Comments

This use case creates a CA Process Automation reservation with a VM-based template.

Users access this use case with the Reservation Manager start request form, which gathers all the available Reservation Manager resource pools, software groups, and VM templates.

After this information is gathered, users can see it in a form from which they can select the desired VM template and the date and time to use for the reservation.

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_SSRM_USERNAME__

T_SSRM_ORGUNIT__

Storage Provision VM Image Process

CA Server Automation Connectors Called

- Query Service Controller
- Storage Provision NFS
- Storage Create NAS Datastore
- Provision VC Image
- VC Job Status
- Publish Indication

Comments

This use case adds additional storage to a vCenter server, and provisions a new VM with the additional storage assigned. Storage is provisioned on NetApp using NFS export. A VMware datastore is then created using the new NFS storage. A VM is then created on the new storage.

The use case can start from the Start Form 'Storage Provision VMware Image'.

This process takes the following connect parameters:

T_SVC_URL__

T_SVC_USER__

T_SVC_PASSWORD__

T_STORAGE_VCENTER__

T_STORAGE_ESX__

T_STORAGE_DATASET__

T_STORAGE_DATASET_SIZE__

T_STORAGE_POLICY__

T_STORAGE_RES_POOL__

T_DATACENTER_NAME__ T_COMPUTE_RESOURCE_NAME__

T_STORAGE_TARGET_LOC__

T_CLONE_NAME__

T_TARGET_USERNAME__

T_TARGET_PASSWORD__

T_VM_NAME_CLONE_FROM__ (mutually exclusive with T_VC_TEMPLATE__)

T_VC_TEMPLATE__ (mutually exclusive with T_VM_NAME_CLONE_FROM__)

T_VC_SPECIFICATION__ (Optional)

T_NETAPP_SERVER_NAME__ (Optional)

T_NETAPP_SERVER_USERNAME__ (Optional)

T_NETAPP_SERVER_PASSWORD__ (Optional)

T_NETAPP_SERVER_PROTOCOL__ (Optional)

T_NETAPP_SERVER_PORT__ (Optional)

T_VC_DATASTORE_PAUSE__

Import Connectors into CA Process Automation

If you skip the installation configuration step, CA Server Automation connectors are stored in an XML file that you can import into CA Process Automation manually.

To import CA Server Automation connectors into CA Process Automation manually

1. Launch the CA Process Automation client.
2. Highlight the root folder in the CA Process Automation management client.
3. Right-click Import.
4. Browse to the file CA_Server_Automation.xml on the CA Server Automation server in the following location:

   `\Program Files\CA\ServerAutomation\CA_Server_Automation.xml`
5. Select all the available options In the Import Object dialog. The options are:

   Set Imported Version As Current

   Make Imported Custom Operators/Sensors Available

The connectors become current and available to users.

To import CA Server Automation connectors into CA Process Automation using the ITPAMexport utility

1. Click Start, CA, CA Server Automation Command Prompt.

   Command Prompt window opens.
2. Run the following command:

   `\Program Files\CA\ServerAutomation\bin\itpamexport`

   The command accepts the following arguments:

   `<ITPAM Server> <user> <password> CA Portal <protocol> <import file>`

   Example:

   `C:\Program Files\CA\ServerAutomation\bin\itpamexport hostname.ca.com pamadmin <password> 8080 http CA_Server_Automation.xml`

In the CA Process Automation client, the latest connectors are located here:

`/Custom Operators/CA Server Automation`

Connector Syntax

Use the following syntax for the CA Server Automation connectors:

- All fields expect an expression (variable) as a parameter value, unless otherwise specified.
- Enclose the examples provided in double quotation marks.
- When providing parameter values in Request Forms, the double quotation marks around parameters are not required. For more information, see the CA Process Automation documentation.
- Unless otherwise specified, the connect parameters for all components are:

  CA Server Automation Manager URL

  Specifies the URL located under the Administration, Configuration tabs in the UI.

  Example: "https://localhost:443/dpm"

  Login User

  Specifies the CA Server Automation user with corresponding permissions.

  Example: "saadmin"

  Login Password

  Specifies the CA Server Automation password associated with the login user.

  Example: "adminpassword"

  Note: In a distributed environment, the first connector in each process should be *Query Service Controller.*

  Each subsequent controller should contain an expression (variable) that the Query Service Controller connector populates with the correct URL.

  The URL for the Query Service Controller connector is located at the bottom of the specified location.

  Example: https://localhost:443/aip (Do not include anything after *aip*)

CA Process Automation Connectors

A complete list of CA Server Automation connectors and descriptions follows organized by their CA Process Automation categorization.

Common Connectors

A complete list of Common connectors follows.

More information:

Add Machine to Service

Adds new machines to the CA Server Automation service.

Component

resourcemgr

Service Machine Name(s)

Specifies the fully qualified name of the target server as displayed in CA Server Automation.

Example: "dev-test.company.com"

Service Name

Specifies the target service name associated with the specified machine name.

Example: "Enterprise\\\\Data Center\\\\TargetSvc"

Component Status SC

Returns status information for a component.

Component

sc

SC Component ID

Specifies the short name ID of the component.

Example: "sda"

SC Host Name

Specifies the server where the component is installed.

Example: “localhost.anycompany.com”

Create Service Group

Creates a service or group in a CA Server Automation server.

Component

resourcemgr

New Service Name

Specifies a name for the new service to create in CA Server Automation. For a nested service use ‘\\\\’ between levels.

Example: “Enterprise\\\\Data Center\\\\Production”

Service Machine List

Specifies the server name to include in the new service.

Example: “localhost.anycompany.com”

Service Lower Threshold

Specifies the lower utilization threshold for the service. This value must be less than the upper threshold.

Example: “20”

Service Upper Threshold

Specifies the upper utilization threshold for the service. This value must be greater than the lower threshold.

Example: “80”

Service Lag

Specifies the lag limits for the service. This value must be 1 second or greater.

Default: “1”

Service Priority

Specifies the priority for the service. This value must be 1 or greater.

Default: “1”

Service ID

(Optional) Specifies a service ID for the service.

Delete Job

Deletes a machine job.

Component

sch

Job ID

Specifies the unique ID of the CA Server Automation job.

Example: “12039384”

Job Username

Specifies the CA Server Automation authorized user.

Example: "username1"

Job Username's Password

Specifies the password for the job user name.

Example: "passw0rd1"

Discovery System

Discovers a system based on the name provided.

Component

AOM

System Name

Specifies the fully qualified name for the target system.

Example: “localhost.mycompany.com”

Discovery Correlation ID

Specifies a user-defined correlation ID to use to identify the discovery job.

Example: “PA-2010-05-03-12-00-00”

Server Automation Service Path

(Optional) Specifies the service path for the CA Server Automation service to add the discovered system to.

Example: “Enterprise\\\\Data Center\\\\New Service”

CCA Service Path

(Optional) Specifies the CA Configuration Automation service path.

Example:
"https:\/\/localhost\/aip\/AOM\/root\/cimv2:CIM_Service.CreationClassName=\"CIM_Service\",Name=\"CCA_Servavc-cca.anycompany.com.8080\",SystemCreationClassName=\"CA_ComputerSystem\",SystemName=\"cca.anycompany.com\""

Get Job

Gets the status of a given system job.

Component

sch

Job ID

Specifies the unique ID of the CA Server Automation job. Use a UUID value.

Example: “00004-23233-232132-340323”

Get Job List

Gets a job list for a given system.

Component

sch

Job Type

(Optional) Specifies a job type filter.

Job Name

(Optional) Specifies a job name filter.

Service ID

(Optional) Specifies a service ID filter.

Job Description

(Optional) Specifies a job description filter.

Job Username

(Optional) Specifies a user assignment filter.

Job Last Return Code

(Optional) Specifies a last return code for the jobs executed.

Job Last Run Date Lower

(Optional) Specifies a lower bound for the last run date.

Example: "01-01-2008"

Job Last Run Date Upper

(Optional) Specifies an upper bound for the last run date.

Example: "12-01-2010"

Job Next Run Date Lower

(Optional) Specifies a lower bound for the next run date.

Example: "01-01-2008"

Job Next Run Date Upper

(Optional) Specifies an upper bound for the next run date.

Example: "12-01-2010"

Get Service List

Gets a listing of the available components for a service.

Component

resourcemgr

Service Name

Specifies the root level to obtain the services from.

Default: "Data Center"

Example: "\\\\Enterprise\\\\Data Center"

Recursive

(Optional) Identifies whether to activate a recursive search for other services.

Example: "1"

Get Service Machine List

Gets a server list for a specified service name.

Component

resourcemgr

Service Name

Specifies the name of the service.

Example: "Enterprise\\\\Production" when targeting a root service.

For a nested service use \\\\ between levels.

Example: "Enterprise\\\\Data Center\\\\Accounting"

Get Version

Gets the version number of the imaging option.

Component

img

No parameters are required.

Imaging Job Status

Returns the status of jobs for AmazonEC2, JumpStart, LPAR, HyperV, RSI, Zones, Software Delivery, and VMware.

Component

img

Type of Imaging Job

Specifies the type of job to return the status of. Available values are:

AmazonEC2

JumpStart

IBMLPAR

MSHyperV

RSI

SolarisZones

SoftwareDelivery

VMware

Image Job ID

Specifies the job ID provided by the CA Process Automation connectors.

Example: "ef044f5b-7fdf-11a0-abcd-1150568605fc”

Publish Indication

Provides the different event indication of a process.

Component

AOM

Indication Description

Describes the event taking place.

Indication Type

Indicates the type of event taking place.

Indication Message ID

Specifies the ID for the message.

Indication Message Arguments 1-4

(Optional) Specifies any message arguments.

Indication Severity

(Optional) Indicates the severity of the message.

Example: "Information"

Indication Provider Name

(Optional) Specifies the generator of the message.

Example: "IT PAM"

Indication Source Description

(Optional) Specifies a description for the generated indication.

Indication Error Code

(Optional) Specifies the error code associated with the generated indication.

Example: "Informational"

Query Service Controller

Queries the status of a service.

Component

sc

No parameters other than the connection parameters are required.

Remove Machine From Service

Removes a system from the CA Server Automation service.

Component

resourcemgr

Machine Name

Specifies the fully qualified name for the target host.

Example: "localhost.anycompany.com"

Service Name

Specifies the service name where the target host is located.

Example: "Enterprise\\\\Data Center\\\\YourService".

Run Job

Runs a system job.

Component

sched

Job ID

Specifies the unique ID of the CA Server Automation job.

Example: 'SAm-1'

AmazonEC2 Connectors

A complete list of Platform Support AmazonEC2 connectors follows.

More information:

AMI Run Instance (see page 114)
AMI Terminate Instance (see page 114)
Provision AMI Image (see page 115)

AMI Run Instance

Performs an AMI instance start.

Component

ec2

Instance ID

Specifies the provisioned AMI image instance identifier that was returned when provisioning was granted.

Example: 'i-a9b9c9d0'

AMI Terminate Instance

Performs a termination of an AMI instance.

Component

ec2

Instance ID

Specifies the provisioned AMI image instance identifier that was returned when provisioning was granted.

Example: 'i-a9b9c9d0'

Provision AMI Image

Creates an Amazon EC2 AMI image.

Component

img

Image ID

Specifies an ID for the image.

Example: “ami-10c66688” (Instances Tab)

Instance Type

Specifies the instance type assigned to the image ID.

Example: “c1.medium” (Reserved Instance Tab)

Minimum Count

Optional parameter.

Maximum Count

Optional parameter.

Group Set

Specifies the Amazon AMI default group set to define this image ID.

Key Pair Name

Specifies the security name for the device that allows access to the web server.

Example: “SAtest-pair”, “test”

Availability Zone

Specifies the zone name where provisioning occurs.

Example: “us-east-1b”

OS Type

Specifies the type of operating system to provision.

User Data

Optional parameter.

AppLogic Connectors

A complete list of AppLogic connectors follows.

More information:

AppLogic App Parameter Details

Returns the parameters for an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Application Parameters

Returns the configuration parameters for an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Copy Application

Creates a copy of an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic New Application

Specifies the name for the new application.

Example: "new sample application"

AppLogic Delete Application

Deletes an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Job Info

Returns the status of an AppLogic job.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Job

Specifies the ID of the job.

Example: "AAAA-BBBB-CCCC-DDDD-EEEE"

AppLogic List App Templates

Returns a list of the available AppLogic templates.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic List Applications

Returns a list of the available AppLogic applications.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic List Grids

Returns a list of the available AppLogic grids.

Component

applogicws

No parameters are required.

AppLogic List Used IPs

Returns a list of the AppLogic grid public IP addresses currently in use.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Migrate Application

Creates a copy of an AppLogic application in another AppLogic grid.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic New Grid

Specifies the name of the grid to copy the application to.

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Modify Application

Changes the value of a parameter for an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Application Parameter

Specifies the name of the application parameter.

Example: "mem"

AppLogic Application Parameter Value

Specifies the new value to apply to the application parameter.

Example: "512"

AppLogic Provision App Unix

Provisions an AppLogic Unix/Linux application to a grid based on an application template with specified resource sizing and boundary values.

Component

img

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application Template

Specifies the name of the application template.

Example: "sample application template"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic App Hostname

Specifies the hostname to assign to the application.

Example: "localhost.ca.com"

AppLogic App Primary IP

Specifies AppLogic public IP to assign to the application.

Example: "127.0.0.1"

AppLogic App Secondary IP

(Optional) Specifies a secondary IP to assign to the application.

Example: "127.0.0.2"

AppLogic App Gateway

Specifies the grid public gateway to assign to the application.

Example: "1.2.2.1"

AppLogic App DNS Server

Specifies the DNS server to assign to the application.

Example: "1.2.2.2"

AppLogic App DNS Server 2-3

(Optional) Specifies additional DNS servers to assign to the application.

Example: "1.2.2.3"

AppLogic App Netmask

Specifies the grid public netmask to assign to the application.

Example: "255.255.255.255"

AppLogic App Root Password

Specifies the root password for the application.

Example: "rootpassword"

AppLogic App User

(Optional) Specifies the default user for the application.

Example: "demouser"

AppLogic App User Password

(Optional) Specifies the password for the default user.

Example: "demopassword"

AppLogic Provision App Windows

Provisions an AppLogic Windows application to a grid based on an application template with specified resource sizing and boundary values.

Component

img

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application Template

Specifies the name of the application template.

Example: "sample application template"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic App Hostname

Specifies the hostname to assign to the application.

Example: "localhost.ca.com"

AppLogic App Primary IP

Specifies AppLogic public IP to assign to the application.

Example: "127.0.0.1"

AppLogic App Secondary IP

(Optional) Specifies a secondary IP to assign to the application.

Example: "127.0.0.2"

AppLogic App Gateway

Specifies the grid public gateway to assign to the application.

Example: "1.2.2.1"

AppLogic App DNS Server

Specifies the DNS server to assign to the application.

Example: "1.2.2.2"

AppLogic App DNS Server 2-3

(Optional) Specifies additional DNS servers to assign to the application.

Example: "1.2.2.3"

AppLogic App Netmask

Specifies the grid public netmask to assign to the application.

Example: "255.255.255.255"

AppLogic App Admin Password

Specifies the admin password for the application.

Example: "rootpassword"

AppLogic App User

(Optional) Specifies the default user for the application.

Example: "demouser"

AppLogic App User Password

(Optional) Specifies the password for the default user.

Example: "demopassword"

AppLogic Provision Application

Provisions an AppLogic application to a grid based on an application template with specified resource sizing and boundary values.

Component

img

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application Template

Specifies the name of the application template.

Example: "sample application template"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic App Config Parameters

Specifies a comma-separated list of configuration parameters required to provision the application.

Example: "hostname=anyhost,usr_ip=127.0.0.1, admin_ip=127.1.0.2"

AppLogic Rename Application

Changes the name of an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic New Application

Specifies the new name for the application.

Example: "new sample application"

AppLogic Restart Application

Restarts an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Start Application

Starts an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Stop Application

Stops an AppLogic application.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Application

Specifies the name of the application.

Example: "sample application"

AppLogic Template Parameters

Returns the configuration parameters for an AppLogic template.

Component

applogicws

AppLogic Grid

Specifies the name of the grid.

Example: "user-grid"

AppLogic Template

Specifies the name of the application template.

Example: "sample application template"

Configuration Management Connectors

A complete list of Configuration Management connectors follows.

More information:

Create Config Mgmt Snapshot (see page 126)
Get All Components (see page 127)
Get All Config Mgmt Snapshots (see page 127)
Get Current Activity (see page 128)
Perform Change Detection (see page 128)
Perform Compare Systems (see page 129)
Run Discovery Profile (see page 130)

Create Config Mgmt Snapshot

Creates a machine snapshot in the CA Configuration Automation system.

Component

ccm

Snapshot Name

Specifies a name to identify the snapshot.

Example: “firstsnapshotname”

Target System Name

Specifies the fully qualified name of the target system.

Example: “localhost.anycompany.com”

Is Gold Standard?

(Optional) Indicates whether to set a gold standard for the snapshot.

Is Base Line?

(Optional) Indicates whether to set a base line for the snapshot.

Is Silver Standard?

(Optional) Indicates whether to set a silver standard for the snapshot.

Is Bronze Standard?

(Optional) Indicates whether to set a bronze standard for the snapshot.

Get All Components

Gets all available components for a specific system.

Component

ccm

Target Machine

Specifies the fully qualified name of the server.

Example: "localhost.anycompany.com"

Get All Config Mgmt Snapshots

Gets all available CA Configuration Automation snapshots for a specific machine.

Component

ccm

Target System Name

Specifies the fully qualified name of the target system.

Example: "localhost.anycompany.com"

Target Snapshot Type

Specifies the snapshot type to retrieve. Available values are:

1 – Gold Standard

2 – Base Line

3 – Silver Standard

4 – Bronze Standard

Example: "2"

Get Current Activity

Gets the current activity of a specific system.

Component

ccm

Host Name

Specifies the fully qualified name of the target system.

Example: “localhost.anycompany.com”

Perform Change Detection

Performs a change detection of a specific system. After you choose the parameters, the system prompts for the CA Server Automation URL, user name, and password. Always surround the URL with quotation marks ("); user name and password only if they contain special characters like # or &.

Component

ccm

Operation Type

Specifies the type of change detection to conduct. Available options are:

- Most Recent Snapshot to Current
- Baseline to Current
- Gold Standard to Current
- Silver Standard to Current
- Bronze Standard to Current

Difference Type

Specifies the type of differences to detect. Available options are:

- Component Only
- All Differences

Example: "Component Only"

System Name

Specifies the fully qualified name for the target system.

Example: “localhost.mycompany.com”

Perform Compare Systems

Performs a system compare between a specific system and a stored snapshot.

Note: Depending on the comparison to perform, one or more of the parameters become optional.

Component

ccm

Source System Name

Specifies the fully qualified name for the system to serve as base reference for comparison.

Example: “localhostBase.anycompany.com”

Type of Difference

Specifies the type of differences to detect. Available options are:

1 – Component only

2 – All differences

Default: "2"

Target System Name

Specifies the fully qualified name of the target system.

Example: “localhost.anycompany.com”

Snapshot ID

Specifies the ID for the snapshot belonging to the source system.

Example: “5b6160ec-717a-4a90-b53c-ce97bf9bd83f”

Snapshot Type

Specifies the type of snapshot. Available options are:

0 – Compare source against the system gold snapshot.

1 – Compare source against the system baseline.

2 – Compare source against the system silver snapshot.

3 – Compare source against the system bronze snapshot.

Default: "1"

Target Snapshot ID

Specifies the snapshot ID that belongs to the target server.

Run Discovery Profile

Runs the discovery profile for a given system.

Component

ccm

System Name

Specifies the fully qualified name for the target system.

Example: “localhost.mycompany.com”

Synchronized Call

Indicates whether the call waits for discovery or only gets a job ID.

Example: “No”

HelpDesk Connectors

A complete list of HelpDesk connectors follows.

More information:

Create Ticket (see page 130)

Create Ticket

Generates a new help desk ticket.

Component

hd

HelpDesk Ticket Security Token

(Optional) Specifies the type of request.

Example: “issue” or “problem”

HelpDesk Ticket Entity

(Optional) Specifies a detailed message for the service desk description field.

Example: “Description from CA Process Automation”

HelpDesk Ticket Type of Request

Specifies the request type.

Example: "issue" or "problem"

HelpDesk Ticket Summary

Specifies a detailed message for the service desk summary field.

Example: "Details from CA Process Automation"

HelpDesk Ticket Template

(Optional) Specifies the template name from service desk to use.

HelpDesk Ticket Affected User

Specifies the user affected by the help desk issue.

Default: "itpam"

Example: "user_lastname, user_firstname"

LPAR Connectors

A complete list of Platform Support LPAR connectors follows.

More information:

LPAR Add LPAR CPU

Adds CPU units to an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Virtual CPU

Specifies the number of CPUs for the LPAR.

Example: "1"

LPAR CPU Adjustment

Specifies the adjustment value for the CPU processor assigned to the LPAR.

Example: "0.2"

LPAR CPU Adjustment Type

Specifies the type of adjustment for the LPAR CPU. Available values are "dynamic" and "all".

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Add LPAR Memory

Adds memory to an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Memory Adjustment

Specifies the adjustment value for the memory assigned to the LPAR in megabytes.

Example: "128"

LPAR Memory Adjustment Type

Specifies the type of adjustment for the LPAR memory. Available values are "dynamic" or "all".

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Attach iSCSI Target

Attaches an iSCSI target to an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

VIO Server Name

Specifies the IBM Virtual I/O Server name.

Example: "lpar_vio"

LPAR iSCSI Local Initiator ID

Specifies the ID of the LPAR iSCSI local initiator.

LPAR iSCSI Target Initiator ID

Specifies the ID of the LPAR iSCSI target initiator.

LPAR iSCSI Target Hostname

Specifies the LPAR iSCSI target hostname.

LPAR iSCSI Target Port Number

Specifies the LPAR iSCSI target port number.

LUN ID

Specifies the ID of the LUN.

Manage Targets File

Specifies the manage targets file.

LPAR iSCSI Target Password Type

(Optional) Specifies the LPAR iSCSI target password type.

LPAR Create Logical Part-IVM

Creates an empty LPAR partition with no OS installed.

Component

img

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

LPAR Min CPU Requested

Specifies the minimum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Max CPU Requested

Specifies the maximum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Num of CPU Desired

Specifies the ideal number of CPUs for this LPAR provisioning.

Example: "1"

LPAR CPU Shared Mode

Indicates whether to provision the CPU for the LPAR in shared mode.

Default: "true"

LPAR Min CPU Unit Requested

Specifies the minimum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Max CPU Unit Requested

Specifies the maximum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Desired Min CPU Unit

Specifies the ideal CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR CPU Capped

Indicates whether the CPU for the provisioned LPAR is capped.

Default: "false"

LPAR CPU Uncapped Weight

Specifies the uncapped weight for the provisioned LPAR.

Example: "128"

LPAR Minimum Memory

Specifies the minimum memory for the provisioned LPAR.

Example: "512"

LPAR Maximum Memory

Specifies the maximum memory for the provisioned LPAR.

Example: "1024"

LPAR Memory Desired

Specifies the ideal memory for the provisioned LPAR.

Example: "512"

LPAR Virtual Ethernet IEEE

Indicates whether the LPAR is virtual ethernet IEEE compliant.

Example: "false"

LPAR Virtual SCSI Remote LPAR

Specifies the remote LPAR associated with the virtual SCSI adapter for the provisioned LPAR.

Example: "lpar_hostname"

LPAR Virtual SCSI Adapter Is Client

Indicates whether the virtual SCSI adapter is a client for the provisioned LPAR.

Default: "true"

LPAR Virtual SCSI Adapter Is Required

Indicates whether the virtual SCSI adapter is required for the provisioned LPAR.

Default: "false"

LPAR Create Logical Partition

Creates an empty LPAR partition with no OS installed.

Component

img

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Profile

Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

LPAR Min CPU Requested

Specifies the minimum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Max CPU Requested

Specifies the maximum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Num of CPU Desired

Specifies the ideal number of CPUs for this LPAR provisioning.

Example: "1"

LPAR CPU Shared Mode

Indicates whether to provision the CPU for the LPAR in shared mode.

Default: "true"

LPAR Min CPU Unit Requested

Specifies the minimum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Max CPU Unit Requested

Specifies the maximum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Desired Min CPU Unit

Specifies the ideal CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR CPU Capped

Indicates whether the CPU for the provisioned LPAR is capped.

Default: "false"

LPAR CPU Uncapped Weight

Specifies the uncapped weight for the provisioned LPAR.

Example: "128"

LPAR Minimum Memory

Specifies the minimum memory for the provisioned LPAR.

Example: "512"

LPAR Maximum Memory

Specifies the maximum memory for the provisioned LPAR.

Example: "1024"

LPAR Memory Desired

Specifies the ideal memory for the provisioned LPAR.

Example: "512"

LPAR Virtual Ethernet IEEE

Indicates whether the LPAR is virtual ethernet IEEE compliant.

Example: "false"

LPAR Virtual SCSI Remote LPAR

Specifies the remote LPAR associated with the virtual SCSI adapter for the provisioned LPAR.

Example: "lpar_hostname"

LPAR Virtual SCSI Adapter Is Client

Indicates whether the virtual SCSI adapter is a client for the provisioned LPAR.

Default: "true"

LPAR Virtual SCSI Adapter Is Required

Indicates whether the virtual SCSI adapter is required for the provisioned LPAR.

Default: "false"

LPAR Create Logical Volume

Creates a logical volume for an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

VIO Server Name

Specifies the IBM Virtual I/O Server name.

Example: "lpar_vio"

LPAR Volume Group List

Specifies the LPAR volume group list.

LPAR Logical Volume Name

Specifies the name of the LPAR logical volume.

LPAR Logical Volume Size

Specifies the size of the LPAR logical volume.

LPAR Delete iSCSI Target

Deletes an iSCSI target for an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

VIO Server Name

Specifies the IBM Virtual I/O Server name.

Example: "lpar_vio"

LPAR Physical Volume Name

Specifies the name of the LPAR physical volume.

Manage Targets File

Specifies the manage targets file.

LPAR iSCSI Target Password Type

(Optional) Specifies the LPAR iSCSI target password type.

LPAR Delete Logical Volume

Deletes a logical volume for an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Logical Volume Name

Specifies the name of the LPAR logical volume.

LPAR Delete LPAR

Deletes an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Delete All Logical Volumes

Indicates whether to delete all logical volumes for the LPAR. Available values are "Yes" or "No".

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR List LPAR Profiles

Returns a list of available LPAR profiles.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR List NIM Images

Returns a list of available NIM images.

Component

img

NIM Master

Specifies the name of the NIM master host where the list of available NIM images is obtained.

Example: "hostnamea"

LPAR NIM Provision Ind Res

Provisions an LPAR partition and installs an operating system based on a NIM image. The selection is made from the individual resource.

Component

img

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Profile

Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

Provision LPAR

Indicates whether to provision an empty LPAR before installing the operating system. Select "false" if the LPAR exists and to install the operating system to it.

NIM Master

Specifies the name of the NIM master host where the list of available NIM images is obtained.

Example: "hostnamea"

NIM Machine Resource Name

Specifies the name of the NIM server used to install the operating system image to the provisioned LPAR.

Example: "machine_to_copy_from"

NIM Installation Type

Specifies the type of NIM installation to perform with the following options:

- rte - run-time environment
- mksysb - template based

Target Username

Specifies the user name for the target machine for operating system installation.

Example: "root"

Target Password

Specifies the password for the target user name.

Example: "password"

Mksysb Image

(Optional) Specifies the name of the image when the installation type is mksysb.

Example: "mksysb_image"

SPOT Resource

Specifies the image Shared Product Object Tree (SPOT) resource.

Example: "7100spot"

License Program Products

(Optional) Specifies licensed program products to deploy with the image when the installation type is rte.

Base Operating System

(Optional) Specifies the base operating system required when the installation type is rte.

Image DNS Resolve Config

(Optional) Specifies the DNS resolve configuration file to use with the image when the installation type is rte.

Example: "sample_dns.conf"

Image Data

(Optional) Specifies image data to use when the installation type is rte.

Example: "image_data"

First Boot Script

(Optional) Specifies the script to execute during the first boot when the installation type is rte.

Example: "first_boot.sh"

Post Installation Script 1-3

(Optional) Specifies post-installation scripts to execute with the image when the installation type is rte.

Example: "post_install.sh"

ITCM Scalability Server

(Optional) Specifies the ITCM/DSM scalability server to use to deploy the software delivery agent.

Example: "itcm_server"

Server Automation Template

(Optional) Specifies the server automation template to deploy to the LPAR.

Example: "software_package_template"

Agents Auto Deploy

(Optional) Indicates whether to deploy CA Configuration Automation, Performance, and Asset Management agents to the LPAR.

Example: "false"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

LPAR Min CPU Requested

Specifies the minimum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Max CPU Requested

Specifies the maximum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Num of CPU Desired

Specifies the ideal number of CPUs for this LPAR provisioning.

Example: "1"

LPAR CPU Shared Mode

Indicates whether to provision the CPU for the LPAR in shared mode.

Default: "true"

LPAR Min CPU Unit Requested

Specifies the minimum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Max CPU Unit Requested

Specifies the maximum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Desired Min CPU Unit

Specifies the ideal CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR CPU Capped

Indicates whether the CPU for the provisioned LPAR is capped.

Default: "false"

LPAR CPU Uncapped Weight

Specifies the uncapped weight for the provisioned LPAR.

Example: "128"

LPAR Minimum Memory

Specifies the minimum memory for the provisioned LPAR.

Example: "512"

LPAR Maximum Memory

Specifies the maximum memory for the provisioned LPAR.

Example: "1024"

LPAR Memory Desired

Specifies the ideal memory for the provisioned LPAR.

Example: "512"

LPAR Virtual Ethernet IEEE

Indicates whether the LPAR is virtual ethernet IEEE compliant.

Example: "false"

LPAR Virtual SCSI Remote LPAR

Specifies the remote LPAR associated with the virtual SCSI adapter for the provisioned LPAR.

Example: "lpar_hostname"

LPAR Virtual SCSI Adapter Is Client

Indicates whether the virtual SCSI adapter is a client for the provisioned LPAR.

Default: "true"

LPAR Virtual SCSI Adapter Is Required

Indicates whether the virtual SCSI adapter is required for the provisioned LPAR.

Default: "false"

LPAR NIM Provision Ind Res-IVM

Provisions an LPAR partition and installs an operating system based on a NIM image. The selection is made from the individual resource. The IVM (Integrated Virtual Manager) is used to provision the image.

Component

img

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

Provision LPAR

Indicates whether to provision an empty LPAR before installing the operating system. Select "false" if the LPAR exists and to install the operating system to it.

NIM Master

Specifies the name of the NIM master host where the list of available NIM images is obtained.

Example: "hostnamea"

NIM Machine Resource Name

Specifies the name of the NIM server used to install the operating system image to the provisioned LPAR.

Example: "machine_to_copy_from"

NIM Installation Type

Specifies the type of NIM installation to perform with the following options:

- rte - run-time environment
- mksysb - template based

Target Username

Specifies the user name for the target machine for operating system installation.

Example: "root"

Target Password

Specifies the password for the target user name.

Example: "password"

Mksysb Image

(Optional) Specifies the name of the image when the installation type is mksysb.

Example: "mksysb_image"

SPOT Resource

Specifies the image Shared Product Object Tree (SPOT) resource.

Example: "7100spot"

License Program Products

(Optional) Specifies licensed program products to deploy with the image when the installation type is rte.

Base Operating System

(Optional) Specifies the base operating system required when the installation type is rte.

Image DNS Resolve Config

(Optional) Specifies the DNS resolve configuration file to use with the image when the installation type is rte.

Example: "sample_dns.conf"

Image Data

(Optional) Specifies image data to use when the installation type is rte.

Example: "image_data"

First Boot Script

(Optional) Specifies the script to execute during the first boot when the installation type is rte.

Example: "first_boot.sh"

Post Installation Script 1-3

(Optional) Specifies post-installation scripts to execute with the image when the installation type is rte.

Example: "post_install.sh"

ITCM Scalability Server

(Optional) Specifies the ITCM/DSM scalability server to use to deploy the software delivery agent.

Example: "itcm_server"

Server Automation Template

(Optional) Specifies the server automation template to deploy to the LPAR.

Example: "software_package_template"

Agents Auto Deploy

(Optional) Indicates whether to deploy CA Configuration Automation, Performance, and Asset Management agents to the LPAR.

Example: "false"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

LPAR Min CPU Requested

Specifies the minimum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Max CPU Requested

Specifies the maximum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Num of CPU Desired

Specifies the ideal number of CPUs for this LPAR provisioning.

Example: "1"

LPAR CPU Shared Mode

Indicates whether to provision the CPU for the LPAR in shared mode.

Default: "true"

LPAR Min CPU Unit Requested

Specifies the minimum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Max CPU Unit Requested

Specifies the maximum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Desired Min CPU Unit

Specifies the ideal CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR CPU Capped

Indicates whether the CPU for the provisioned LPAR is capped.

Default: "false"

LPAR CPU Uncapped Weight

Specifies the uncapped weight for the provisioned LPAR.

Example: "128"

LPAR Minimum Memory

Specifies the minimum memory for the provisioned LPAR.

Example: "512"

LPAR Maximum Memory

Specifies the maximum memory for the provisioned LPAR.

Example: "1024"

LPAR Memory Desired

Specifies the ideal memory for the provisioned LPAR.

Example: "512"

LPAR Virtual Ethernet IEEE

Indicates whether the LPAR is virtual ethernet IEEE compliant.

Example: "false"

LPAR Virtual SCSI Remote LPAR

Specifies the remote LPAR associated with the virtual SCSI adapter for the provisioned LPAR.

Example: "lpar_hostname"

LPAR Virtual SCSI Adapter Is Client

Indicates whether the virtual SCSI adapter is a client for the provisioned LPAR.

Default: "true"

LPAR Virtual SCSI Adapter Is Required

Indicates whether the virtual SCSI adapter is required for the provisioned LPAR.

Default: "false"

LPAR NIM Provision Res Grp

Provisions an LPAR partition and installs an operating system based on a NIM image. The selection is made from a resource group.

Component

img

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Profile

Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

Provision LPAR

Indicates whether to provision an empty LPAR before installing the operating system. Select "false" if the LPAR exists and to install the operating system to it.

NIM Master

Specifies the name of the NIM master host where the list of available NIM images is obtained.

Example: "hostnamea"

NIM Machine Resource Name

Specifies the name of the NIM server used to install the operating system image to the provisioned LPAR.

Example: "machine_to_copy_from"

NIM Installation Type

Specifies the type of NIM installation to perform with the following options:

- rte - run-time environment
- mksysb - template based

NIM Resource Group

Specifies the name of the NIM resource group to use to provision the LPAR.

Example: "resource_group"

Target Username

Specifies the user name for the target machine for operating system installation.

Example: "root"

Target Password

Specifies the password for the target user name.

Example: "password"

Server Automation Template

(Optional) Specifies the server automation template to deploy to the LPAR.

Example: "software_package_template"

Agents Auto Deploy

(Optional) Indicates whether to deploy CA Configuration Automation, Performance, and Asset Management agents to the LPAR.

Example: "false"

ITCM Scalability Server

(Optional) Specifies the ITCM/DSM scalability server to use to deploy the software delivery agent.

Example: "itcm_server"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

LPAR Min CPU Requested

Specifies the minimum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Max CPU Requested

Specifies the maximum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Num of CPU Desired

Specifies the ideal number of CPUs for this LPAR provisioning.

Example: "1"

LPAR CPU Shared Mode

Indicates whether to provision the CPU for the LPAR in shared mode.

Default: "true"

LPAR Min CPU Unit Requested

Specifies the minimum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Max CPU Unit Requested

Specifies the maximum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Desired Min CPU Unit

Specifies the ideal CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR CPU Capped

Indicates whether the CPU for the provisioned LPAR is capped.

Default: "false"

LPAR CPU Uncapped Weight

Specifies the uncapped weight for the provisioned LPAR.

Example: "128"

LPAR Minimum Memory

Specifies the minimum memory for the provisioned LPAR.

Example: "512"

LPAR Maximum Memory

Specifies the maximum memory for the provisioned LPAR.

Example: "1024"

LPAR Memory Desired

Specifies the ideal memory for the provisioned LPAR.

Example: "512"

LPAR Virtual Ethernet IEEE

Indicates whether the LPAR is virtual ethernet IEEE compliant.

Example: "false"

LPAR Virtual SCSI Remote LPAR

Specifies the remote LPAR associated with the virtual SCSI adapter for the provisioned LPAR.

Example: "lpar_hostname"

LPAR Virtual SCSI Adapter Is Client

Indicates whether the virtual SCSI adapter is a client for the provisioned LPAR.

Default: "true"

LPAR Virtual SCSI Adapter Is Required

Indicates whether the virtual SCSI adapter is required for the provisioned LPAR.

Default: "false"

LPAR NIM Provision Res Grp-IVM

Provisions an LPAR partition and installs an operating system based on a NIM image. The selection is made from a resource group. The IVM (Integrated Virtual Manager) engine is used to provision the image.

Component

img

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Profile

Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

Provision LPAR

Indicates whether to provision an empty LPAR before installing the operating system. Select "false" if the LPAR exists and to install the operating system to it.

NIM Master

Specifies the name of the NIM master host where the list of available NIM images is obtained.

Example: "hostnamea"

NIM Machine Resource Name

Specifies the name of the NIM server used to install the operating system image to the provisioned LPAR.

Example: "machine_to_copy_from"

NIM Installation Type

Specifies the type of NIM installation to perform with the following options:

- rte - run-time environment
- mksysb - template based

NIM Resource Group

Specifies the name of the NIM resource group to use to provision the LPAR.

Example: "resource_group"

Target Username

Specifies the user name for the target machine for operating system installation.

Example: "root"

Target Password

Specifies the password for the target user name.

Example: "password"

Server Automation Template

(Optional) Specifies the server automation template to deploy to the LPAR.

Example: "software_package_template"

Agents Auto Deploy

(Optional) Indicates whether to deploy CA Configuration Automation, Performance, and Asset Management agents to the LPAR.

Example: "false"

ITCM Scalability Server

(Optional) Specifies the ITCM/DSM scalability server to use to deploy the software delivery agent.

Example: "itcm_server"

LPAR Virtual Slots

Specifies the maximum number of virtual slots to use for the LPAR provisioning.

Default: "10"

LPAR Min CPU Requested

Specifies the minimum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Max CPU Requested

Specifies the maximum number of CPUs for this LPAR provisioning.

Example: "1"

LPAR Num of CPU Desired

Specifies the ideal number of CPUs for this LPAR provisioning.

Example: "1"

LPAR CPU Shared Mode

Indicates whether to provision the CPU for the LPAR in shared mode.

Default: "true"

LPAR Min CPU Unit Requested

Specifies the minimum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Max CPU Unit Requested

Specifies the maximum CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR Desired Min CPU Unit

Specifies the ideal CPU unit slice for the LPAR provisioning.

Example: "0.2"

LPAR CPU Capped

Indicates whether the CPU for the provisioned LPAR is capped.

Default: "false"

LPAR CPU Uncapped Weight

Specifies the uncapped weight for the provisioned LPAR.

Example: "128"

LPAR Minimum Memory

Specifies the minimum memory for the provisioned LPAR.

Example: "512"

LPAR Maximum Memory

Specifies the maximum memory for the provisioned LPAR.

Example: "1024"

LPAR Memory Desired

Specifies the ideal memory for the provisioned LPAR.

Example: "512"

LPAR Virtual Ethernet IEEE

Indicates whether the LPAR is virtual ethernet IEEE compliant.

Example: "false"

LPAR Virtual SCSI Remote LPAR

Specifies the remote LPAR associated with the virtual SCSI adapter for the provisioned LPAR.

Example: "lpar_hostname"

LPAR Virtual SCSI Adapter Is Client

Indicates whether the virtual SCSI adapter is a client for the provisioned LPAR.

Default: "true"

LPAR Virtual SCSI Adapter Is Required

Indicates whether the virtual SCSI adapter is required for the provisioned LPAR.

Default: "false"

LPAR Remove LPAR CPU

Removes CPU units from an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Virtual CPU

Specifies the number of CPUs for the LPAR.

Example: "1"

LPAR CPU Adjustment

Specifies the adjustment value for the CPU processor assigned to the LPAR.

Example: "0.2"

LPAR CPU Adjustment Type

Specifies the type of adjustment for the LPAR CPU. Available values are "dynamic" and "all".

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Remove LPAR Memory

Removes memory from an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Memory Adjustment

Specifies the adjustment value for the memory assigned to the LPAR in megabytes.

Example: "128"

LPAR Memory Adjustment Type

Specifies the type of adjustment for the LPAR memory. Available values are "dynamic" or "all".

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Restart LPAR

Restarts an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Restart Type

Specifies the type of restart to perform. Available values are "immediate", "os_shutdown", or "os_shutdown_immediate".

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Shutdown LPAR

Shuts down an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Shutdown Type

Specifies the type of shutdown to perform. Available values are "delayed", "immediate", "os_shutdown", or "os_shutdown_immediate".

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

LPAR Start LPAR

Starts an LPAR partition.

Component

lpar

HMC/IVM Server Name

Specifies the IBM PowerVM (HMC/IVM) server name.

Example: "hostnamea"

Managed System Name

Specifies the IBM LPAR management system name.

Example: "lpar.any.com"

LPAR Partition Name

Specifies the name of the LPAR partition.

Example: "lpar_a"

LPAR Activation Override Key

Specifies the type of manual activation overrides available for the LPAR. Available values are:

1 – normal – After power-on, the system operates in normal (unattended) mode

2 – manual – After power-on, the system operates in manual (attended) mode using the control panel to direct the system

3 – donotoverride – Use the current partition settings

Example: "1"

LPAR Activation Boot Mode

Specifies the type of boot the LPAR performs. Available options are:

1 – normal

Example: "1"

LPAR Profile

(Optional) Specifies the profile to use for the LPAR partition.

Example: "Default"

RSI Connectors

A complete list of RSI connectors follows.

More information:

RSI Get Image List (see page 163)
RSI Job Status (see page 163)
RSI Operation Status (see page 164)
RSI OS Type List (see page 164)
RSI Perform Image Capture (see page 165)
RSI Perform Image Deployment (see page 166)
RSI Perform Image Removal (see page 167)
RSI Validate Server Access (see page 168)

RSI Get Image List

Retrieves captured images stored in the RSI server.

Component

rsi

RSI Host

Specifies the machine name hosting the RSI server.

Example: "RacemiBoot"

Dynacenter Username

Specifies the user name to log in to the RSI server.

Example: "admin"

Dynacenter Password

Specifies the password for the DynaCenter user name.

Example: "dynacenter"

RSI Job Status

The RSI Job Status connector is superseded by the Imaging Job Status (see page 112) connector.

RSI Operation Status

Retrieves the status of services running on an RSI server.

Component

rsi

RSI Target Host

Specifies the machine name hosting the RSI server.

Example: "RacemiBoot"

RSI Target Username

Specifies a user name to log in to the host machine.

Example: “root”

Dynacenter Username

Specifies the user name to log in to the RSI server.

Example: “admin”

Dynacenter Password

Specifies the password for the DynaCenter user name.

Example: “dynacenter”

RSI OS Type List

Retrieves a list of valid OS types.

Component

rsi

RSI Target Host

Specifies the machine name hosting the RSI server.

Example: "RacemiBoot"

RSI Target Username

Specifies a user name to log in to the host machine.

Example: “root”

Dynacenter Username

Specifies the user name to log in to the RSI server.

Example: “admin”

Dynacenter Password

Specifies the password for the DynaCenter user name.

Example: “dynacenter”

RSI Perform Image Capture

Captures system and network information from the target server.

Component

rsi

RSI Host

Specifies the machine name hosting the RSI server.

Example: “RacemiBoot”

Dynacenter Username

Specifies the user name to log in to the RSI server.

Example: “admin”

Dynacenter Password

Specifies the password for the DynaCenter user name.

Example: “dynacenter”

Image Name

Specifies the image name.

Example: “userid_RedHat_Linux_EL_4_ITPAM”

Image Description

(Optional) Specifies a description for the deployed image.

OS Type

Specifies a valid operating system type in the RSI server.

Example: “RedHat_Linux-El_4_*-*”

Mac Address To Capture From

Specifies the Mac Address of the server to capture.

Example: "00:50:56:AB:16:BC"

Server Id

(Optional) Specifies a unique ID for the server identified for the Mac Address. If a value is provided, the 'Mac Address to Capture From' parameter is ignored.

Image Host

(Optional) Specifies the host name or IP address of the server.

Example: "RacemiDeploy2"

Image Profile

(Optional) Specifies the profile source to use.

Example: The URL or '/tmp/profile' location.

Capture Type

(Optional) Specifies the type of capture. Available values are "live" or "offline".

Default: "live"

Boot Network

(Optional) Specifies the boot network to use for the target server.

RSI Perform Image Deployment

Deploys images.

Component

rsi

RSI Host

Specifies the machine name hosting the RSI server.

Example: "RacemiBoot"

Dynacenter Username

Specifies the user name to log in to the RSI server.

Example: "admin"

Dynacenter Password

Specifies the password for the DynaCenter user name.

Example: "dynacenter"

Image Name

Specifies the image name.

Example: "userid_RedHat_Linux_EL_4_ITPAM"

Image Description

(Optional) Specifies a description for the deployed image.

OS Type

Specifies a valid operating system type in the RSI server.

Example: “RedHat_Linux-El_4_*-*”

MAC Address

Specifies the MAC address of the newly deployed server.

Example: "00:50:56:AB:16:BC"

Server Id

(Optional) Specifies a unique ID for the server identified for the Mac Address.

Image Host

(Optional) Specifies the host name or IP address of the server.

Example: “RacemiDeploy2”

Image Profile

(Optional) Specifies the profile source to use.

Example: The URL or ‘/tmp/profile’ location.

Image Scale

Indicates whether to perform image scaling. Available values are "yes" or "no".

Default: "yes"

Boot Network

(Optional) Specifies the boot network to use for the target server.

RSI Perform Image Removal

Deletes a previously captured image from the RSI store.

Component

rsi

RSI Target Host

Specifies the machine name hosting the RSI server.

Example: "RacemiBoot"

RSI Target Username

Specifies a user name to log in to the host machine.

Example: “root”

RSI Target Password

Specifies a password for the RSI user name.

Example: “dynacenter”

Dynacenter Username

Specifies the user name to log in to the RSI server.

Example: "admin"

Dynacenter Password

Specifies the password for the DynaCenter user name.

Example: "dynacenter"

Image Name

Specifies the image name.

Example: "userid_RedHat_Linux_EL_4_ITPAM"

RSI Validate Server Access

Validates the provided credentials against an RSI server.

Component

rsi

RSI Host

Specifies the machine name hosting the RSI server.

Example: "RacemiBoot"

Dynacenter Username

Specifies the user name to log in to the RSI server.

Example: "admin"

Dynacenter Password

Specifies the password for the DynaCenter user name.

Example: "dynacenter"

Software Delivery Connectors

A complete list of Software Delivery connectors follows.

More information:

Add Personality (see page 169)
Check Software Delivery Status (see page 170)
Get Software Package List (see page 170)
Get Software Package Procedure (see page 170)
ITCM Get Software Job Status (see page 171)
ITCM OS Image List (see page 171)
ITCM OS Imaging Parameters (see page 171)
ITCM Server Info (see page 171)
Provision OSIM Image (see page 172)

Add Personality

Adds a new personality.

Component

sda

SD Host Name

Specifies the target server where the software package is delivered.

Example: “localhost.anycompany.com”

SD Package Name

Specifies the name of the package as it appears in the CA Server Automation UI. Use this name without the version number.

Example: “CCA Agent Solaris”

SD Package Version

Specifies the version for the software package.

Example: “r5.0”

SD Procedure Type

Specifies the type of procedure that the connector performs.

Example: “INSTALL”

SD Target Machine User

Specifies a user with administrator privileges for the target server.

Example: “root” or “Administrator”

SD Target Machine Password

Specifies the password for the user.

Example: “rootpassword” or “Adminpassword”

SD Operating System Type

Specifies the operating system installed in the target server.

Example: “Solaris”

SD Procedure

Specifies the procedure that the connector performs.

Example: “Install”

SD Scalability Server

(Optional) Specifies the scalability server related to the software delivery server.

Example: “remotehost.anycompany.com”

Check Software Delivery Status

Verifies the status of a software delivery package.

Component

sd

SD Job ID

Specifies the unique identifier for a software package job ID generated by the Add Personality connector.

Example: “1234”

Get Software Package List

Gets the available software package list.

Component

sd

No parameters are required.

Get Software Package Procedure

Gets the available software package procedures.

Component

sd

SD Software Package ID

Specifies the unique identification for a package.

Example: “EC30FD7386F34E1F8F09B084F5FBF7CB”

ITCM Get Software Job Status

The ITCM Get Software Job Status connector is superseded by the Imaging Job Status (see page 112) connector.

ITCM OS Image List

Gets the available list of OS images in an ITCM server.

Component

sda

No parameters are required.

ITCM OS Imaging Parameters

Gets the OS images parameters in an ITCM server.

Component

sda

SD OS Image Name

Specifies the software delivery OS image name.

Example: "ITCM Windows 2008"

SD OS Category Id

Specifies the ID of the ITCM OS category to use.

Example: “ITCM Windows 2008"

ITCM Server Info

Gets information about the ITCM server.

Component

sda

No parameters are required.

Provision OSIM Image

Creates an OSIM image.

Component

img

Auto Deploy Agent

Indicates whether to trigger agent delivery after provisioning.

Example: “Yes”, “No”

Target User Name

Specifies the administrator name for the new machine.

Example: “administrator”

Target Password

Specifies a password for the administrator user.

Example: “password”

Scalability Server

(Optional) Specifies a fully qualified name.

Example: “localhost.anycompany.com”

Target Machine

Specifies the name of the machine to provision.

Example: “TestServer”

Image Name

Specifies the template name to use to provision.

Target Machine MAC Address

Specifies the MAC address for the new machine.

Example: “000d5612B1FF”

UCS Connectors

A complete list of Platform Support UCS connectors follows.

More information:

UCS Associate Service Profile

Associates a UCS service profile with a given UCS blade.

Component

ucsws

UCS Manager Host

Specifies the fully qualified name of the host where the UCS manager is located.

Example: "localhostUCS.anyco.com"

UCS Blade

Specifies the full path provided by either the user or the UCS Blade List connector.

Example: "sys/chassis-1/blade-1"

UCS Service Profile

Specifies the full name of the service profile.

Example: "org-root/ls-dev-boot-from-san-esx"

UCS Blade Power Off

Shuts down a UCS blade.

Component

ucsws

UCS Manager Host

Specifies the fully qualified name of the host where the UCS manager is located.

Example: "localhostUCS.anyco.com"

UCS Blade

Specifies the full path provided by either the user or the UCS Blade List connector.

Example: "sys/chassis-1/blade-1"

UCS Blade Power On

Starts a UCS blade.

Component

ucsws

UCS Manager Host

Specifies the fully qualified name of the host where the UCS manager is located.

Example: “localhostUCS.anyco.com”

UCS Blade

Specifies the full path provided by either the user or the UCS Blade List connector.

Example: “sys/chassis-1/blade-1”

UCS Blade Reset

Resets a UCS blade.

Component

ucsws

UCS Manager Host

Specifies the fully qualified name of the host where the UCS manager is located.

Example: “localhostUCS.anyco.com”

UCS Blade

Specifies the full path provided by either the user or the UCS Blade List connector.

Example: “sys/chassis-1/blade-1”

UCS Disassociate Service Profile

Disassociates a UCS service profile from a UCS blade.

Component

ucsws

UCS Manager Host

Specifies the fully qualified name of the host where the UCS manager is located.

Example: “localhostUCS.anyco.com”

UCS Service Profile

Specifies the full name of the service profile.

Example: “org-root/ls-dev-boot-from-san-esx”

UCS List Blades

Provides a list of blades in a UCS chassis.

Component

ucsws

UCS Manager Host

Specifies the fully qualified name of the host where the UCS manager is located.

Example: “localhostUCS.anyco.com”

UCS Power Operations

Provides individualized Cisco UCS blade power operations.

Component

ucsws

UCS Manager Host

Specifies the fully qualified name of the host where the UCS manager is located.

Example: “localhostUCS.anyco.com”

UCS Blade

Specifies the full path provided by either the user or the UCS Blade List connector.

Example: “sys/chassis-1/blade-1”

Operation

Specifies the operation to perform on a UCS blade. Available values are:

2 – Cycle-immediate

3 – Cycle-wait

4 – Hard-reset-immediate

5 – Hard-reset-wait

6 – Soft-shut-down

7 – Shut-down

8 – Boot-up

Example: “Cycle-immediate”

UCS Service Profile Operation

Provides Cisco UCS service profile operations.

Component

ucsws

Operation

Specifies the operation to perform. Available values are:

SelectiveImportNoQueue

SelectiveExportNoQueue

SelectiveImportDelNoQueue

UCS Manager

Specifies the name of the configured UCS manager.

Example: “localhostUCS.anyco.com”

Service Profile Name

Specifies the name of the service profile.

Example: "pam_test_sp"

When using the selective export operation the name should be "*UCS manager name|service profile name*".

Example: "localhostUCS.anyco.com|pam_test_sp"

Organization Distinguished Name

Specifies the corresponding name of an organization. The name should start with "org-".

Example: "org-name"

VMware Connectors

A complete list of Platform Support VMware connectors follows.

More information:

Get Machine Status VC

Gets the status of a vCenter machine.

Component

img

Data Center Name

Specifies the name of the data center located in the VC server.

Example: "North-2/CA Server Automation"

Virtual Machine Name

Specifies the fully qualified name of the target VM.

Example: "localhost.anycompany.com"

VC Server

Specifies the Virtual Center server.

Example: "server.mycompany.com"

Get VC Image List

Gets a list of available vCenter machines.

Component

img

No parameters are required.

Get VM Properties

Gets VMware VM properties.

Component

vc

Data Center Name

Specifies the name of the data center where the server is located.

Example: “DC Datacenter”

VC Server

Specifies the fully qualified name for the computer where the Virtual Center Server is located.

Example: “localhostVC.anycompany.com”

Power On VC

Starts a vCenter system.

Component

img

Data Center Name

Specifies the name of the data center where the server is located.

Example: “DC Datacenter”

Virtual Machine Name

Specifies the fully qualified name of the target VM.

Example: “localhost.anycompany.com”

VC Server

Specifies the fully qualified name for the computer where the Virtual Center Server is located.

Example: “localhostVC.anycompany.com”

Power Off VC

Turns off a vCenter system.

Component

img

Data Center Name

Specifies the name of the data center where the server is located.

Example: “DC Datacenter”

Virtual Machine Name

Specifies the fully qualified name of the target VM.

Example: “localhost.anycompany.com”

VC Server

Specifies the fully qualified name for the computer where the Virtual Center Server is located.

Example: “localhostVC.anycompany.com”

Provision VC Image

Creates a vCenter image.

Component

img

Data Center Name

Specifies the name of the data center where the server is located.

Example: “DC Datacenter”

Auto Deploy Agents

(Optional) Indicates whether to deploy agents after provisioning.

Example: “No”

Target User Name

Specifies the administrator name for the new machine.

Example: “administrator”

Target Password

Specifies a password for the administrator user.

Example: “password”

Clone Name

Specifies the name of the new machine to provision.

Example: "LocalhostTest"

Datastore Name

Specifies the data store name that corresponds to the data center.

Example: "storage1 (2)"

Scalability Server

Specifies the fully qualified name of the scalability server.

Example: "remotehost.anycompany.com"

Compute Resource Name

Specifies the fully qualified name of the server where this compute resource is located.

Example: "server1.anycompany.com"

ESX Host Server

Specifies the fully qualified name of the ESX server.

Example: "localhost.anycompany.com"

Template Name

Specifies the template to use for provisioning.

Example: "DCRMFolder/Base"

Name of VM used instead of template

(Optional) Specifies a VM machine to use instead of a template for provisioning.

Resource Pool Name

Specifies the pool name to use.

Example: "Resources/DcrmPool_test"

Specification Name

Specifies a name for the template to use. Consult the VC server or CA Server Automation for valid values.

Example: "VolumeSPEC"

VM Computer Name

Specifies the name of the VM to clone. This value is the same as the Clone Name.

Image Host

(Optional) Specifies the host where the image resides.

Imaging Server Type

Specifies the type of provisioning desired. Always use “VC".

Provision VM Image Linux

Creates a vCenter image on Linux.

Component

img

Imaging Server Type

Specifies the type of provisioning desired. Always use “VC".

VC Server

Specifies the name of the VC center.

VC Data Center

Specifies the name of the data center where this host is located.

Example: “DC Datacenter”

VC Compute Resource

Specifies the fully qualified name of the server where this compute resource is located.

Example: “server1.anycompany.com”

VC ESX Host Server

Specifies the fully qualified name of the server where the ESX server is installed.

Example: “localhost.anycompany.com”

VC Datastore Name

Specifies the data store name that corresponds to the data center.

Example: “storage1 (2)”

VC Target Location (resource pool)

Specifies the resource pool name to use.

Example: “Resources/DcrmPool_test”

Hostname/VM Name

Specifies the name of the new host to provision.

Example: “LocalhostTest”

VC User Name

Specifies the administrator name for the new host.

Example: “administrator”

VC User Password

Specifies the password for the user name.

Example: “password”

VC Virtual Machine

(Optional) Specifies a VM machine to use instead of a template for provisioning.

VC Template Name

Specifies the template to use for provisioning.

Example: “DCRMFolder/Base”

VC Specification Name

Specifies a name for the template to use. Consult the VC server or CA Server Automation for valid values.

Example: “VolumeSPEC”

VM OS Type

Specifies the default operating system type to use for template creation.

Example: “Linux”

Memory Size

Specifies the amount of allocated virtual memory for the new VM image. This value is defined in megabytes, for example, 1024MB = 1GB.

Virtual Processors

Specifies the number of virtual CPUs to allocate to the new VM image. Valid values are 1, 2, or 4.

Disk 1 Size

Specifies the default size of the virtual hard disk for the VM image. This value is in gigabytes, for example 6GB.

Disk 1 Datastore

Specifies the name of the VMware datastore to use to create the virtual disk for the new image. Use the same datastore where the VM image is going to be created.

Virtual Hard Disk Controller Key

Specifies the SCSI controller key to use to create the virtual disk for the new VM image.

NIC IP Address

Specifies the network IP address (v4) for the new system.

NIC Default Gateway

(Optional) Specifies the default gateway for the new VM image.

NIC Alt Gateway

(Optional) Specifies the IP address of the alternate gateway.

NIC Subnet Mask

(Optional) Specifies the default network subnet mask to use.

Example: "255.255.255.0"

Network Connection

(Optional) Specifies the name of the network interface.

DNS Search Suffix

(Optional) Specifies the global DNS search suffix.

Primary DNS Server

(Optional) Specifies the global primary DNS entry.

Secondary DNS Server

(Optional) Specifies the global secondary DNS entry.

Tertiary DNS Server

(Optional) Specifies the third global DNS entry.

Global Domain Name

(Optional) Specifies the global domain name.

SD Scalability Server

(Optional) Specifies the software delivery scalability server to use.

Deployment Template

(Optional) Specifies the deployment template to use.

Auto Deploy Agents

(Optional) Indicates whether to deploy agents after provisioning.

Example: "No"

Deploy SD Agents

(Optional) Indicates whether to deploy Software Delivery agents automatically.

Default: "No"

Provision VM Image Windows

Creates a vCenter image on Windows.

Component

img

Imaging Server Type

Specifies the type of provisioning desired. Always use "VC".

VC Server

Specifies the name of the VC center.

VC Data Center

Specifies the name of the data center where this host is located.

Example: "DC Datacenter"

VC Compute Resource

Specifies the fully qualified name of the server where this compute resource is located.

Example: "server1.anycompany.com"

VC ESX Host Server

Specifies the fully qualified name of the server where the ESX server is installed.

Example: "localhost.anycompany.com"

VC Datastore Name

Specifies the data store name that corresponds to the data center.

Example: "storage1 (2)"

VC Target Location (resource pool)

Specifies the resource pool name to use.

Example: “Resources/DcrmPool_test”

Hostname/VM Name

Specifies the name of the new host to provision.

Example: “LocalhostTest”

VC User Name

Specifies the administrator name for the new host.

Example: “administrator”

VC User Password

Specifies the password for the user name.

Example: “password”

VC Virtual Machine

(Optional) Specifies a VM machine to use instead of a template for provisioning.

VC Template Name

Specifies the template to use for provisioning.

Example: “DCRMFolder/Base”

VC Specification Name

Specifies a name for the template to use. Consult the VC server or CA Server Automation for valid values.

Example: “VolumeSPEC”

VM OS Type

Specifies the default operating system type to use for template creation.

Example: "Windows"

Memory Size

Specifies the amount of allocated virtual memory for the new VM image. This value is defined in megabytes, for example, 1024MB = 1GB.

Virtual Processors

Specifies the number of virtual CPUs to allocate to the new VM image. Valid values are 1, 2, or 4.

Disk 1 Size

Specifies the default size of the virtual hard disk for the VM image. This value is in gigabytes, for example 6GB.

Disk 1 Datastore

Specifies the name of the VMware datastore to use to create the virtual disk for the new image.

Virtual Hard Disk Controller Key

Specifies the SCSI controller key to use to create the virtual disk for the new VM image.

NIC IP Address

Specifies the network IP address (v4) for the new system.

NIC Default Gateway

(Optional) Specifies the default gateway for the new VM image.

NIC Alt Gateway

(Optional) Specifies the IP address of the alternate gateway.

NIC Subnet Mask

(Optional) Specifies the default network subnet mask to use.

Example: "255.255.255.0"

Network Connection

(Optional) Specifies the name of the network interface.

DNS Search Suffix

(Optional) Specifies the global DNS search suffix.

WINS Primary

(Optional) Specifies the primary WINS entry.

WINS Secondary

(Optional) Specifies the secondary WINS entry.

DNS Server

(Optional) Specifies the preferred DNS entry.

Alt. DNS

(Optional) Specifies an alternate DNS entry.

SD Scalability Server

(Optional) Specifies the software delivery scalability server to use.

Template Name

(Optional) Specifies the software delivery template to use.

Auto Deploy Agents

(Optional) Indicates whether to deploy agents after provisioning.

Example: "No"

Deploy SD Agents

(Optional) Indicates whether to deploy Software Delivery agents automatically.

Default: "No"

Shutdown VC Image

Shuts down a Virtual Center image.

Component

img

Data Center Name

Specifies the name of the data center where the server is located.

Example: "DC Datacenter"

Virtual Machine Name

Specifies the fully qualified name of the target VM.

Example: "localhost.anycompany.com"

VC Server

Specifies the fully qualified name for the computer where the Virtual Center Server is located.

Example: "localhostVC.anycompany.com"

Validate VC Imaging Server

Validates the VC imaging server.

Component

img

Target Username

Specifies the user to log in to the server.

Example: “administrator”

Target Username Password

Specifies the password for the target user name.

Example: “Password”

Image Host

Specifies the fully qualified name for the server.

Example: “localhost.ca.com”

Image Host Port

Specifies the assigned port.

Example: “4443”

Protocol

Specifies the protocol used.

Example: “HTTPS”

Proxy Port

Specifies the port used for the proxy.

Example: “48008”

VC Add Virtual NIC

Adds an additional virtual NIC to a VMware VM.

Component

vc

VMware Virtual Center Name

Specifies the name of the Virtual Center server where the VM is located.

Example: "vmwareVC"

VMware VM Name

Specifies the name of the VM to update.

Example: "vmware vm"

VMware Network Device Type

Specifies the VMware network device type to add.

Example: "E1000"

VMware Network

Specifies the VMware network to add the new NIC to.

Example: "VMNetwork"

VC Add VM Disk

Adds an additional disk to a VMware VM.

Component

vc

VMware VC Name

Specifies the name of the Virtual Center server where the VM is located.

Example: "vmwareVC"

VMware VM Name

Specifies the name of the VM to update.

Example: "vmware vm"

VMware Datastore Name

Specifies the name of the datastore where the VM resides.

Example: "storage1 (1)"

Disk Size

Specifies the size of the disk to add in megabytes.

Example: "150"

Thin Provisioning

Indicates whether to provision as a thin disk.

Default: "false"

VC Job Status

The VC Job Status connector is superseded by the Imaging Job Status (see page 112) connector.

VC Remove Virtual NIC

Removes a virtual NIC from a VMware VM.

Component

vc

VMware Virtual Center Name

Specifies the name of the Virtual Center server where the VM is located.

Example: "vmwareVC"

VMware VM Name

Specifies the name of the VM to update.

Example: "vmware vm"

VMware Virtual NIC Key

Specifies the key of the VMware virtual NIC to remove.

Example: "4001"

VC Update VM CPU

Updates the vCPUs in a VMware VM.

Component

vc

VMware Datacenter Name

Specifies the data center name where the VM resides.

Example: "vmware datacenter"

VMware ESX Hostname

Specifies the ESX hostname where the VM resides.

Example: "esxhost.company.com"

VMware VM Name

Specifies the name of the VM to update.

Example: "vmware vm"

CPU Value

Specifies the number of CPUs to use. Available values are 1, 2, 4, or 8.

Example: "1"

VC Update VM Memory

Updates the memory in a VMware VM.

Component

vc

VMware Datacenter Name

Specifies the data center name where the VM resides.

Example: "vmware datacenter"

VMware ESX Hostname

Specifies the ESX hostname where the VM resides.

Example: "esxhost.company.com"

VMware VM Name

Specifies the name of the VM to update.

Example: "vmware vm"

CPU Value

Specifies the number of CPUs to use. Available values are 1, 2, 4, or 8.

Example: "1"

SSRM Connectors

A complete list of SSRM connectors follows.

More information:

SSRM Cancel Reservation (see page 192)
SSRM Check System Availability (see page 192)
SSRM Create Reservation (see page 193)
SSRM Extend Reservation (see page 195)
SSRM Get Data Software (see page 196)
SSRM Get Data Template (see page 197)
SSRM Get Resrc Pool (see page 198)
SSRM Get System Requirements (see page 199)
SSRM Get VM Res Name (see page 200)
SSRM Res Status (see page 200)
SSRM Return Res System (see page 201)
SSRM Verify User (see page 201)

SSRM Cancel Reservation

Cancels an existing reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: “ssrm_user”

SSRM Reservation ID

Specifies the ID of the reservation.

Example: “2”

SSRM Check System Availability

Verifies the system availability for a Reservation Manager reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: “ssrm_user”

SSRM Org Unit

Specifies the Reservation Manager organizational unit.

Example: “Public”

SSRM Start DateTime

Specifies the date and time to start verifying the availability window. Define the time in a 24 hour clock, and provide the Zulu timezone.

Example: “2010-10-15 12:00:00Z”

SSRM End DateTime

Specifies the date and time to stop verifying the availability window. Define the time in a 24 hour clock, and provide the Zulu timezone.

Example: “2010-10-16 12:00:00Z”

SSRM System Image ID

Specifies the system image ID (VM template) defined in the Reservation Manager system.

Example: “4939434-49493843-sd”

SSRM Resource Count

Specifies the number of systems to create during this reservation.

Example: "1"

SSRM Min CPUs

Specifies the number of CPUs to assign to each reserved system. Available values are 1, 2, or 4.

Example: "1"

SSRM Min Memory

Specifies the minimum amount of memory assigned to the reserved system. This value is controlled based on the Reservation Manager image value.

Example: "1024"

SSRM Min Disk

Specifies the minimum size of reserved system disk. This value is controlled based on the Reservation Manager image value.

Example: "6GB"

SSRM Resource Pool

Specifies the resource pool where the reservation resides.

Example: “ssrm-demo”

SSRM Create Reservation

Creates a reservation in Reservation Manager.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: “ssrm_user”

SSRM Org Unit Name

Specifies the Reservation Manager organizational unit.

Example: “Public”

SSRM Start DateTime

Specifies the date and time to start verifying the availability window. Define the time in a 24 hour clock, and provide the Zulu timezone.

Example: “2010-10-15 12:00:00Z”

SSRM End DateTime

Specifies the date and time to stop verifying the availability window. Define the time in a 24 hour clock, and provide the Zulu timezone.

Example: “2010-10-16 12:00:00Z”

SSRM System Image ID

Specifies the system image ID (VM template) defined in the Reservation Manager system.

Example: “4939434-49493843-sd”

SSRM Project ID

(Optional) Specifies a user created project ID from the project management tool.

SSRM Notes

(Optional) Specifies notes for this reservation.

SSRM Send Notifications

Specifies the email address of the user to receive notification when this reservation is ready.

SSRM Save As Template

(Optional) Indicates whether to save this reservation as a template for future use.

SSRM Template Description

(Optional) Specifies a description for the template when the reservation is saved as a template.

SSRM Resource Count

Specifies the number of systems to create during this reservation.

Example: "1"

SSRM Min CPUs

Specifies the number of CPUs to assign to each reserved system. Available values are 1, 2, or 4.

Example: "1"

SSRM Min Memory

Specifies the minimum amount of memory assigned to the reserved system. This value is controlled based on the Reservation Manager image value.

Example: "1024"

SSRM Min Disk

Specifies the minimum size of reserved system disk. This value is controlled based on the Reservation Manager image value.

Example: "6GB"

SSRM Resource Pool

Specifies the resource pool where the reservation resides.

Example: “ssrm-demo”

SSRM Software Group

Specifies the Reservation Manager system where software packages can be deployed in a reservation.

SSRM Add Disk 2-7

(Optional) Specifies information for additional disks. Values are represented as 100MB or 10GB.

SSRM Extend Reservation

Extends an existing reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: “ssrm_user”

SSRM Is User Admin

Indicates whether the user is an administrator. Available values are "0" (false) or "1" (true).

Example: "0"

SSRM Reservation ID

Specifies the ID of the reservation.

Example: “2”

SSRM Res Extension Time

Specifies the date and time to extend the reservation to.

Example: "27-06-201112:00:00"

SSRM Get Data Software

Get the necessary software groups for a Reservation Manager reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: "ssrm_user"

SSRM Org Unit Name

Specifies the Reservation Manager organizational unit.

Example: "Public"

Caller Type

Specifies the type of call made to obtain the available software groups defined in the Reservation Manager system. The default value cannot be changed.

Default: "user"

Callback ID

Specifies the callback ID for this connector. The default value cannot be changed.

Default: "1"

CLS

Specifies the action to filter on. The default value cannot be changed.

Default: "SoftwareGroup"

Filter Type

Specifies the filter type to use in this connector. The default value cannot be changed.

Default: "SelectorSet"

SSRM Get Data Template

Gets the necessary system templates for a Reservation Manager reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: "ssrm_user"

SSRM Org Unit Name

Specifies the Reservation Manager organizational unit.

Example: "Public"

Caller Type

Specifies the type of call made to obtain the available software groups defined in the Reservation Manager system. The default value cannot be changed.

Default: "user"

Callback ID

Specifies the callback ID for this connector. The default value cannot be changed.

Default: "1"

CLS

Specifies the action to filter on. The default value cannot be changed.

Default: "Systemimage"

Filter Type

Specifies the filter type to use in this connector. The default value cannot be changed.

Default: "SelectorSet"

SSRM Get Resrc Pool

Gets the resource pools for a Reservation Manager reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: "ssrm_user"

SSRM Org Unit Name

Specifies the Reservation Manager organizational unit.

Example: "Public"

Caller Type

Specifies the type of call made to obtain the available software groups defined in the Reservation Manager system. The default value cannot be changed.

Default: "user"

Callback ID

Specifies the callback ID for this connector. The default value cannot be changed.

Default: "1"

CLS

Specifies the action to filter on. The default value cannot be changed.

Default: "ResourcePool"

Filter Type

Specifies the filter type to use in this connector. The default value cannot be changed.

Default: "SelectorSet"

SSRM Get System Requirements

Gets the necessary system requirements for a Reservation Manager reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: "ssrm_user"

SSRM Org Unit Name

Specifies the Reservation Manager organizational unit.

Example: "Public"

Caller Type

Specifies the type of call made to obtain the available software groups defined in the Reservation Manager system. The default value cannot be changed.

Default: "user"

Callback ID

Specifies the callback ID for this connector. The default value cannot be changed.

Default: "1"

CLS

Specifies the action to filter on. The default value cannot be changed.

Default: "MaxAllowedResources"

Filter Type

Specifies the filter type to use in this connector. The default value cannot be changed.

Default: "SelectorSet"

SSRM Image ID

Specifies the system ID assigned to the VM templates.

Example: "3940930-3949349-3948934"

SSRM Get VM Res Name

Returns the VM name for a reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: "ssrm_user"

SSRM Org Unit Name

Specifies the Reservation Manager organizational unit.

Example: "Public"

SSRM Reservation ID

Specifies the ID of the reservation.

Example: "2"

SSRM Res Status

Gets Reservation Manager reservation creation status.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: "ssrm_user"

SSRM Org Unit Name

Specifies the Reservation Manager organizational unit.

Example: "Public"

Caller Type

Specifies the type of call made to obtain the available software groups defined in the Reservation Manager system. The default value cannot be changed.

Default: "user"

Callback ID

Specifies the callback ID for this connector. The default value cannot be changed.

Default: "1"

CLS

Specifies the action to filter on. The default value cannot be changed.

Default: “ReservationDetail”

Filter Type

Specifies the filter type to use in this connector. The default value cannot be changed.

Default: “SelectorSet”

SSRM Reservation ID

Specifies the ID of the reservation.

Example: “2”

SSRM Return Res System

Returns a reserved system for either a single or multi system reservation.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: “ssrm_user”

SSRM Reservation ID

Specifies the ID of the reservation.

Example: “2”

SSRM Reservation System Name

Specifies the name of the system in a multi-VM reservation to return.

Example: "vm_name"

SSRM Verify User

Verifies that a user can create Reservation Manager reservations.

Component

caresman

SSRM Username

Specifies the name of the Reservation Manager user.

Example: “ssrm_user”

Storage Connectors

A complete list of Storage connectors follows.

More information:

Storage Create NAS Datastore

Creates a NAS (NFS, CIFS, MixedMode) datastore.

Component

vc

vCenter Name

Specifies the fully qualified domain name (FQDN) of the VMware vCenter server.

Example: “vpm-vc01.domain.com”

ESX Name

Specifies the fully qualified domain name (FQDN) of the ESX Server.

Example: “esxhost.com”

Datastore Name

Specifies the name of the datastore.

Example: “mystore10”

Access Mode

Specifies the access mode for the mount point. Available values are “readOnly” or “readWrite”.

Example: “readOnly”

Host Name

Specifies the host running the NFS server.

Example: “nfsserverhost01”

Host Path

Specifies the remote path of the NFS mount point.

Example: “/vol/nfsmountpoint_01/nfsmountpoint”

Type

Specifies the type of NAS volume. Available values are "CIFS" or "NFS".

Default: "NFS"

Storage Create SAN Datastore

Creates a SAN (FCP, SCSI) datastore.

Component

vc

vCenter Name

Specifies the fully qualified domain name (FQDN) of the VMware vCenter server.

Example: “vpm-vc01.domain.com”

ESX Name

Specifies the fully qualified domain name (FQDN) of the ESX Server.

Example: “esxhost.com”

Datastore Name

Specifies the name of the datastore.

Example: “mystore10”

Device Path

Specifies the storage path.

Example: "/vmfs/devices/disks/naa.01"

Block Size

Specifies the block size of the datastore. Available values are 1, 2, 4, or 8.

Example: “2”

Storage Deprovision

Deprovisions a dataset.

Component

spm

Dataset Name

Specifies the name of the dataset.

Example: “mydataset-01”

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: “host01.domain.com”

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Discover

Discovers storage objects such as datasets and resource pools.

Component

spm

Type

Specifies the type of object to discover. The available values are:

0 – Dataset

1 – Resource Pool

2 – Provision Policy

3 – Storage Service

4 – vFiler

6 – Storage System

Default: "0"

Filter

Specifies the XPath type filter to specify the criteria for data returned where a subset of all data is desired. This filter is a key and value pairing that works with two possible keys (name or type). For instance, if you want a list of all datasets that start with "QAtest" enter name=QAtest*. If you want a list of SAN or NAS provisioning policies, enter type=san or type=nas.

Example: “name=test*”

Detail

Specifies the level of detail to return. The available values are:

0 – Basic

1 – Simple

2 – Complete

Storage Discover has several objects it can discover and the details of each object are different. Basic returns minimal information (such as name and description). Simple includes more information than Basic. Complete returns the most detailed information.

Default: “0”

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: “host01.domain.com”

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: “user01”

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: “password01”

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: “8088”

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Get Available SCSI Disks

Get a list of the available SCSI disks.

Component

vc

vCenter Name

Specifies the fully qualified domain name (FQDN) of the VMware vCenter server.

Example: "vpm-vc01.domain.com"

ESX Name

Specifies the fully qualified domain name (FQDN) of the ESX Server.

Example: "esxhost.com"

Datastore Index

(Optional) Specifies the name of the datastore index.

Example: "datastore-0001"

To obtain the datastore index:

1. Navigate to Start -> Programs -> CA -> CA Server Automation - > CA Server Automation Command Prompt.
2. Enter the following command:

   ```
   C:\CA\ServerAutomation\bin>caaipaomwsclient /enumerate=CA_Datastore
   /queryFilter="Select * from CA_Datastore where ElementName='datastore'"
   /user=SA_user /password=SA_password
   ```

   Replace *datastore* with the name of the datastore, replace *SA_user* and *SA_password* with credentials for a CA Server Automation authorized user.

Use the value returned by 'Index' as the datastore index.

Storage Get Host HBA

Get a list of the host bus adapters (HBAs).

Component

vc

vCenter Name

Specifies the fully qualified domain name (FQDN) of the VMware vCenter server.

Example: “vpm-vc01.domain.com”

ESX Name

Specifies the fully qualified domain name (FQDN) of the ESX Server.

Example: “esxhost.com”

Storage Lun Break

Breaks a Lun through Snapmirror (the NetApp data replication option).

Component

spm

Destination Volume

Specifies a volume or Lun to receive information.

Example: “uservol01”

Destination Filer

Specifies a storage host for vFiler and Lun.

Example: “destfiler”

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: “host01.domain.com”

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: “user01”

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: “password01”

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Lun Status

Gets the status of an update Lun action.

Component

spm

Destination Filer

Specifies a storage host for vFiler and Lun.

Example: "destfiler"

Destination Volume

Specifies a volume or Lun to receive information.

Example: "uservol01"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Move

Performs a completed move (including both vFiler and Lun) and runs the following actions:

- Storage Lun Break
- Storage Lun Status
- Storage Move Lun
- Storage vFiler Active
- Storage vFiler Resync
- Storage vFiler Status
- Storage vFiler Stop
- Storage vLan Interface
- Storage Volume Offline

Note: Set up the Disaster Recovery Lun and vFiler before running this command.

Component

spm

Source Volume

Specifies the volume or Lun.

Example: "uservol1"

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: “sourcefiler01”

Destination Volume

Specifies a volume or Lun to receive information.

Example: “uservol01”

Destination Filer

Specifies a storage host for vFiler and Lun.

Example: “destfiler”

vFiler Name

Specifies the name of the vFiler.

Example: “somevfiler”

vFiler VLAN

Specifies the virtual LAN ID of the vFiler.

Example: “e0b”

Remote Filer User Name

Specifies a user name to log in to the remote storage host.

Example: “username”

Remote Filer Password

Specifies the password for the user name.

Example: “password01”

Synchronous

Indicates whether to perform a synchronous Snapmirror transfer (Snapmirror must be licensed). Available values are "true" or "false".

Default: “false”

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: “host01.domain.com”

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: “user01”

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Move Lun

Moves storage (vFiler and Lun) from one to another through Snapmirror (the NetApp data replication option).

Note: This connector uses the NetApp Disaster Recovery option. The destination vFiler and Lun must be set up and ready to accept information. Set up the Disaster Recovery Lun before running this command.

Component

spm

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: "sourcefiler01"

Source Volume

Specifies the volume or Lun.

Example: "uservol1"

Destination Filer

Specifies a storage host for vFiler and Lun.

Example: "destfiler"

Destination Volume

Specifies a volume or Lun to receive information.

Example: "uservol01"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision and Attach CIFS

Provision CIFS dataset and attach it to a physical or virtual machine.

Component

spm

Target Server ModelPath

Specifies the ModelPath of the server from CA Server Automation.

Example:
“https://localhost/aip/AOM/root/cimv2:CA_ComputerSystem.CreationClassName="CA_ComputerSystem",Name="3c16db00-405a-11e0-9207-0800200c9a66"

To obtain the target server ModelPath:

1. Navigate to Start -> Programs -> CA -> CA Server Automation - > CA Server Automation Command Prompt.
2. Enter the following command:

   ```
   C:\CA\ServerAutomation\bin>caaipaomwsclient /enumerate=CA_ComputerSystem
   /queryFilter="Select Name from CA_ComputerSystem where
   ElementName='server_name'" /user=SA_user /password=SA_password
   ```

   Replace *server_name* with the name of the server. Replace *SA_user* and *SA_password* with credentials for the CA Server Automation authorized user.
3. Use the value returned by 'Name' in the target server ModelPath.

Name

Specifies the name to use for the dataset and datastore.

Example: “mydataset-01”

Size in MB

Specifies the size of storage to create in megabytes.

Example: “5000”

Description

(Optional) Specifies a text description for the dataset.

Example: “my dataset”

Owner

(Optional) Specifies the owner of the dataset.

Example: “OwnerName”

Provisioning Policy

Specifies the NetApp policy to use when creating the dataset.

Example: “NAS_policy”

Resource Pool Name

Specifies the resource pool to use for the dataset.

Example: "lodnetapp10z"

Resource Pool ID

(Optional) Specifies the ID of the resource pool.

Example: "1234"

Group

(Optional) Specifies a group name to use to group the provisioned storage.

Example: "group01"

Container

(Optional) Specifies the name of the storage container in which to add the dataset.

Example: "stor_container01"

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: "lodnetapp40"

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: "123.456.7.89"

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: "23"

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: "0"

CIFS Domain

Specifies the name of the CIFS domain.

Example: "somedomain.com"

CIFS User

Specifies the CIFS user name.

Default: "everyone"

CIFS User Permissions

Specifies CIFS user permissions. Available values are:

0 – No Access

1 – Full Control

2 – Read

3 – Change

Default: “1”

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: “GMT-5”

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

Call Asynchronously

Indicates whether calls return immediately or wait for commands to return. When "true", calls return immediately and you must verify the JobStatus.

Default: “false”

Attach Username

(Optional) Specifies the user name to connect to the host to mount storage.

Example: "username"

Attach Password

(Optional) Specifies the user password to connect to the host to mount storage.

Example: "password"

Attach Location

(Optional) Specifies the directory on which to mount storage.

Example: "/vol/userdir/userdir01"

CIFS Attach Username

Specifies the CIFS user name authorized in the domain.

Example: “user23”

CIFS Attach Password

Specifies the password for the CIFS user name.

Example: "password"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision and Attach FCP

Provision an FCP dataset and attach it to the datastore.

Component

spm

Target Server ModelPath

Specifies the ModelPath of the server from CA Server Automation.

Example:
“https://localhost/aip/AOM/root/cimv2:CA_ComputerSystem.CreationClassName="CA_ComputerSystem",Name="3c16db00-405a-11e0-9207-0800200c9a66"

To obtain the target server ModelPath:

1. Navigate to Start -> Programs -> CA -> CA Server Automation - > CA Server Automation Command Prompt.
2. Enter the following command:

   ```
   C:\CA\ServerAutomation\bin>caaipaomwsclient /enumerate=CA_ComputerSystem
   /queryFilter="Select Name from CA_ComputerSystem where
   ElementName='server_name'" /user=SA_user /password=SA_password
   ```

   Replace *server_name* with the name of the server. Replace *SA_user* and *SA_password* with credentials for the CA Server Automation authorized user.
3. Use the value returned by 'Name' in the target server ModelPath.

Name

Specifies the name to use for the dataset and datastore.

Example: “mydataset-01”

Size in MB

Specifies the size of storage to create in megabytes.

Example: “5000”

Description

(Optional) Specifies a text description for the dataset.

Example: “my dataset”

Owner

(Optional) Specifies the owner of the dataset.

Example: “OwnerName”

Provisioning Policy

Specifies the NetApp policy to use when creating the dataset.

Example: “NAS_policy”

Resource Pool Name

Specifies the resource pool to use for the dataset.

Example: “lodnetapp10z”

Resource Pool ID

(Optional) Specifies the ID of the resource pool.

Example: “1234”

Group

(Optional) Specifies a group name to use to group the provisioned storage.

Example: “group01”

Container

(Optional) Specifies the name of the storage container in which to add the dataset.

Example: “stor_container01”

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: “lodnetapp40”

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: “123.456.7.89”

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: “23”

Initial Snapshot Size

Specifies the initial size of the snapshot in megabytes.

Default: "0"

Max Snapshot Size

Specifies the maximum size of the snapshot in megabytes.

Default: "0"

World Wide Port Host

(Optional) Specifies the host that owns the initiator.

Example: “somehost”

World Wide Port ID

(Optional) Specifies the initiator ID on a host to which to map a LUN.

Example: 00:00:00:00:00:00:00:00

World Wide Port OS

(Optional) Specifies the operating system type of the initiator host. Available values are:

0 – VMWARE

1 – AIX

2 – HP-UX

3 – Linux

4 – Solaris

5 – Netware

6 – Windows

7 – Windows_2008

8 – Windows_GPT

Default: “0”

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: “0”

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: “GMT-5”

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

Call Asynchronously

Indicates whether calls return immediately or wait for commands to return. When "true", calls return immediately and you must verify the JobStatus.

Default: “false”

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision and Attach NAS

Provision a NAS dataset and attach it to the datastore.

Component

spm

Target Server ModelPath

Specifies the ModelPath of the server from CA Server Automation.

Example:
“https://localhost/aip/AOM/root/cimv2:CA_ComputerSystem.CreationClassName="CA_ComputerSystem",Name="3c16db00-405a-11e0-9207-0800200c9a66"

To obtain the target server ModelPath:

1. Navigate to Start -> Programs -> CA -> CA Server Automation - > CA Server Automation Command Prompt.
2. Enter the following command:

   ```
   C:\CA\ServerAutomation\bin>caaipaomwsclient /enumerate=CA_ComputerSystem
   /queryFilter="Select Name from CA_ComputerSystem where
   ElementName='server_name'" /user=SA_user /password=SA_password
   ```

   Replace *server_name* with the name of the server. Replace *SA_user* and *SA_password* with credentials for the CA Server Automation authorized user.
3. Use the value returned by 'Name' in the target server ModelPath.

Name

Specifies the name to use for the dataset and datastore.

Example: “mydataset-01”

Size in MB

Specifies the size of storage to create in megabytes.

Example: “5000”

Description

(Optional) Specifies a text description for the dataset.

Example: “my dataset”

Owner

(Optional) Specifies the owner of the dataset.

Example: “OwnerName”

Provisioning Policy

Specifies the NetApp policy to use when creating the dataset.

Example: “NAS_policy”

Resource Pool Name

Specifies the resource pool to use for the dataset.

Example: “lodnetapp10z”

Resource Pool ID

(Optional) Specifies the ID of the resource pool.

Example: “1234”

Group

(Optional) Specifies a group name to use to group the provisioned storage.

Example: “group01”

Container

(Optional) Specifies the name of the storage container in which to add the dataset.

Example: “stor_container01”

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: “lodnetapp40”

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: “123.456.7.89”

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: “23”

Disable Set UID

Indicates whether to ignore attempts to enable the SETUID.

Default: “false”

User

Specifies the user for root access. If the client accessing the export is not present in the root access list for the export, the effective root user is the specified value. Default value is 65534 which maps to user "nobody". Valid values are user names not more than 255 characters or an integer ranging 0-65534.

Example: “someuser”

NFS Export Path Security

Specifies the security type supported on exports. Available values are:

0 – None

1 – Unix Style (SYS)

2 – Kerberos v5

3 – KRB5I - Kerberos v5

4 – KRB5P - Kerberos v5

Default: "0"

Max Data Size

Specifies the maximum storage space in megabytes.

Example: "4000"

Read Only

Indicates whether all hosts get read-only permissions on NFS exports. Available values are "true" and " false".

Default: "false"

Hostname

Specifies the host to grant NFS permissions.

Example: "somehost"

Host Privilege

Specifies the permission granted to the host on the NFS export path. Available values are:

0 – Read-only

1 – Write

2 – Root

Example: "2"

Host Exception

Indicates whether the host specified in the NFS host information is the exception to the host privilege specified. Available values are "true" or "false".

Example: "true"

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: "0"

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: "GMT-5"

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

Call Asynchronously

Indicates whether calls return immediately or wait for commands to return. When "true", calls return immediately and you must verify the JobStatus.

Default: "false"

NFS User Name

Specifies the user name to connect to the host to mount storage.

Example: "user01"

NFS User Password

Specifies the password for the NFS user name.

Example: "password01"

NFS Mount Point

Specifies the directory on which to mount storage.

Example: "/vol/userdir/userdir01"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision and Attach SCSI

Provision a SCSI dataset and attach it to the datastore.

Component

spm

Target Server ModelPath

Specifies the ModelPath of the server from CA Server Automation.

Example:
“https://localhost/aip/AOM/root/cimv2:CA_ComputerSystem.CreationClassName="CA_ComputerSystem",Name="3c16db00-405a-11e0-9207-0800200c9a66"

To obtain the target server ModelPath:

1. Navigate to Start -> Programs -> CA -> CA Server Automation - > CA Server Automation Command Prompt.
2. Enter the following command:

   ```
   C:\CA\ServerAutomation\bin>caaipaomwsclient /enumerate=CA_ComputerSystem
   /queryFilter="Select Name from CA_ComputerSystem where
   ElementName='server_name'" /user=SA_user /password=SA_password
   ```

   Replace *server_name* with the name of the server. Replace *SA_user* and *SA_password* with credentials for the CA Server Automation authorized user.
3. Use the value returned by 'Name' in the target server ModelPath.

Name

Specifies the name to use for the dataset and datastore.

Example: “mydataset-01”

Size in MB

Specifies the size of storage to create in megabytes.

Example: “5000”

Description

(Optional) Specifies a text description for the dataset.

Example: “my dataset”

Owner

(Optional) Specifies the owner of the dataset.

Example: “OwnerName”

Provisioning Policy

Specifies the NetApp policy to use when creating the dataset.

Example: “NAS_policy”

Resource Pool Name

Specifies the resource pool to use for the dataset.

Example: “lodnetapp10z”

Resource Pool ID

(Optional) Specifies the ID of the resource pool.

Example: “1234”

Group

(Optional) Specifies a group name to use to group the provisioned storage.

Example: “group01”

Container

(Optional) Specifies the name of the storage container in which to add the dataset.

Example: “stor_container01”

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: “lodnetapp40”

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: “123.456.7.89”

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: “23”

Initial Snapshot Size

Specifies the initial size of the snapshot in megabytes.

Default: "0"

Maximum Snapshot Size

Specifies the maximum size of the snapshot in megabytes.

Default: "0"

Initiator Host

Specifies the host that owns the initiator.

Example: “somehost.ca.com”

Initiator ID

Specifies an initiator ID on a host to which to map a LUN.

Example: "iqn.01.com.vmware:host01"

Initiator Host Type

Specifies the operating system type of the initiator host. Available values are:

0 – VMWARE

1 – AIX

2 – HP-UX

3 – Linux

4 – Solaris

5 – Netware

6 – Windows

7 – Windows_2008

8 – Windows_GPT

Default: “0”

Call Asynchronously

Indicates whether calls return immediately or wait for commands to return. When "true", calls return immediately and you must verify the JobStatus.

Default: “false”

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: “0”

Attach Username

(Optional) Specifies the user name to connect to the host to mount storage.

Example: "username"

Attach Password

(Optional) Specifies the user password to connect to the host to mount storage.

Example: "password"

Attach Location

(Optional) Specifies the directory on which to mount storage.

Example: "/vol/userdir/userdir01"

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: “GMT-5”

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: “host01.domain.com”

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: “user01”

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: “password01”

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: “8088”

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: “0”

Storage Platform Type

Specifies the storage platform type. The only possible value is “1” (NetApp).

Default: “1”

Storage Provision CIFS

Provision a new CIFS dataset.

Component

spm

Dataset Name

Specifies the name of the dataset.

Example: “mydataset-01”

Storage Size in MB

Specifies the size of storage to create in megabytes.

Example: “5000”

Storage Description

(Optional) Specifies a text description for the dataset.

Example: “This is my dataset.”

Storage Owner

(Optional) Specifies the owner of the dataset.

Example: “OwnerName”

Storage Provisioning Policy

Specifies the NetApp Policy to use during creation of the dataset.

Example: “NAS_policy”

Storage Resource Pool

Specifies the resource pool to use for the dataset.

Example: “lodnetapp10z”

Storage Resource ID

(Optional) Specifies the resource pool ID.

Example: “1234”

Storage Group

(Optional) Specifies a name used to group the provisioned storage.

Example: “group01”

Storage Container

(Optional) Specifies the name of the storage container to which to add the dataset.

Example: “stor_container01”

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: “lodnetapp40”

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: “123.456.7.89”

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: “23”

CIFS Domain

Specifies the name of the CIFS domain.

Example: “somedomain.com”

CIFS User

Specifies the CIFS user name.

Default: "everyone"

CIFS User Permissions

Specifies CIFS user permissions. Available values are:

0 – No Access

1 – Full Control

2 – Read

3 – Change

Default: “1”

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: “0”

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: “GMT-5”

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision FCP

Provision a new Fiber Channel dataset.

Component

spm

Dataset Name

Specifies the name of the dataset.

Example: “mydataset-01”

Storage Size in MB

Specifies the size of storage to create in megabytes.

Example: “5000”

Storage Description

(Optional) Specifies a text description for the dataset.

Example: “This is my dataset.”

Storage Owner

(Optional) Specifies the owner of the dataset.

Example: “OwnerName”

Storage Provisioning Policy

Specifies the NetApp Policy to use during creation of the dataset.

Example: “NAS_policy”

Storage Resource Pool

Specifies the resource pool to use for the dataset.

Example: “lodnetapp10z”

Storage Resource ID

(Optional) Specifies the resource pool ID.

Example: “1234”

Storage Group

(Optional) Specifies a name used to group the provisioned storage.

Example: “group01”

Storage Container

(Optional) Specifies the name of the storage container to which to add the dataset.

Example: “stor_container01”

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: “lodnetapp40”

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: “123.456.7.89”

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: “23”

Initiator Host

Specifies the host that owns the initiator.

Example: “somehost”

Initiator ID

Specifies an initiator ID on a host to which to map a LUN.

Example: "iqn.01.com.vmware:host01"

Initiator Host Type

Specifies the operating system type of the initiator host. Available values are:

0 – VMWARE

1 – AIX

2 – HP-UX

3 – Linux

4 – Solaris

5 – Netware

6 – Windows

7 – Windows_2008

8 – Windows_GPT

Default: “0”

Initial Snapshot Size

Specifies the initial size of the snapshot in megabytes.

Default: "0"

Max Snapshot Size

Specifies the maximum size of the snapshot in megabytes.

Default: "0"

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: "0"

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: "GMT-5"

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision MixedMode

Provision a new MixedMode dataset.

Component

spm

Dataset Name

Specifies the name of the dataset.

Example: "mydataset-01"

Storage Size in MB

Specifies the size of storage to create in megabytes.

Example: "5000"

Storage Description

(Optional) Specifies a text description for the dataset.

Example: "This is my dataset."

Storage Owner

(Optional) Specifies the owner of the dataset.

Example: "OwnerName"

Storage Provisioning Policy

Specifies the NetApp Policy to use during creation of the dataset.

Example: "NAS_policy"

Storage Resource Pool

Specifies the resource pool to use for the dataset.

Example: "lodnetapp10z"

Storage Resource ID

(Optional) Specifies the resource pool ID.

Example: "1234"

Storage Group

(Optional) Specifies a name used to group the provisioned storage.

Example: "group01"

Storage Container

(Optional) Specifies the name of the storage container to which to add the dataset.

Example: "stor_container01"

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: "lodnetapp40"

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: "123.456.7.89"

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: "23"

Hostname

Specifies the host to grant NFS permissions.

Example: "somehost"

Host Privilege

Specifies the permission granted to the host on the NFS export path. Available values are:

0 – Read-only

1 – Write

2 – Root

Example: "2"

Host Exception

Indicates whether the host specified in the NFS host information is the exception to the host privilege specified. Available values are "true" or "false".

Example: "true"

Disable Set UID

Indicates whether to ignore attempts to enable the SETUID.

Default: “false”

User Root Access

Specifies a user for root access.

Example: “someuser”

NFS Export Path Security

Specifies the security type supported on exports. Available values are:

0 – None

1 – Unix Style (SYS)

2 – Kerberos v5

3 – KRB5I - Kerberos v5

4 – KRB5P - Kerberos v5

Default: “0”

Max Data Size

Specifies the maximum storage space in megabytes.

Example: “4000”

Read Only

Indicates whether all hosts get read-only permissions on NFS exports. Available values are “true” and “ false”.

Default: “false”

Domain

Specifies the name of the MixedMode domain.

Example: “somedomain.com”

Domain User

Specifies the domain user name.

Default: "everyone"

Domain Permission

Specifies domain user permissions. Available values are:

0 – No Access

1 – Full Control

2 – Read

3 – Change

Default: "1"

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: "0"

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: "GMT-5"

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision NFS

Provision a new NFS dataset.

Component

spm

Dataset Name

Specifies the name of the dataset.

Example: "mydataset-01"

Storage Size in MB

Specifies the size of storage to create in megabytes.

Example: "5000"

Storage Description

(Optional) Specifies a text description for the dataset.

Example: "This is my dataset."

Storage Owner

(Optional) Specifies the owner of the dataset.

Example: "OwnerName"

Storage Provisioning Policy

Specifies the NetApp Policy to use during creation of the dataset.

Example: "NAS_policy"

Storage Resource Pool

Specifies the resource pool to use for the dataset.

Example: "lodnetapp10z"

Storage Resource ID

(Optional) Specifies the resource pool ID.

Example: "1234"

Storage Group

(Optional) Specifies a name used to group the provisioned storage.

Example: "group01"

Storage Container

(Optional) Specifies the name of the storage container to which to add the dataset.

Example: "stor_container01"

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: "lodnetapp40"

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: "123.456.7.89"

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: "23"

Hostname

Specifies the host to grant NFS permissions.

Example: "somehost"

Host Privilege

Specifies the permission granted to the host on the NFS export path. Available values are:

0 – Read-only

1 – Write

2 – Root

Example: "2"

Host Exception

Indicates whether the host specified in the NFS host information is the exception to the host privilege specified. Available values are “true” or “false”.

Example: “true”

Disable Set UID

Indicates whether to ignore attempts to enable the SETUID.

Default: “false”

User

Specifies the user for root access. If the client accessing the export is not present in the root access list for the export, the effective root user is the specified value. Default value is 65534 which maps to user "nobody". Valid values are user names not more than 255 characters or an integer ranging 0-65534.

Example: “someuser”

NFS Export Path Security

Specifies the security type supported on exports. Available values are:

0 – None

1 – Unix Style (SYS)

2 – Kerberos v5

3 – KRB5I - Kerberos v5

4 – KRB5P - Kerberos v5

Default: “0”

Max Data Size

Specifies the maximum storage space in megabytes.

Example: “4000”

Read Only

Indicates whether all hosts get read-only permissions on NFS exports. Available values are “true” and “ false”.

Default: “false”

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: "0"

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: "GMT-5"

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Provision SCSI

Provision a new SCSI dataset.

Component

spm

Dataset Name

Specifies the name of the dataset.

Example: "mydataset-01"

Storage Size in MB

Specifies the size of storage to create in megabytes.

Example: "5000"

Storage Description

(Optional) Specifies a text description for the dataset.

Example: "This is my dataset."

Storage Owner

(Optional) Specifies the owner of the dataset.

Example: "OwnerName"

Storage Provisioning Policy

Specifies the NetApp Policy to use during creation of the dataset.

Example: "NAS_policy"

Storage Resource Pool

Specifies the resource pool to use for the dataset.

Example: "lodnetapp10z"

Storage Resource ID

(Optional) Specifies the resource pool ID.

Example: “1234”

Storage Group

(Optional) Specifies a name used to group the provisioned storage.

Example: “group01”

Storage Container

(Optional) Specifies the name of the storage container to which to add the dataset.

Example: “stor_container01”

vFiler Name

(Optional) Specifies the name of the vFiler.

Example: “lodnetapp40”

vFiler IP Address

(Optional) Specifies the IP address of the vFiler.

Example: “123.456.7.89”

vFiler Network Mask

(Optional) Specifies the network mask of the vFiler.

Example: “23”

Initiator Host

Specifies the host that owns the initiator.

Example: “somehost.ca.com”

Initiator ID

Specifies an initiator ID on a host to which to map a LUN.

Example: "iqn.01.com.vmware:host01"

Initiator Host Type

Specifies the operating system type of the initiator host. Available values are:

0 – VMWARE

1 – AIX

2 – HP-UX

3 – Linux

4 – Solaris

5 – Netware

6 – Windows

7 – Windows_2008

8 – Windows_GPT

Default: “0”

Initial Snapshot Size

Specifies the initial size of the snapshot in megabytes.

Default: "0"

Maximum Snapshot Size

Specifies the maximum size of the snapshot in megabytes.

Default: "0"

Dry Run Mode

Specifies whether to perform a test of provisioning. Available values are:

0 – None (performs the provision without performing Dry Run)

1 – Dry Run Only (pseudo provisioning (failure prediction) that verifies the specified parameters)

2 – Provision with Dry Run (performs pseudo provisioning and provisioning)

Default: “0”

Time Zone

(Optional) Specifies a time zone for the dataset.

Example: “GMT-5”

Contact

(Optional) Specifies email addresses for contact information. Separate multiple email addresses with commas.

Example: user1@company.com, user2@company.com

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Remove Datastore

Removes an existing datastore.

Component

vc

vCenter Name

Specifies the fully qualified domain name (FQDN) of the VMware vCenter server.

Example: "vpm-vc01.domain.com"

ESX Name

Specifies the fully qualified domain name (FQDN) of the ESX Server.

Example: "esxhost.com"

Datastore Index

(Optional) Specifies the name of the datastore index.

Example: “datastore-0001”

To obtain the datastore index:

1. Navigate to Start -> Programs -> CA -> CA Server Automation - > CA Server Automation Command Prompt.
2. Enter the following command:

   ```
   C:\CA\ServerAutomation\bin>caaipaomwsclient /enumerate=CA_Datastore
   /queryFilter="Select * from CA_Datastore where ElementName='datastore'"
   /user=SA_user /password=SA_password
   ```

   Replace *datastore* with the name of the datastore, replace *SA_user* and *SA_password* with credentials for a CA Server Automation authorized user.
3. Use the value returned by 'Index' as the datastore index.

Storage Rescan Host HBA

Rescans all host HBAs or a specific host HBA.

Component

vc

vCenter Name

Specifies the fully qualified domain name (FQDN) of the VMware vCenter server.

Example: “vpm-vc01.domain.com”

ESX Name

Specifies the fully qualified domain name (FQDN) of the ESX Server.

Example: “esxhost.com”

Device HBA Name

(Optional) Specifies the device HBA (Host Bus Adapter) name.

Example: “devicehba01”

Storage Resize

Resizes an existing dataset.

Component

spm

Dataset Name

Specifies the name of the dataset.

Example: “mydataset-01”

New Size

Specifies the new size of the dataset.

Example: “15000”

Maximum Capacity

Specifies the new maximum capacity value for a flexible volume.

Example: ”2000”

Snapshot Reserve Percentage

Specifies the percentage of volume space reserved for snapshot copies.

Default: “0%”

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: “host01.domain.com”

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: “user01”

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: “password01”

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: “8088”

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage vFiler Active

Activates the vFiler.

Note: Set up the Disaster Recovery vFiler before running this command.

Component

spm

Destination Filer

Specifies a storage host for vFiler and Lun.

Example: "destfiler"

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: "sourcefiler01"

vFiler Name

Specifies the name of the vFiler.

Example: "somevfiler"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage vFiler Resync

Resyncs the vFiler.

Note: Set up the Disaster Recovery vFiler before running this command.

Component

spm

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: "sourcefiler01"

Destination Filer

Specifies a storage host for vFiler and Lun.

Example: "destfiler"

vFiler Name

Specifies the name of the vFiler.

Example: "somevfiler"

vFiler VLAN

Specifies the virtual LAN ID of the vFiler.

Example: “e0b”

Remote Filer User

Specifies a user name to log in to the remote storage host.

Example: “someuser”

Remote Filer Password

Specifies the password for the user name.

Example: “password01”

Synchronous

Indicates whether to perform a synchronous Snapmirror transfer (Snapmirror must be licensed). Available values are "true" or "false".

Default: “false”

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: “host01.domain.com”

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: “user01”

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: “password01”

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: “8088”

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage vFiler Status

Gets the status of the resync vFiler action.

Note: Set up the Disaster Recovery Lun and vFiler before running this command.

Component

spm

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: "sourcefiler01"

Destination Filer

Specifies a storage host for vFiler and Lun.

Example: "destfiler"

vFiler Name

Specifies the name of the vFiler.

Example: "somevfiler"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage vFiler Stop

Stops a vFiler.

Note: Set up the Disaster Recovery Lun and vFiler before running this command.

Component

spm

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: "sourcefiler01"

vFiler Name

Specifies the name of the vFiler.

Example: "somevfiler"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage vLan Interface

Sets the network interface for vLan.

Note: Set up the Disaster Recovery Lun and vFiler before running this command.

Component

spm

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: "sourcefiler01"

vFiler VLAN

Specifies the virtual LAN ID of the vFiler.

Example: "e0b"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Storage Volume Offline

Takes the volume offline.

Note: Set up the Disaster Recovery Lun and vFiler before running this command.

Component

spm

Source Volume

Specifies the volume or Lun.

Example: "uservol1"

Source Filer

Specifies the original storage host containing vFiler and Lun.

Example: "sourcefiler01"

vFiler Name

Specifies the name of the vFiler.

Example: "somevfiler"

NetApp Server Name

(Optional) Specifies the name of the NetApp DataFabric Manager.

Example: "host01.domain.com"

NetApp Server Username

(Optional) Specifies the name of the user authorized to connect to the NetApp DataFabric Manager.

Example: "user01"

NetApp Server User Password

(Optional) Specifies the password for the user authorized to connect to the NetApp DataFabric Manager.

Example: "password01"

NetApp Server Port

(Optional) Specifies the port number the NetApp DataFabric Manager uses.

Example: "8088"

NetApp Server Protocol

(Optional) Specifies the protocol used to access the NetApp DataFabric Manager. Available values are:

0 – HTTP

1 – HTTPS

Default: "0"

Storage Platform Type

Specifies the storage platform type. The only possible value is "1" (NetApp).

Default: "1"

Event Forwarding

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

Configure Windows for SNMP

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

To configure Windows for SNMP

1. Click Start, Control Panel, Administrative Tools, Services.

   The Services dialog appears.

2. Perform *one* of the following:

   - If SNMP Service and SNMP Trap Service are listed, continue the Event Forwarding configuration.
   - If SNMP Service and SNMP Trap Service are *not* listed, continue with step 3.

3. Click Start, Control Panel, Add or Remove Programs, and Click Add/Remove Windows Components.

   The Windows Components Wizard dialog appears.

4. Scroll down and select Management and Monitoring Tools, and click Next.

   You are prompted for the Windows installation media.

5. Follow the on-screen instructions to complete the installation.

6. Repeat step 1 to verify that SNMP Service and SNMP Trap Service are running.

Configure SNMPv1 Traps by Editing the sysedge.cf File

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

To configure SNMPv1 traps by editing the sysedge.cf file

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

   The trap destination section contains a brief description of trap destinations and communities.
3. Add a line at the end of the trap destination section for each management system to which you want to send SNMPv1 traps. Use the following syntax:

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

   community-name

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

   IP-address

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

   Default: 127.0.0.1

   hostname

   (Optional) Specifies the name of the target system.

   Default: localhost

   port-number

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

   Default: 162
4. Save and close the file.
5. Do one of the following:
   - Restart SystemEDGE to activate the changes by restarting the SystemEDGE service from the Windows service control.
   - Restart SystemEDGE to activate the changes by running the following commands from a Windows command prompt.

     ```
     net stop sysedge
     net start sysedge
     ```
   - Restart SystemEDGE to activate the changes by running the following commands from a UNIX or Linux terminal window.

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

Example

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

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

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

Configure CA Server Automation to Forward Events

Configure CA Server Automation to forward events to a CA or third-party Event Management server.

To configure CA Server Automation to forward events

1. Open the CA Server Automation user interface.
2. Click Administration.

   The Administration page appears.
3. Click Configuration.

   The Configuration page appears.
4. Click Event in the left pane.

   The Event pane appears.
5. Click + (Add).

   The Forwarding and Type fields are automatically populated.

   Note: If these fields do not populate, restart Apache Tomcat.
6. Select SNMP v1 or v3 from the Type drop-down list, then select Enabled from the Forwarding drop-down list.

   If you select SNMP v3 Trap, a panel for additional configuration parameters appears.
7. Enter the management server name in the Server field.
8. Enter a different port number for SNMP or leave the default port 162, which is automatically populated.
   - If you selected SNMP v3, continue with step 9.
   - If you selected SNMP v1, go to step 13.

9. Select Security Level from the drop-down list. The following types are valid options:

   NoAuthNoPriv

   No additional security parameters.

   AuthNoPriv

   Authentication settings only.

   AuthPriv

   Authentication and encryption.

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

    MD5

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

    SHA

    Defines the Secure Hash algorithm cryptographic hash function.

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

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

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

13. Click OK.

    A confirmation message appears.

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

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

SNMP V3 Engine ID

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

For CA Server Automation, the hexadecimal engine ID is built using a combination of the computer name and the string DCAMTrap as the seed. For example, if the computer name is COMP999, the seed of the engine ID is COMP999-DCAMTrap. CA Server Automation uses the seed in an algorithm to compute the engine ID. The computer name is used to help ensure uniqueness. The hexadecimal value is written to an output file under the CA Server Automation installation directory. The name of the file is dem_snmp3_engine_id.dat. The file is created after you configure CA Server Automation to forward events to another management server.

Configure SNMP Management Servers

Configure your management server to receive SNMP traps:

CA NSM

For information about how to configure CA NSM to receive SNMP traps, see the *CA NSM Inside Event Management and Alert Management* guide.

At a minimum, copy the CaIndication.mib file from the CA Server Automation\samples directory to the CA NSM manager host computer. You can also import the file using the CA NSM Trap Manager feature.

CA SPECTRUM

For information about how to configure CA SPECTRUM to receive SNMP traps, see the *CA SPECTRUM SNMPv3 User Guide*.

Chapter 5: Managing Resources

This section contains the following topics:

Discover Resources

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

Note: CA Server Automation discovery requires hostname resolution. If the IP address for a discovered server changes, CA Server Automation 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 Using CA Configuration Automation

You can specify a single system for discovery, management, and optionally assignment to a service or a CA Configuration Automation server. You can have multiple CCA servers in your environment and decide which CCA server manages each individual server. Each system in CA Server Automation can only have one CA Configuration Automation source. At least one CCA server must be designated as the default CCA server.

Note: CA Server Automation discovery requires hostname resolution. If the IP address for a discovered server changes, CA Server Automation 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.

To discover a system

1. Click Resources, Manage, Discover System.

2. Select System in the Discovery Type drop-down list.
3. Specify the server name or IP address in the System Name field.
4. (Optional) Using the Add to Service list, select a service to add the discovered system to.
5. Click Next.

   The Enhanced Discovery and SNMP Information page appears.
6. (Optional) Select Enhanced Discovery or Override SNMP Defaults, and supply the required credentials.
7. Click Next.

   The Assignment section appears.
8. Select a server from the CCA Server column to assign to the discovered system, and click Finish.

   All discovered servers are assigned to the specified CCA server. When you select the Assignment by OS Type option, the OS Types column is disabled.

   Discovered servers are managed automatically, but servers in a subnet discovery are not managed.

Discover a Network

You can specify a segment of networks to discover. You can have multiple CCA servers in your environment and decide which CCA server manages each individual server. Each system in CA Server Automation can only have one CA Configuration Automation source. At least one CCA server must be designated as the default CCA server.

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. Click Next.

   The Assignment section appears.

5. Select a server from the CCA Server column to assign to the discovered network, then click Finish.

   All discovered servers are assigned to this CCA server. When you select the Assignment by OS Type option, the OS Types column is disabled.

   Discovered servers are automatically managed, but servers in a subnet discovery are not managed.

Managed Virtual Environments

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

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

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

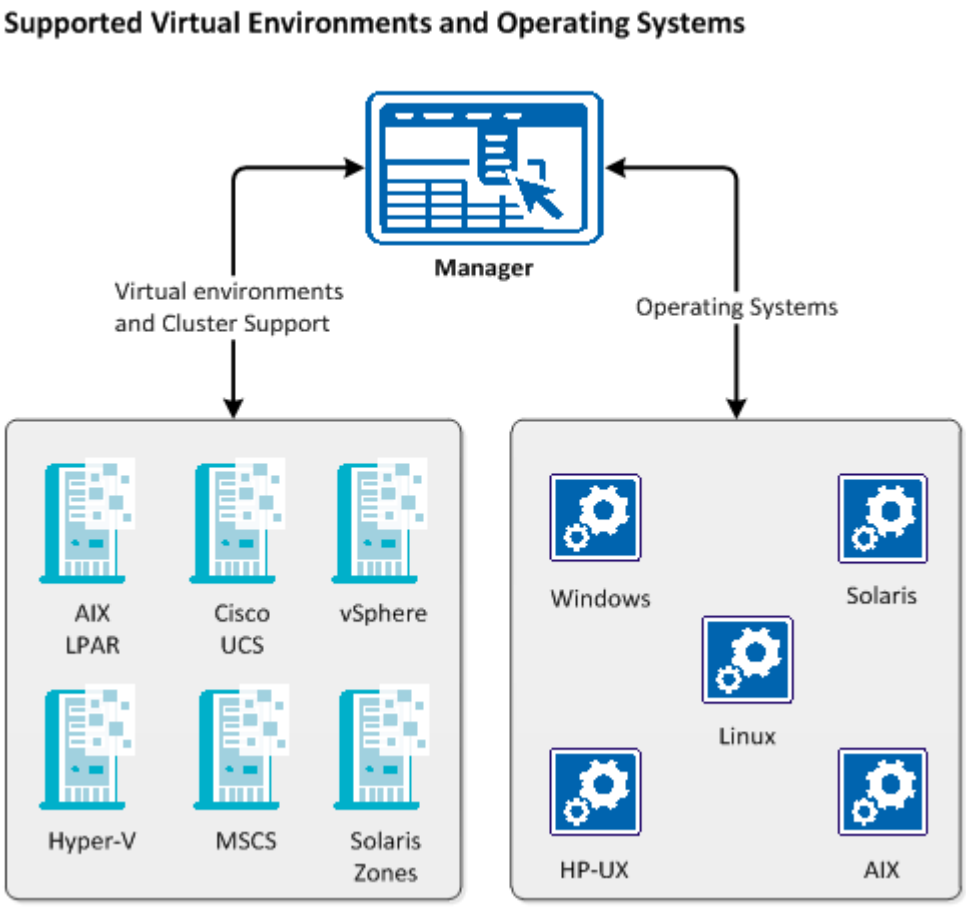

More Information

VMware vCenter Server (see page 269)
Solaris Zones (see page 284)
Microsoft Hyper-V Server (see page 286)
IBM PowerVM (LPAR) (see page 288)
Microsoft Cluster Service (see page 292)
Cisco UCS (see page 295)

VMware vCenter Server

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

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

Interactions Between vCenter Server Management Components

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

Interaction Between vCenter Server Management Components

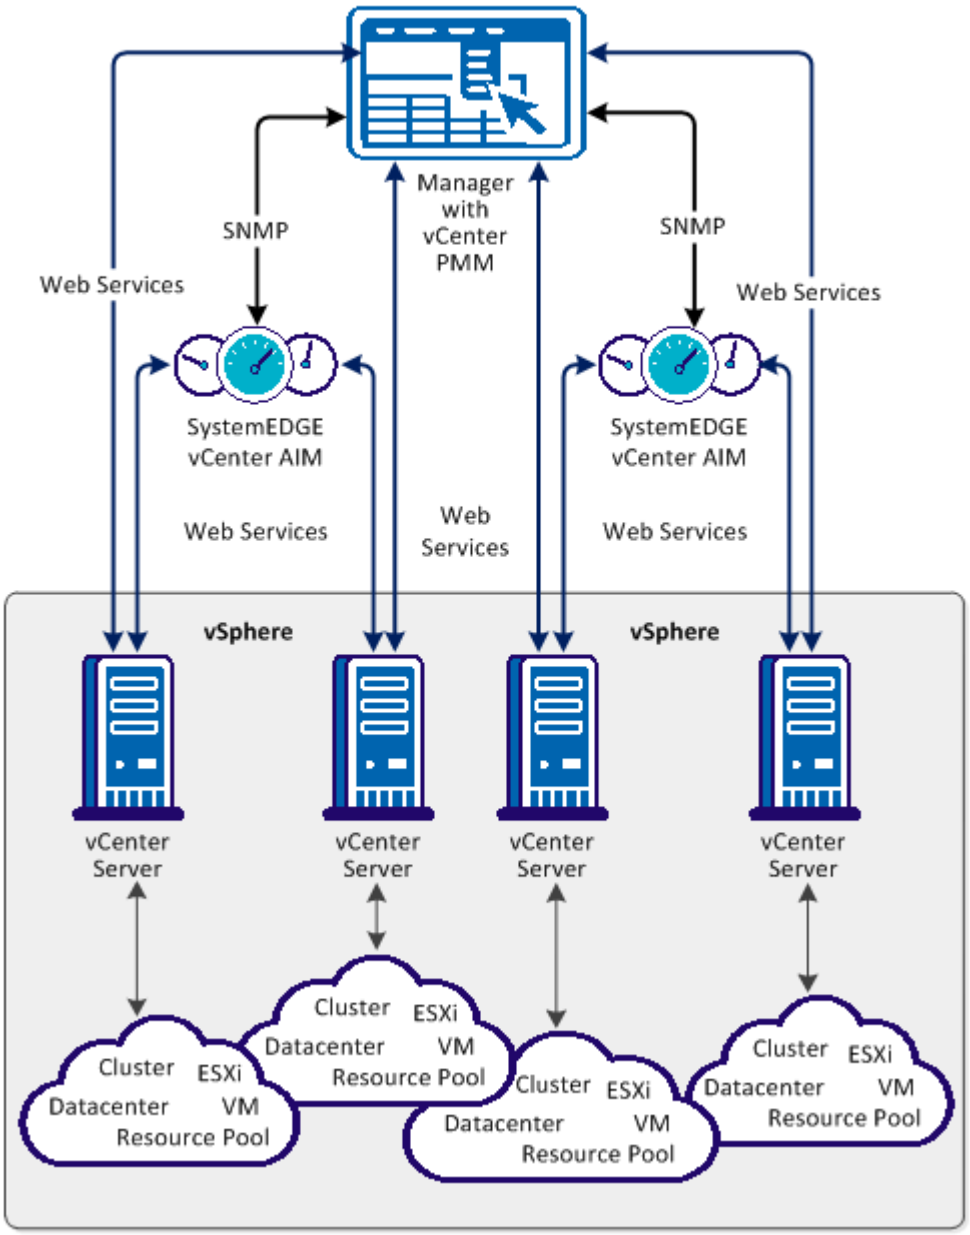

You can configure vCenter Server management during installation or through the Administration tab of the user interface.

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

More Information

Manage Multiple vCenter AIMs (see page 46)

Remote and Multi-instance vCenter Server Support

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

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

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

Note: When you run the vCenter Server AIM without the CA Server Automation manager in an eHealth, [set the sp variable for your book], or CA NSM environment, the AIM supports single-instance mode only.

Hot-plug Support for VMs

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

- Adding vCPU
- Adding vRAM

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

Device Management for VMs

Device management includes the following:

- Adding and removing vDisks
- Adding and removing vNICs

vNetwork Standard Switches (vSwitch)

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

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

Distributed Virtual Switches (DVS)

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

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

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

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

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

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

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

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

Monitor Distributed Virtual Switches Through Events

You can monitor Distributed Virtual Switches through the following events:

- Add switch:

  Distributed Virtual Switch VM-dvSwitch added to Datacenter MyDC. vSphere: vcserver.mycomp.com

- Delete switch:

  Distributed Virtual Switch VM-dvSwitch removed from Datacenter MyDC. vSphere: vcserver.mycomp.com

- Add Port Group:

  Distributed Virtual Port Group VM dvPortGroup added to Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com

- Remove Port Group:

  Distributed Virtual Port Group VM dvPortGroup removed from Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com

- Add Uplink:

  Distributed Virtual Uplink VM DVUplink added to Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com

- Remove Uplink:

  Distributed Virtual Uplink VM DVUplink removed from Distributed Virtual Switch VM-dvSwitch. Datacenter: MyDC, vSphere: vcserver.mycomp.com

To monitor Distributed Virtual Switches through events

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

   The Filter panel opens.

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

   The Events panel lists the filtered events.

vApp Support

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

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

CA Server Automation supports the following operations on vApps:

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

Monitor vApps Through Events

You can monitor vApps through the following events:

- Add vApp:

  vApp *MyvApp* added to parent resource pool resources. vSphere *vcserver.mycomp.com*

- Delete vApp:

  vApp *MyvApp* removed from parent resource pool resources. vSphere, vcserver.mycomp.com

The following traps are available:

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

To monitor vApps through events

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

   The Filter panel opens.

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

   The Events panel lists the filtered events.

vCenter Server in a Cluster

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

Fault Tolerance

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

Regarding VM management, CA Server Automation 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 Server Automation gathers FT VM data on various levels in the environment.

Fault Tolerance Requirements

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

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

VM Fault Tolerance Properties

For each VM CA Server Automation displays:

Fault Tolerance Status

Indicates the VM fault tolerance status.

Not Fault Tolerant

Indicates that the VM is not fault tolerant.

Protected

Indicates that the VM is fault tolerant and protected.

Not Protected (Starting)

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

Not Protected (Need Secondary VM)

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

Not Protected (Disabled)

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

Not Protected (VM Not Running)

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

Secondary VM Location

Identifies the secondary host location.

ESX Host Fault Tolerance Attributes

ESX Host Fault Tolerance attributes are as follows:

Fault Tolerance

Identifies whether the host has the fault tolerance enabled.

Fault Tolerance version

Identifies the version of Fault Tolerance running on the host.

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

Total Primary VMs (calculated by the AIM)

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

Total Secondary VMs (calculated by the AIM)

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

Powered on Primary VMs (calculated by the AIM)

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

Powered on Secondary VMs (calculated by the AIM)

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

Logical Volume Management in VMs

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

VMware vCenter Server Architecture

The following major components comprise the VMware vCenter Server architecture:

VMware vCenter Server

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

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

Database

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

VMware vCenter Server Agent

Connects ESX Servers with a vCenter Server.

Virtual Infrastructure Client

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

Virtual Infrastructure Web Access

Connects with the vCenter Server without installing a client.

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

The AIM for vCenter Server monitors the following resource types:

vCenter Server

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

Data Center

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

ESX Host

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

Resource Pool

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

Virtual Machine

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

Datastore

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

Virtual Disk

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

Virtual Switch

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

Physical NIC

Specifies a physical Ethernet adapter on an ESX Server.

Virtual NICs

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

Hardware Sensors

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

Resource Allocation

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

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

Resource Allocation Shares

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

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

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

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

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

Resource Allocation Reservation

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

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

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

Resource Allocation Limit

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

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

Resource Allocation Best Practices

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

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

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

Solaris Zones

A Solaris Zone defines a virtualized operating system that provides an isolated, secure environment in which to run applications. This environment allows allocation of resources among applications and services, and helps ensure that processes do not affect other zones. Solaris manages each zone as one entity. A *container* is a zone that also uses the resource management of the operating system. The Solaris Zones PMM provides health monitoring, management, and provisioning of Solaris Zones environments.

Solaris Zones Container resources can be managed at three levels:

Solaris Zones Zone Management

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

Solaris Zones Project Management

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

Solaris Zones Resource Pool Management

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

Interaction Between Solaris Zones Management Components

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

Interaction Between Solaris Zones Management Components

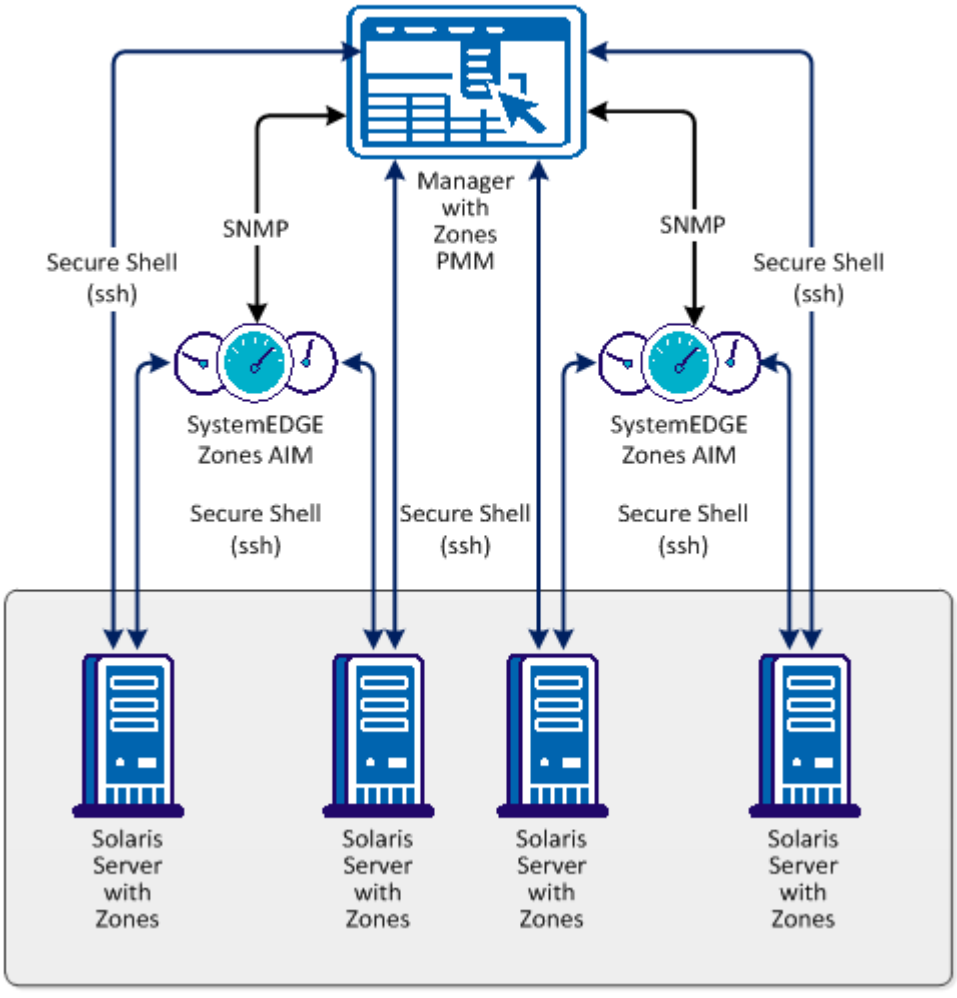

To add the required connection information for each Solaris Zones Server, use a custom installation, the Administration tab of the user interface, or the NodeCfgUtil.exe utility on the managed node. The connection information is written to the configuration file on the managed node. The AIM polls the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

Microsoft Hyper-V Server

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

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

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

Hyper-V Server

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

Virtual Machine

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

Interactions Between Hyper-V Server Management Components

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

Interaction Between Hyper-V Server Management Components

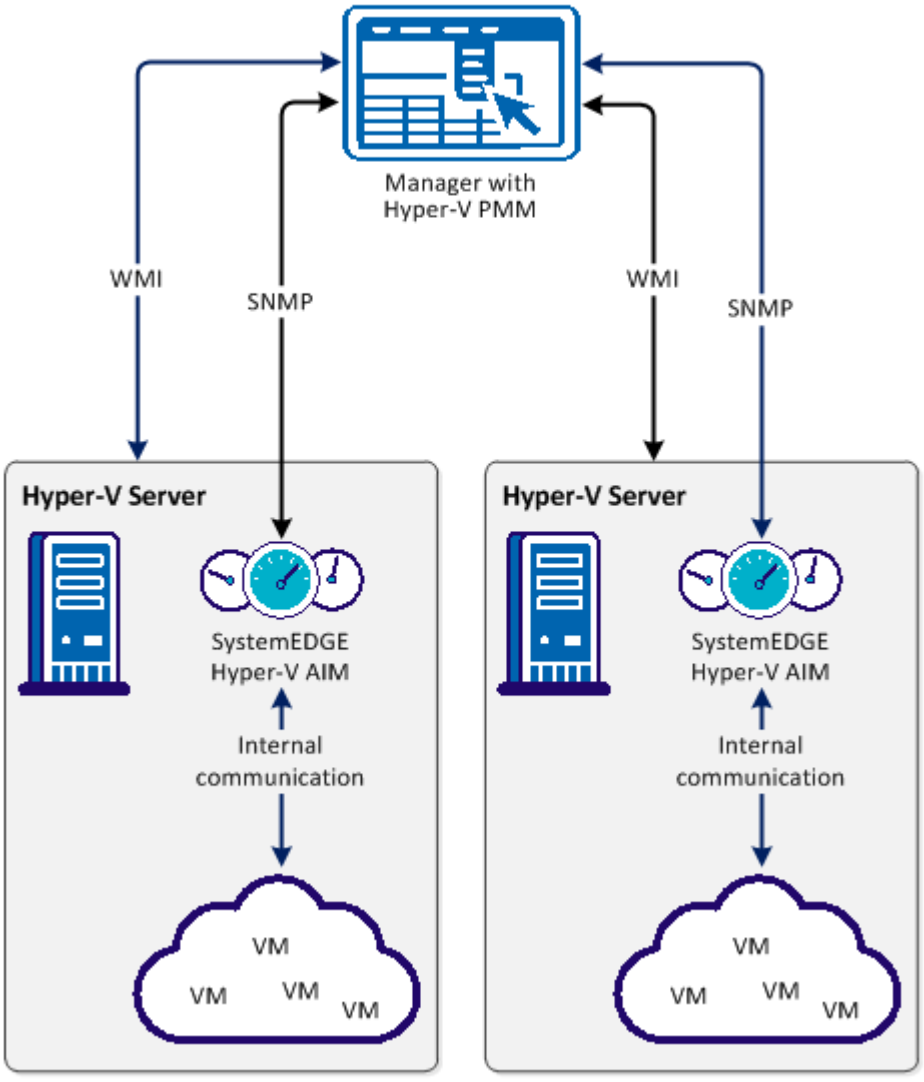

You can configure Hyper-V management during installation or through the Administration tab of the user interface.

IBM PowerVM (LPAR)

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

The LPAR AIM enables SystemEDGE to monitor PowerVM resources.

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

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

Server level

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

- Provision LPARs
- Delete LPARs

Power operations level

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

- Activate LPARs
- Shutdown LPARs
- Restart LPARs

Resource adjustments level

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

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

Interaction Between AIX LPAR Management Components

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

Interaction Between PowerVM Management Components

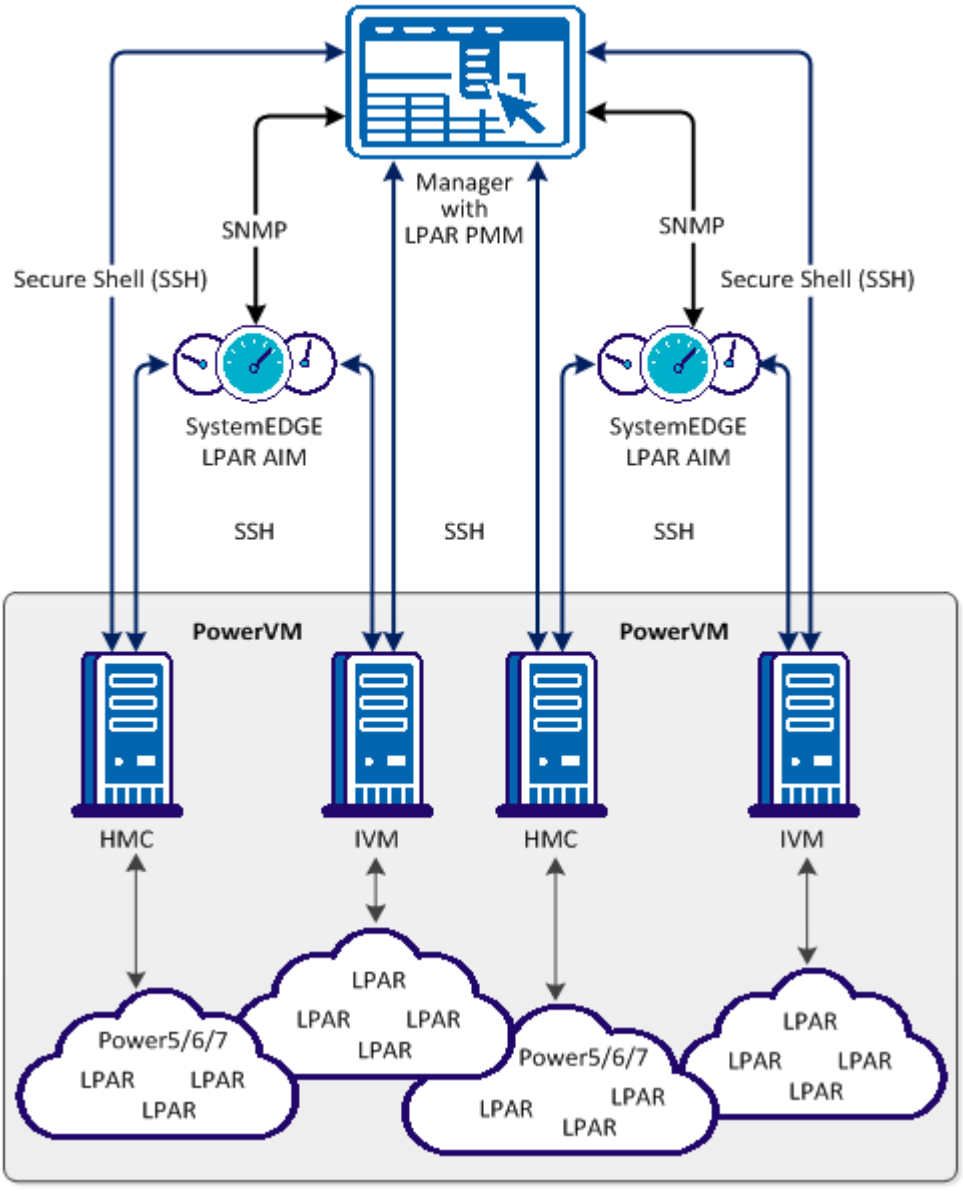

To add the required connection information for each required HMC/IVM and Virtual I/O Server, use one of the following methods:

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

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

LPAR Monitoring

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

Example

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

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

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

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

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

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

Microsoft Cluster Service

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

Interaction Between MSCS Management Components

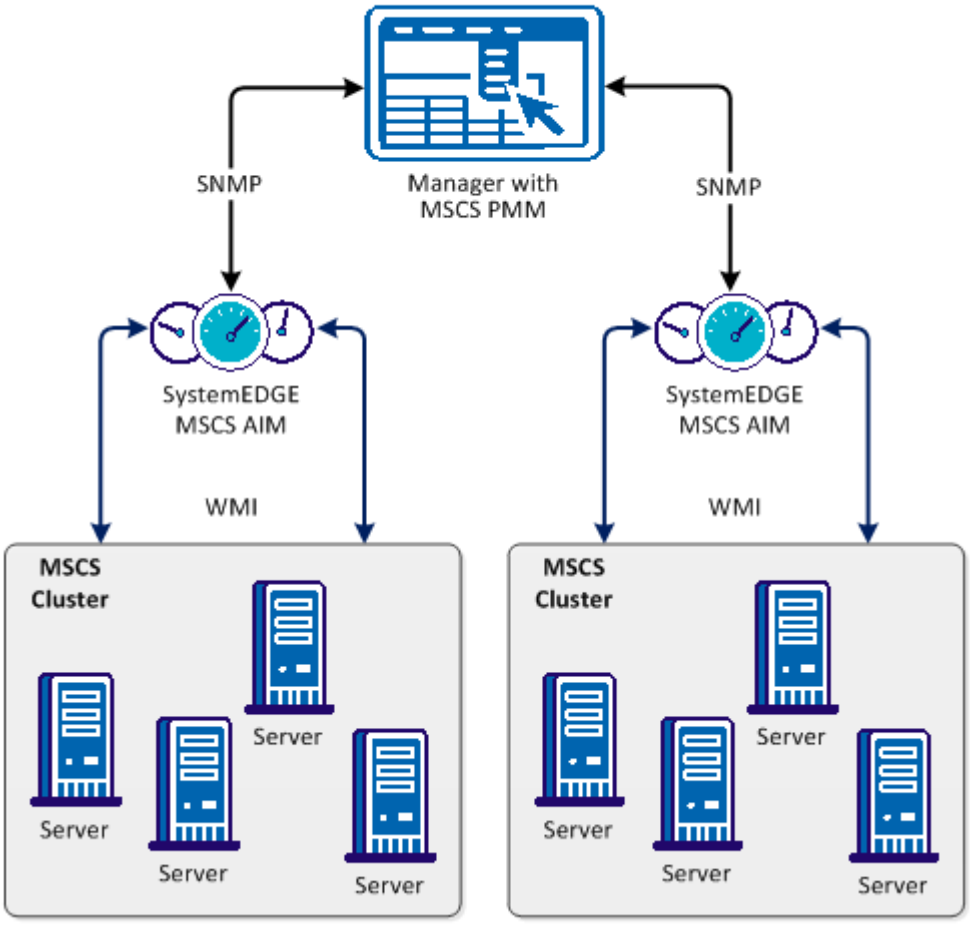

Performance monitoring requires remote access to clusters and individual cluster nodes for metric collection such as CPU and memory use. The cluster-specific information is available on each node. The MSCS AIM uses WMI (port 135) to communicate with clusters.

Microsoft Cluster Overview

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

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

More Information

Register a Cluster (see page 293)
Remove a Cluster (see page 294)
Modify Cluster Properties (see page 294)
Register a Cluster (see page 293)
Remove a Cluster (see page 294)
Modify Cluster Properties (see page 294)

Register a Cluster

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

To register a Microsoft cluster from the user interface

1. Click Administration.

   The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Microsoft Clusters.

   The Microsoft Clusters section appears on the right.
3. Click + (Add) on the Registered Microsoft Clusters toolbar.

   The Register New Cluster dialog appears.
4. Enter the required cluster name and access identification information, and click OK.

   The Microsoft cluster is registered.

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

Remove a Cluster

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

To remove a cluster

1. Click Administration.

   The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Microsoft Clusters.

   The Microsoft Clusters section appears on the right.
3. Select the cluster that you want to remove.
4. Click - (Delete) on the Registered Microsoft Clusters toolbar.

   A confirmation prompt appears.
5. Click OK.

   The cluster is removed.

Modify Cluster Properties

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

To modify cluster properties

1. Click Administration.

   The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Microsoft Clusters.

   The Microsoft Clusters section appears on the right.
3. Select the cluster that you want to edit.
4. Click the Edit icon on the Registered Microsoft Clusters toolbar.

   The Modify Cluster Properties dialog appears.
5. Edit the required properties and click OK.

   The cluster properties are modified.

Cisco UCS

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

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

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

More information:

Interaction Between Cisco UCS Management Components (see page 296)

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 Server Automation to gain access to the hardware, statistics, pools (UUID, MAC, WWPN, WWNN), and service profiles information provided by the UCS Manager.

Interaction Between Cisco UCS Management Components

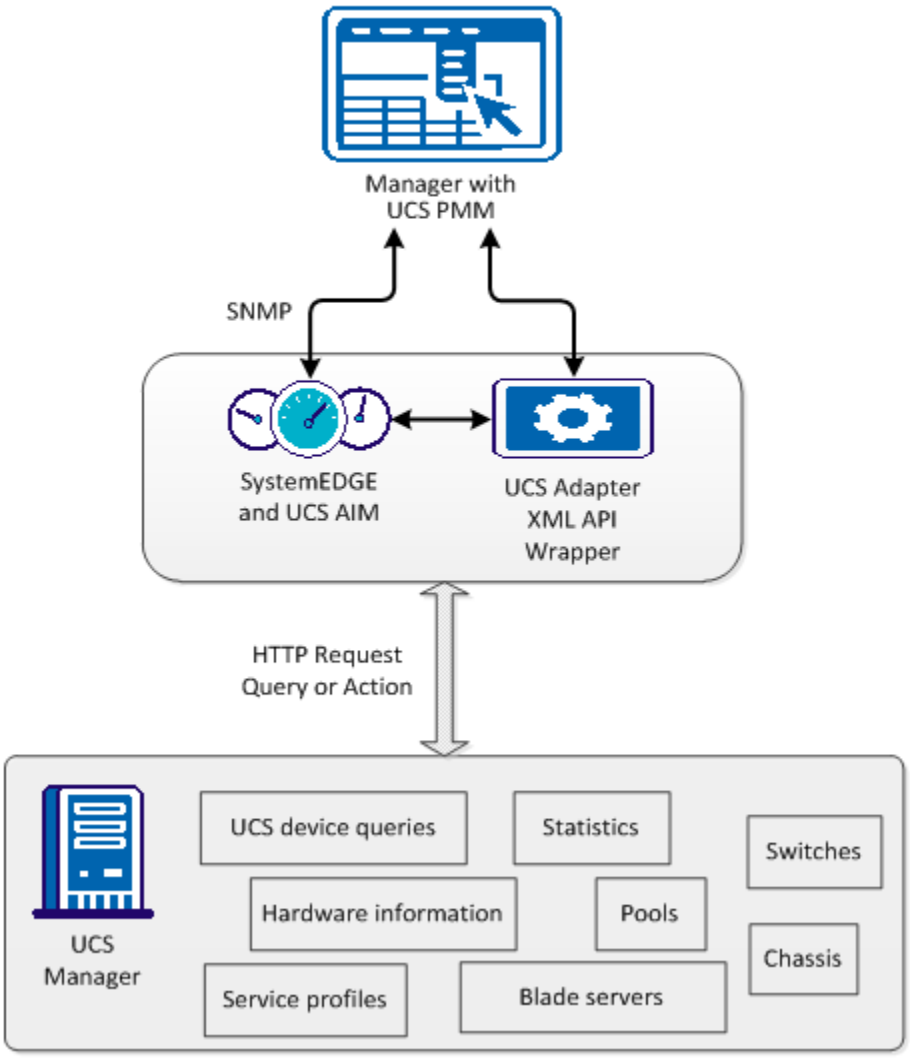

The XML API also provides the ability to configure certain device properties and perform pools and service profile management. Pools and Service Profile Management are one of the use cases that CA Server Automation manages across multiple UCS Managers to detect pool range conflicts.

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

Cisco UCS Overview

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

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

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

More information:

Add a Cisco UCS Manager Server (see page 298)
Register a UCS AIM Server (see page 298)
Remove a UCS Server (see page 299)
Configure the SNMP Data Poller (see page 299)
Configure the Service Poller (see page 300)
UCS Trap Management (see page 300)
Service Profiles (see page 301)
How To Export or Import Service Profiles (see page 301)
How To Use Centralized Service Profiles (see page 302)
Port Profiles (see page 303)
Configure the SNMP Data Poller (see page 299)
Configure the Service Poller (see page 300)
UCS Trap Management (see page 300)
Service Profiles (see page 301)
How to Create or Update a Service Profile (see page 361)

Add a Cisco UCS Manager Server

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

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

To add a UCS Manager from the user interface

1. Click Administration.

   The Administration page appears.
2. Click Configuration.

   The Configuration page appears.
3. Click Cisco UCS Servers.

   The Cisco UCS Servers page appears.
4. Click + (Add) on the Cisco UCS Servers toolbar.

   The Add Cisco UCS Server dialog appears.
5. Enter the required server identification information, and click OK.

   The server is added.

Register a UCS AIM Server

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

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

To register a UCS AIM server from the user interface

1. Click Administration.

   The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Cisco UCS Servers.

   The UCS AIM Servers section appears on the right.
3. Click + (Add) on the UCS AIM Servers toolbar.

   The Add Cisco UCS AIM Server dialog appears.
4. Enter the required server and SNMP access identification information, and click OK.

   The UCS AIM server is registered.

Remove a UCS Server

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

To remove a UCS server

1. Click Administration.

   The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Cisco UCS Servers.

   The Cisco UCS Servers page appears.
3. Select the server that you want to remove.
4. Click - (Delete) on the server toolbar.

   A confirmation prompt appears.
5. Click OK.

   The server is removed.

Configure the SNMP Data Poller

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

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

To set the polling interval

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

   <property name="CONFIG_KEY_UCS_AIM_POLL_INTERVAL">

   <!-- UCS AIM polling interval -->

   <value>300</value>

   <displayName>UCS AIM Polling Interval</displayName>

   </property>
2. Save the file.

Configure the Service Poller

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

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

The default service polling interval is 300 seconds.

To reset the polling interval

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

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

2. Save the file.

UCS Trap Management

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

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

Service Profiles

A *service profile* contains configuration information about Cisco UCS hardware, including interfaces, fabric connectivity, and network and server identity. You can create a service profile for a specific UCS manager or centrally in the CA Server Automation 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 Export or Import Service Profiles

You can clone a service profile to different UCS domains for use in additional provisioning operations.

On import, specify the service profiles to import and whether to update the profiles during the import process.

On export, specify the service profiles to export, whether to update any during the transfer process, and the target UCS domains.

Example: Export a Service Profile to a New Domain

You have a multiple UCS domain configuration in which individual UCS systems are configured with shared storage. Select a UCS system and service profile. Then, associate the profile to a blade to which an application image is provisioned and begins to run.

Later, you want to transfer the application workload to another UCS domain within the shared storage environment.

Suspend the application, disassociate the service profile from its blade, and export the service profile from the source UCS system to the target UCS system. In the new system, the service profile is associated with another blade, and the application is resumed on the target system.

How To Use Centralized Service Profiles

Central service profiles that reside in the CA Server Automation Management DB provide an efficient way to manage configuration information across multiple UCS domains. Use the CA Server Automation 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.

Port Profiles

A Cisco UCS *port profile* includes properties and settings for configuring virtual interfaces for Cisco UCS VN-Link in hardware. The VMware vCenter environment uses the term *port groups*.

The Cisco UCS Manager requires a configured *port profile client* to push a port profile to a VMware vCenter with which it is connected. A port profile client associates port profiles with one or more distributed virtual switches. When a distributed virtual switch is enabled, its associated port profiles are pushed to vCenter automatically.

CA Server Automation lets you create the necessary pieces for Cisco UCS to finish the integration with VMware vCenter. To import a port profile, use the vCenter Layout dialog to define the following VMware vCenter network topology:

- Datacenter
- Network folder
- Distributed virtual switch (DVS)
- Profile client

Use the vCenter Creation wizard to create, update, and delete items in the port profile network topology.

To manage port profiles, right-click a Cisco UCS Manager in the Explore tree and select VMware to launch the vCenter Layout dialog.

AppLogic Overview

CA AppLogic is a turnkey cloud computing platform for composing, running, and scaling distributed applications. AppLogic uses advanced virtualization technologies to be compatible with existing operating systems, middleware, and web applications. AppLogic operates on the logical structure of the application, enabling you to package an entire N-tier application into a logical entity and manage it as a single system. This approach also makes it easy to assemble, deploy, monitor, control, and troubleshoot applications visually in a browser.

Each application on AppLogic includes everything needed run on a grid of commodity servers. All infrastructure components such as fire walls, load balancers, network configurations, and database servers are combined with the application code and data forming a “single entity”. Advanced monitoring and metering tools, and prepackaged operational procedures are also combined with the application. Binding software to hardware enables applications running on AppLogic to be replicated on demand, on the same grid or in multiple locations, without any code modifications.

CA Server Automation integrates with AppLogic to provide the automated, policy-based provisioning capabilities of CA Server Automation in AppLogic environments.

Administrators can add AppLogic servers and discover grids using either the Administration user interface or the dpmutil CLI command.

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

Add AppLogic Servers

To enable the management of AppLogic grids and applications in CA Server Automation, add the AppLogic Web Services API server that manages the grids.

Follow these steps:

1. View the Administration tab, and select AppLogic Server from the Provisioning menu.

   The AppLogic Server page opens showing the currently configured servers.

2. Click + (Add) in the AppLogic Server toolbar.

   The New AppLogic Server panel opens.

3. Enter the server connection parameters, select a proxy server to use, and click OK.

   Note: You may not require a proxy if the AppLogic grid resides within your network.

   The server connection validates and the resources associated with the AppLogic server are added to Explore pane in the Resources tab.

More information:

Add Proxy Servers (see page 40)

Discover AppLogic Grids

During the setup of an AppLogic environment and its integration with CA Server Automation, you can manually refresh the AppLogic resources displayed in the CA Server Automation UI.

Note: An automated task performs this discovery job on a regular basis to refresh AppLogic resource information during normal operations.

Follow these steps:

1. View the Resources tab, and in the Explore pane, right-click AppLogic Virtual Clouds.
2. Select Management, Discover, Discover AppLogic Grid.

   All configured AppLogic servers are checked and the resource information for all grids, appliances, templates, and grid servers updates in the CA Server Automation UI.

Using Multiple Multi-instance AIMs

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

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

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

Example

Three vCenter Servers are managed through two AIMs:

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

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

Systems Management

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

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

Available management components for managed nodes are:

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

Supported Virtual Environments and Operating Systems

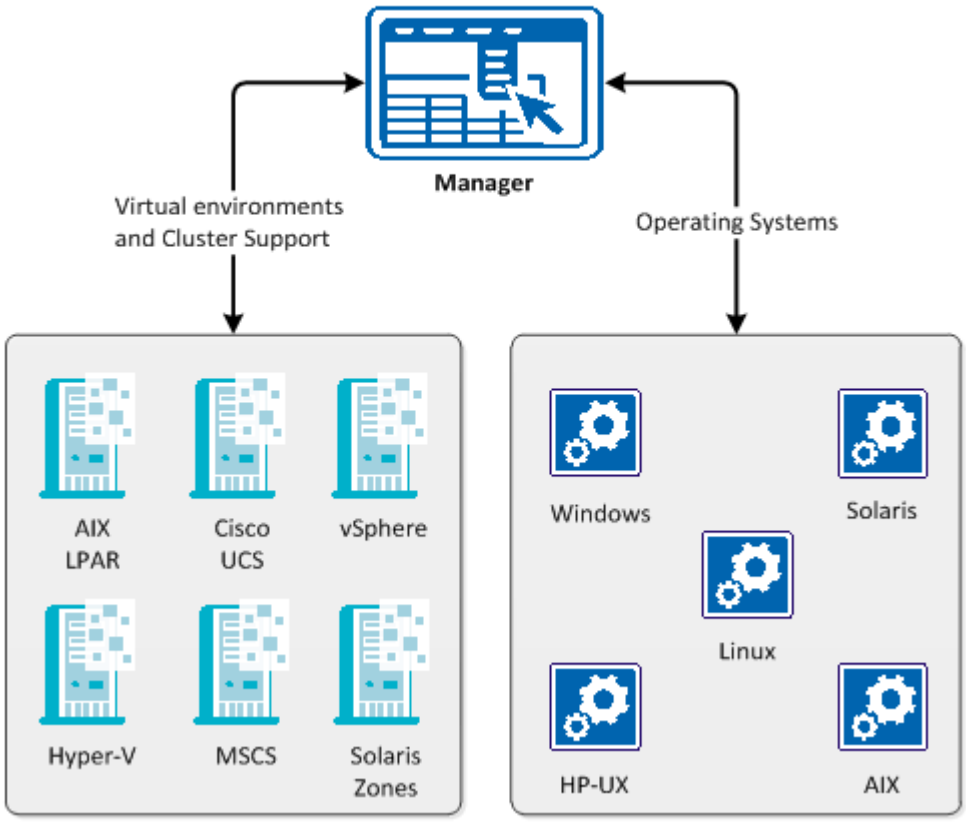

SystemEDGE is the base for systems management in CA Server Automation and provides the following benefits:

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

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

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

More Information

SystemEDGE Features (see page 308)
Service Response Monitoring (see page 318)
Remote Monitoring (see page 322)
Agent Configuration (see page 323)
Agent Visualization (see page 324)
Security and Maintenance (see page 327)

SystemEDGE Features

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

SystemEDGE supports monitoring metrics from the following MIBs:

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

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

Self monitoring

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

Process and service monitoring

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

Process group monitoring

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

Log file and directory monitoring

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

Windows event monitoring

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

History collection

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

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

More Information

Systems Management MIB (see page 310)
State Management Model (see page 312)
Configure Object Aggregation (see page 313)
Stateless Monitoring (see page 314)
Managed Mode and Legacy Mode (see page 315)
SystemEDGE AIMs (see page 316)

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 Server Automation Management Model. The agent aggregates multiple monitors of different severities into a single Managed Object. This object has a state corresponding to the breached monitor with the worst severity.

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

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

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

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

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

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

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

Configure Object Aggregation

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

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

Follow these steps:

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

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Click Control Settings.

   The Controls page appears.

4. Click Aggregate Monitors.

   The Aggregate Monitors page appears.

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

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

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

   Send legacy traps for all aggregated monitors

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

   Execute commands of all aggregated monitors

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

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

More Information:

Define SystemEDGE Policy Control Settings (see page 366)

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 Server Automation user interface, but you cannot visualize the resultant data. You must rely on traps that the agent sends when one of the following is detected based on defined monitors:

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

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

Managed Mode and Legacy Mode

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

You can also operate SystemEDGE in legacy mode, or without a CA Server Automation 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 Server Automation, 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 Server Automation, provide a CA Server Automation Manager node for the agent to run in managed mode.

SystemEDGE AIMs

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

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

You can configure and deploy these AIMs and SystemEDGE from the CA Server Automation user interface.

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

Remote Configuration and Deployment

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

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

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

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

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

Service Response Monitoring

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

The following diagram illustrates these relationships.

Interaction Between Service Response Monitoring Components

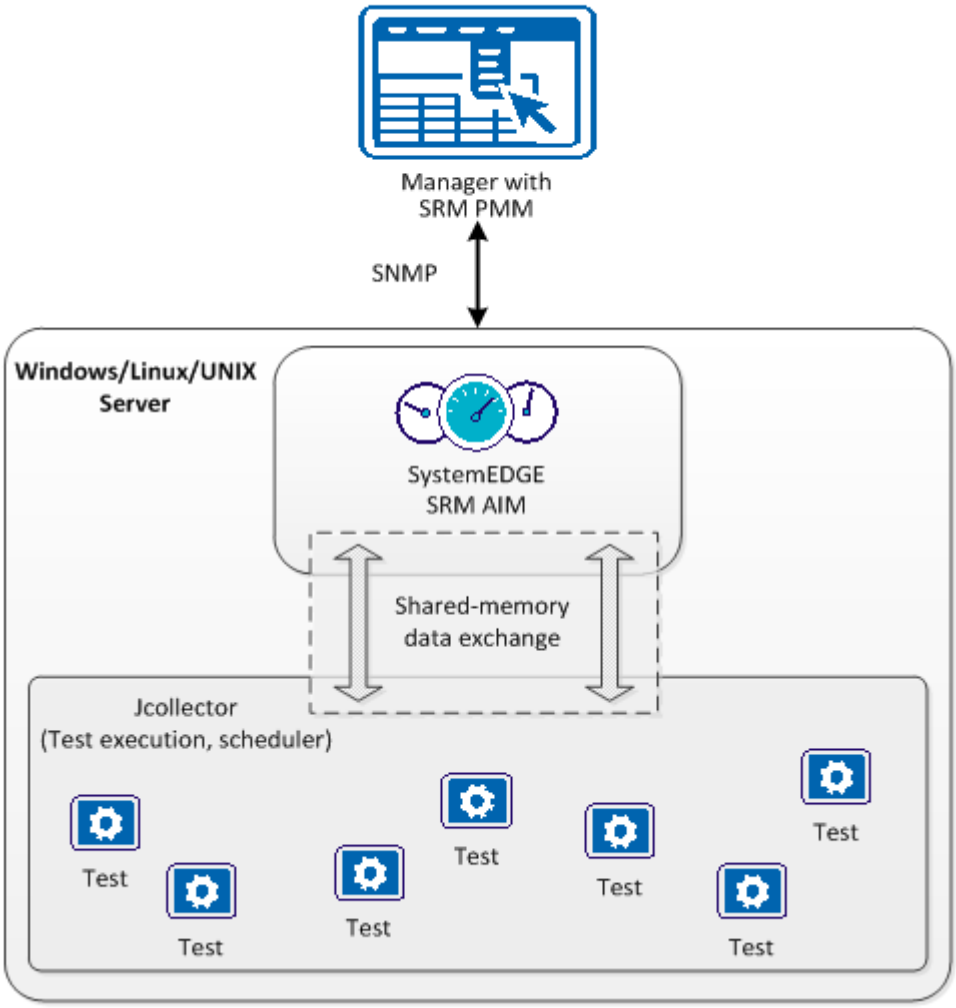

The svcrsp.cf configuration file contains the test specifications. The SRM AIM reads this configuration file and makes the test specifications available in the shared memory segment. The SRM Jcollector component reads each test configuration from the shared memory. Jcollector executes the tests, collects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

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 Server Automation 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 319)

SRM Tests

The SRM AIM provides the following response time tests:

Active Directory

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

Custom

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

DHCP

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

DNS

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

File I/O

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

FTP and TFTP

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

HTTP and HTTPS

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

LDAP

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

NIS

Verifies that NIS map requests are being processed.

NNTP

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

Ping

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

Email

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

SNMP

Verifies that SNMP agents are responding to SNMPv1 GET requests.

SQL Query

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

TCP

Verifies that systems are listening for and processing connection requests.

Virtual User

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

Remote Monitoring

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

Remote Monitoring collects and provides Windows system information (Win32_ComputerSystem) performing WMI queries (port 135) on the monitored RM systems. WMI uses port 135 (default).

The following diagram illustrates these relationships.

Interaction Between Remote Monitoring Components

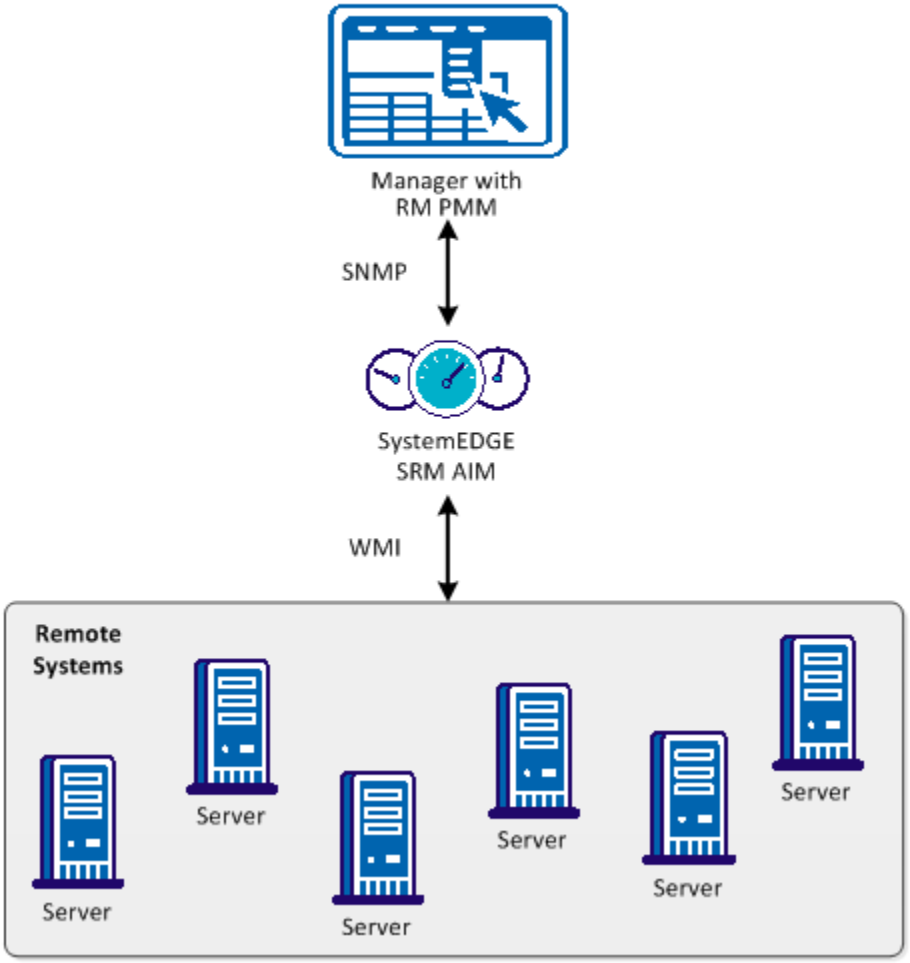

Agent Configuration

The following two types of SystemEDGE configuration are available from the CA Server Automation 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 Server Automation 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 Server Automation policy configuration overrides these changes after the policy is applied.

More Information

Perform a Point Agent Configuration (see page 323)
Policy Configuration (see page 363)

Perform a Point Agent Configuration

CA Server Automation 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 Server Automation 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.

Agent Visualization

The CA Server Automation 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 Server Automation 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 Server Automation 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 325)
View SystemEDGE Monitors (see page 325)
View SRM Tests (see page 326)
Dashboard Status Views (see page 327)

View Managed Object States

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

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

The system Summary page appears in the right pane.

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

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

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

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

View SystemEDGE Monitors

The CA Server Automation user interface displays all defined self and process monitors for systems running SystemEDGE in managed mode. You can view details about each monitor and perform point configuration (see page 323) such as adding, deleting, modifying, or copying a monitor.

To view SystemEDGE monitors

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

   The system Summary page appears in the right pane.

2. Click Configuration, and click Self Monitors or Process Monitors.

   The Self Monitors or Process Monitors pane appears.

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

- Index
- State
- Status

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

- Class, Instance, and Attribute of the object

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

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

View SRM Tests

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

To view SRM tests

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

   The system Summary page appears in the right pane.

2. Click Details, and click SRM Tests.

   The SRM Tests pane appears.

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

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

Dashboard Status Views

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

SystemEDGE Machines Status

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

SystemEDGE Objects Status

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

Security and Maintenance

CA Server Automation 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 Server Automation
- Notification of changes performed outside of CA Server Automation, and the opportunity to override or reject unwanted changes

More Information

Enable Maintenance Mode (see page 327)

Enable Maintenance Mode

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

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

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Expand Managed, and select a system.

3. Click the Resources tab, and then Monitoring Software.

   The Machine Details pane appears.

4. Set the Maintenance Mode option to Enabled, and click Apply.

   The agent performs a warm start and enables maintenance mode.

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

CA Network Automation

CA Server Automation integrates with CA Network Automation to provide the following functionality to Cisco routers and switches that run the Cisco Internetwork Operating System (Cisco IOS):

- Network device discovery
- VLAN listing
- VLAN creation and deletion

How to Manage Network Devices

Use the following process to integrate CA Server Automation with CA Network Automation and manage your network devices.

1. Configure the CA Network Automation server.
2. Configure CA Network Automation to discover network devices.
3. Obtain CA Server Automation scripts for use with CA Network Automation; scripts are located in the *install-directory*\nma_scripts folder.

   To import these scripts into the NetMRI system, navigate to the Configuration Management, Job Management page, and click Import.
4. In the CA Server Automation user interface, refresh the Network Automation devices list.
5. Click Resources, and open the Explore pane.
6. Right-click Data Center, and select Management, Network Management Automation. Select Refresh from the Actions drop-down menu.
7. Select the network device on which to perform network device functions using CA Network Automation.

Chapter 6: Managing Policies

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

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

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

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

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

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

For example, CA Server Automation is integrated with CA Service Desk Manager, 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 Service Desk Manager. This arrangement is useful for actions requiring third-party approval. After the third party approves your ticket in CA Service Desk Manager, 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 Service Desk Manager

For CA Service Desk Manager releases before Version 12.5, configure CA Service Desk Manager 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 Server Automation and CA Service Desk Manager need not be the same, as long as the two products do not share a database.

To configure CA Service Desk Manager

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

   http://*servicedesk_servername*:8080

   The CA Service Desk Manager splash screen appears.

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

   The CA Service Desk Manager 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 Service Desk Manager setup is complete and you can now automatically open requests when an action is triggered.

Configure the CA SDM Ticket Status Setting

CA Service Desk Manager versions before 12.5 used default status code settings of APP (Approved) and REJ (Rejected) for help desk tickets. CA Server Automation 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 Service Desk Manager Version 12.5, new ticket status codes are supported. PRBAPP (Approved) and PRBREJ (Rejected) must be associated to the existing approval codes in CA Server Automation. 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 Server Automation *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 Server Automation is configured to use the CA Service Desk Manager 12.5 status codes.

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

Substitution Parameters

In predefined and custom action types and in action sequences, you can use substitution parameters which are replaced by a value when the action is run. Parameters can be selected from drop-down menus of those options where they can be used.

Note: You can use only one substitution parameter for an option.

The following substitution parameters are available:

%SERVER%

The name of the current server. This parameter can be used only for actions assigned to a rule that is created on the server level.

%SERVICE%

The names of all services to which the server belongs, separated by commas. This parameter can be used only for actions assigned to a rule that is created on the service level.

%ACTIONNAME%

The name of the current action. Valid only for actions assigned to a rule.

%RULENAME%

The name of the rule to which the action is assigned. Valid only for actions assigned to a rule.

%EVENTSOURCE%

The source of the event (a system or component). Valid only for actions assigned to the event type rule evaluation.

%EVENTMESSAGE%

The message that the event generates. Valid only for actions assigned to the event type rule evaluation.

The %EVENTMESSAGE[P,0]% and %EVENTMESSAGE[“Regex”]% formats are used for extracting values from messages.

%VMNAME%

The virtual machine name used in string substitution. This parameter can only be used for a virtual machine node of the VMware vCenter Server.

%DATACENTER%

The Datacenter name used in string substitution. This parameter can only be used for a virtual machine node of the VMware vCenter Server.

%HOSTSYSTEM%

The host system used in string substitution. This parameter can only be used for a virtual machine node of the VMware vCenter Server.

%VCSERVER%

The VMware vCenter Server name. This parameter can only be used for a virtual machine node of the VMware vCenter Server.

%STDOUT%

Standard output. This parameter can only be used for an action sequence.

%STDERR%

Standard error. This parameter can only be used for an action sequence.

%EXITCODE%

The exit code returned when the action is run. This parameter can only be used for an action sequence.

%AutoIncrement(0)%, %AutoDecrement(0)%

These parameters are used for automatically generated text. They increase or decrease the counter by 1 in each run. The initial value (represented by the number in parentheses) is set by the user.

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 Server Automation 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 Service Desk Manager 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 Service Desk Manager 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.

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.

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 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.

Manage Actions Monitoring (CA Process Automation)

You can manage CA Process Automation processes to monitor actions.

Follow these steps:

1. In the Explore pane, select the Data Center node.
2. Click Resources, Policy, and then click the Action Monitoring tab.

   The CA Process Automation Processes page appears.
3. Select a process.
4. Click an item in the Actions drop-down menu.

   The monitoring process is placed in the selected state.

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 347)
Use Case: Adding a New Rule to a Service (see page 347)
Use Case: Defining an Action (see page 348)

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 Server Automation user interface.
3. Verify the results of adding the server to the service:
   - The action has been created with the description provided by the user.
   - The action is now available for rules.
   - The action is now available for job scheduling.

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

Configuring Data Collection

You can control how the following data is collected in your data center:

- The time intervals for metrics collection to occur.
- The systems from which you want to collect metrics (filtering).
- The metrics that you want to collect for each server.
- Data aging and the expiration of data (how long to retain data).

More information:

Key Points About Metrics Collection (see page 349)
Configure Data Collection for a Data Center (see page 351)
Configure Data Collection for a Server (see page 352)
Configure Data Collection for a Virtual Resource (see page 354)
Configure Performance Thresholds (see page 356)
Configure the Metric Filter (see page 357)

Key Points About Metrics Collection

To make informed decisions when you select metrics, review these points to understand CA Server Automation performance and application metrics collection:

- How does CA Server Automation collect metrics data? CA Server Automation communicates with the CA Performance Agent or with the SystemEDGE agent on the remote computer to collect the specified system metrics. The CA Performance Agent is a lightweight version of the CA NSM Performance Agent. If you have installed the CA NSM Performance Agent, CA Server Automation can also poll that agent.

  Note: The CA Performance Agent works differently from the CA NSM Performance Agent, so the metrics available for collection on Linux can be different, depending on which agent you use.

  The CA Performance Agent or the SystemEDGE agent must be installed on any server from which you want to collect the base system metrics, unless you are already using CA NSM performance agents or SystemEDGE agents. If CA NSM performance agents or SystemEDGE agents are present, then the CA Performance Agent is not required. If necessary, you can install the SystemEDGE agent using the product user interface. All performance metrics are stored in the Performance DB.

- How is overall utilization calculated? Overall utilization is an aggregate calculation of all the metrics that are currently being collected for servers managed by CA Server Automation. The calculation is based on the value of the metrics and the user-defined thresholds that define the parameters for normal operation. Any new metric that you select for collection is not used in the overall utilization calculation unless you select Include for Overall Calculation in the Policy, Metrics, Thresholds section of the user interface. When you do this, CA Server Automation does not provide false results when evaluating the state of the servers.

- How is overall utilization impacted by metric evaluations? The metric details provided in the tables help you understand how CA Server Automation evaluates the different metrics. Each metric has a method property set to either *exact* or *complement*. A higher exact value is a worse scenario than a lower exact value because it indicates an increase in overall utilization. A higher complement value is a positive scenario because it indicates a decrease in overall utilization. Generally, a high exact value negatively impacts overall utilization and a low exact value positively affects overall utilization. By contrast, a high complement value positively impacts overall utilization, and a low complement value negatively affects overall utilization. For example, if the value of Memory: Percentage Committed Bytes In Use increases, overall utilization of the system increases. If the value of Memory: Available MB increases, overall utilization decreases.

- What are the default metrics? The default metric definitions are located in the metric list in the Filter section for all supported platforms. You can find the default metrics indicator on the metric list with the value Yes in the Default column. CA Server Automation uses this list to obtain the metric definitions when you add a new server. You can configure platforms, types, subtypes, instances, and the type of data to collect in the Filter section. The metric filter and definitions for each server are stored in the Performance DB.
- Is performance data currently available for my systems? By default, if data cannot be collected, CA Server Automation does not negatively affect server state because lack of data does not reflect server criticality. By reviewing the Events list or selecting a specific system, you can determine whether metric data is being collected. However, if a more immediate means of determining this is needed, or if performance data is critical, CA Server Automation can be configured to change the state of a system automatically to Warning or Critical if performance data cannot be collected. To enable easy identification of systems where performance data is not available, modify the caaipconf.cfg file located in the CA Server Automation *install_path*\conf directory. Open the file with a text editor and locate the health state property as follows:

```
<property name="CONFIG_KEY_DEFAULT_HEALTH_STATE">
            <!-- Valid values: 0 (Unknown); 5 (OK); 10 (Warning); 15 (Minor Failure);
20 (Major Failure); 25 (CriticalFailure) -->
            <!-- Changes the value of HealthState for the CA_CollectionState object
associated to the CA_ComputerSystem -->
            <!-- If set to 30, CE will not set the HealthState. -->
            <value>5</value>
            <displayName>Default node health state when problem encountered in metric
or data collection</displayName>
        </property
```

By modifying the value surrounded by the value XML elements to one of the other supported values such as 5 or 10, for OK or Warning respectively, CA Server Automation reflects the desired state when performance data cannot be collected. For example:

```
<property name="CONFIG_KEY_DEFAULT_HEALTH_STATE">
            <!-- Valid values: 0 (Unknown); 5 (OK); 10 (Warning); 15 (Minor Failure);
20 (Major Failure); 25 (CriticalFailure) -->
            <!-- Changes the value of HealthState for the CA_CollectionState object
associated to the CA_ComputerSystem -->
            <!-- If set to 30, CE will not set the HealthState. -->
            <value>10</value>
            <displayName>Default node health state when problem encountered in metric
or data collection</displayName>
        </property
```

Because <value> was changed to "10", systems that do not have performance data available are displayed in a warning state in the CA Server Automation user interface.

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

Configure Data Collection for a Data Center

You can configure data collection at the Data Center level. The Data Center level policy takes effect immediately.

Follow these steps:

1. Click Resources, and select the Data Center folder in the Explore pane.
2. Right-click, and select Policy, Configure Collection Settings.

   The Settings dialog appears.
3. Complete the following fields in the Collection Setting section:

   Data recording interval (seconds)

   Defines how often the data is collected and stored in the Performance DB.

   Default: 300 seconds

   Polled data retention (days)

   Defines how long to store the polled data in the Performance DB. Consider the number of managed systems, services, and metrics collected when defining this number. The stored polled data objects accumulate over time and can impact performance. If performance issues arise, decrease the number of retention days.

   Default: 14 days

   Daily rollup data retention (days)

   Defines how long to store the average of the daily data in the Performance DB.

   Default: 365 days
4. Enter the threshold limits in the Thresholds section and then click Save.

   Your settings are saved.

Configure Data Collection for a Server

You can configure data collection for individual servers. Use this procedure to configure specific servers to collect data for the data center. You can also select metrics to monitor, set threshold values for individual metrics, and include metrics in overall utilization.

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Open the Data Center folder, and select the service to which the server belongs.
3. Right-click and select Policy.

   The Policy submenu appears.

4. Click Metrics.

   The Metrics wizard opens.

5. Select the server for which you want to configure data collection.
6. Complete the following fields in the Set Interval dialog:

   Use Default

   Specifies the data center level as the default when selected. If you leave the check box cleared, the values that you specify are used instead.

Data recording interval (seconds)

Defines how often the data is collected and stored in the Performance DB.

Default: 300 seconds

Daily rollup data retention (days)

Defines how long to store the average of the daily data in the Performance DB.

Default: 365 days

Polled data retention (days)

Defines how long to store the polled data in the Performance DB. Consider the number of managed systems, services, and metrics collected when defining this number. The stored polled data objects accumulate over time and can impact performance. If performance issues arise, decrease the number of retention days.

Default: 14 days

7. Select the metrics to monitor from the Available Metrics section and then click the down-arrow.

   The selected metrics are moved to the Selected Metrics to Collect section.

   Note: If you disable the default metrics (CPU and memory) and enable others, you will not see an overall utilization until you modify the thresholds of the newly selected metrics.

8. You can configure which performance metrics to monitor for each server and set threshold boundaries for each metric. Select the metric for which you want to set thresholds and complete the following fields:

   Upper Threshold

   Defines the upper limit of utilization for the selected metric group.

   Default: 80%

   Lower Threshold

   Defines the lower limit of utilization for the selected metric group.

   Default: 20%

   Include for Overall

   Specifies that you want the selected metrics to be included in the overall utilization calculation and evaluated by CA Server Automation.

9. Click Finish to save your settings.

Configure Data Collection for a Virtual Resource

You can configure data collection for virtual platforms and the virtual resources created and managed on those platforms. Use this procedure when you want to configure specific virtual machines or other resources to collect data at an interval that differs from the default for the data center. You can also select metrics to monitor, set threshold values for individual metrics and include metrics in overall utilization.

You can configure data collection for the following virtual platform objects:

- vCenter Server
- vCenter Data Center
- vCenter ESX Server
- vCenter Virtual Machine
- Hyper-V
- Microsoft Clusters
- Microsoft Cluster Nodes
- IBM AIX LPAR Server
- IBM Logical Partition
- Solaris Zones Server
- Solaris Zone

To configure data collection for a virtual resource

1. Click Resources, and open the Explore pane.
2. Expand the Data Center or MS Cluster Service folder, then any subfolder, and select the object that you want to configure.

   Subtabs for that object appear in the right pane.

   Note: If you select a top-level folder (such as VMware vCenter Server) or an object for which no data is collected (such as a vCenter cluster), you must select the specific object contained within the folder or object for which to configure data collection.

   Note: If you select MS Cluster Service as the top-level folder, then you see clusters and their nodes.

3. Right-click and select Policy, Configure Server Metrics Collection.

   Note: If you select the top-level folder for Solaris Zones, the Hardware Class column in the System section always shows the value Other.

4. Select the metrics that you want to monitor from the Available Metrics section and then click the down arrow.

   The metrics you select move to the Selected Metrics to Collect section.

   Note: If you disable the default metrics (CPU and memory) and enable others, you will not see an overall utilization until you modify the thresholds of the newly selected metrics.

5. Click Save to apply the selected metrics.
6. Right-click the resource, and select Policy, Configure Collection Settings.
7. Complete the following fields in the Collection Setting section:

   Use Default

   Specifies the data center level as the default when selected. If you leave the check box cleared, the values that you specify are used instead.

   Data recording interval (seconds)

   Defines how often the data is collected and stored in the Performance DB.

   Default: 300 seconds

   Daily rollup data retention (days)

   Defines how long to store the average of the daily data in the Performance DB.

   Default: 365 days

   Polled data retention (days)

   Defines how long to store the polled data in the Performance DB. Consider the number of managed systems, services, and metrics collected when defining this number. The stored polled data objects accumulate over time and can impact performance. If performance issues arise, decrease the number of retention days.

   Default: 14 days

8. Click Save to save your settings.

   Note: The default thresholds are used. If you want to modify thresholds, you must do this separately.

Configure Performance Thresholds

You can configure which performance metrics to monitor for each server and set threshold boundaries for each metric.

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Expand the Data Center folder and any subfolder, then select the server that you want to configure. Navigate to a virtual server to select a specific virtual resource, such as a virtual machine or logical partition.

3. Right-click and select Policy.

   The Policy submenu appears.

4. Click Configure Threshold Settings.

   The Configure Threshold Settings appears.

5. Select the metric for which you want to set thresholds and complete the following fields:

   Upper Threshold (%)

   Defines the upper limit of utilization for the selected metric group.

   Default: 80%

   Lower Threshold (%)

   Defines the lower limit of utilization for the selected metric group.

   Default: 20%

   Include for Overall Utilization Calculation

   Specifies that you want the selected metrics to be included in the overall utilization calculation and evaluated by CA Server Automation.

6. Click Modify to save your settings.

Configure the Metric Filter

You may want to add or delete metrics to or from the metric filter for the Data Center, depending on which performance metrics you want to monitor.

Note: For information about the performance metrics available for CA NSM Performance Agent users, see the CA NSM *Systems Performance Metrics Help* in the CA NSM documentation set.

To configure the metric filter

1. Select the Data Center folder in the Explore pane.
2. Right-click and select Policy, Configure Collection Criteria.

   The Collection Criteria dialog appears.
3. Do *one* of the following:
   - Select the check box for an existing metric to modify an existing entry. The information for the selected metric populates the fields of the Details section. Make any changes and click Update.
   - Select an OS, and complete the fields in the Details section to add a new metric, then click Add.

   The metric is saved.

   The Details section contains the following fields:

   OS

   Defines the operating system for the metric being monitored.

   Type

   Defines the type of metric being monitored.

   Example:

   Type: CA Disk Group

   Sub Type: Writes per second (average)

   Sub Type

   Defines which aspect of the metric is being monitored.

   Example:

   Type: CA Disk Group

   Sub Type: Reads per second (average)

Instance

Defines the instance of the managed object in the MIB hierarchy.

Example:

Type: vmvcaim.StatClusterEffectiveCPU

Sub Type: 1.3.6.1.4.1.546.16.52.2.7.2.1.14

Instance: %3 [%2]

%<n> where <n> is the numeric value listed under Instance matched to any value corresponding to the n^{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 the collection engine collects the metrics specified by the filter by default. Unless a metrics filter is set as default, the collection engine will not automatically collect the metrics specified.

Include for Overall Utilization Calculation

Specifies that you want the selected metrics to be included in the overall utilization calculation and evaluated by CA Server Automation.

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.

Enhanced Storage Policies

An enhanced storage policy is a policy you define for use with the CA Server Automation Storage Provisioning Manager. Creating and using enhanced storage policies lets you perform the following tasks:

- Reuse Storage Provisioning Manager parameters

  The Storage Provisioning Manager requires you to supply parameters each time you provision storage using the storage service or provisioning policy methods. Using enhanced storage policies lets you create a set of predefined parameters you can quickly apply during storage provisioning.

- Automate storage provisioning tasks at the enterprise level

  Storage vendors provide policies for managing storage, but these policies are not server aware, making it difficult to automate provisioning at the enterprise level. CA Server Automation enhanced storage policies combine storage operating environment information with IT infrastructure information about servers and services to provide protocol-based policies that are server aware.

- Distribute storage provisioning tasks to nontechnical members of your organization

  Enhanced storage policies let you create policies that mask the complexity of storage provisioning then distribute storage provisioning tasks to nontechnical personnel.

UCS Action Types for Policies

CA Server Automation provides Cisco UCS action types for blade power operations and service profile operations.

More information:

Create a Blade Power Action (see page 360)
Service Profiles (see page 301)
How to Create or Update a Service Profile (see page 361)
Associate Service Profiles with Blades (see page 362)

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.

How to Create or Update a Service Profile

CA Server Automation 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 Server Automation 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.

Create a Blade Provisioning Action

You can define an action for provisioning a Cisco UCS blade server.

To create a Cisco UCS blade provisioning action

1. Select the Resources tab, click Data Center in the Explore pane, select the Policy tab, and select the Actions sub-tab.
2. Click + (Add new action).
3. Select Provision Machine on the Type drop-down list.
4. Select Cisco UCS on the Environment drop-down list.
5. Employ the wizard as you would to provision a blade under Imaging.

   Note: If help desk approval is required, enter information as needed.

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

   The action is created and ready for use.

Rapid Server Imaging Support for UCS

You can specify a supported RSI image transition (physical-physical, physical-virtual, virtual-physical, or virtual-virtual) using a Cisco UCS blade as either an underlying type of physical server resource or as a VM host.

Example: Move Application from x86 System to Cisco UCS Blade

To migrate a Windows application from a legacy x86 system to a Cisco UCS system, place both systems under CA Server Automation control. Employ the RSI capability to capture the application image on the legacy x86 system, and transfer the image to a blade associated with an appropriate service profile on the Cisco UCS system.

Policy Configuration

You can configure managed agents and apply the configuration to multiple systems in one operation using centralized remote deployment from the CA Server Automation 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 Server Automation 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 Server Automation, CA NSM, SystemEDGE, and so on), and the ability to accept or reject changes through CA Server Automation
- 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 364)
How to Create SRM Policy (see page 374)
Apply Policy to Machines (see page 377)
Review Policy Application Progress (see page 378)
Configure and View Applied Policies (see page 379)
Agent Policy Dashboard Views (see page 381)
Example: How to Monitor User-specific Metrics (MIB Extensions) (see page 381)
Example: How to Monitor a Specific Windows Performance Registry Metric (see page 383)

How to Create SystemEDGE Policy

You create SystemEDGE policy to define a set of monitors, AIMs to load, configuration preferences, and other settings that control how the agent runs and what it monitors. Once you create a policy, you can apply it to any number of systems running SystemEDGE agents in managed mode. Policy lets you perform all configuration operations that you can manually configure locally with the benefit of a consolidated interface, pick lists, and dynamic deployment to remote systems.

The following process describes how to create SystemEDGE policy:

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

   The SystemEDGE pane appears.

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

   The New SystemEDGE Policy dialog appears.

3. Enter a name and description for the policy and whether to base it on an existing policy and click OK.

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

4. Define monitors to include (see page 365).
5. Define control settings (see page 366).
6. Define SNMP settings (see page 371).
7. Define MIB extensions (see page 374).

8. Click Save Policy.

   The policy is saved.

Monitor Definition

When configuring policy, click the Monitors tab to view a summary of all monitors currently defined. In the Monitors tab, subtabs exist for defining the following monitors:

- Self monitors
- Process monitors
- Log file monitors
- Windows event monitors
- History monitors
- Process group monitors

From each of these subtabs, you can create a new monitor, delete or modify existing monitors, or copy a monitor. Monitor definition using configuration policy provides the following functional benefits that are not available through local file manipulation:

- Intuitive field names
- Automatic index creation
- A platform designation that lets you limit a monitor to a specific platform
- Pick lists for object class and attribute for self monitors to eliminate the need to remember and enter individual OID numbers
- Pick lists for key properties in other monitors
- A monitor type designation for log file monitoring that lets you specify whether you are monitoring a log file or directory and adjusts the available fields dynamically
- Individual check boxes for setting flags to eliminate the need to provide a hexadecimal string for flag settings

  Note: See the sysedge.cf configuration file or the *SystemEDGE User Guide* for descriptions of these hexadecimal flags.
- Maintenance window settings integrated into the monitor

For more information about defining individual monitors for configuration policy, see the *Online Help*.

Define SystemEDGE Policy Control Settings

You can control the following agent behavior using the SystemEDGE policy control settings:

- Security settings
- SNMP settings
- MIB table population
- UNIX settings
- Performance monitoring settings

You can apply the control settings defined in the policy to all machines.

Follow these steps:

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

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Click Control Settings.

   The Controls page appears.

4. (Optional) Click Use Defaults.

   The default selections pane appears. You can change the default settings.

5. Configure the following control settings:

   SNMP

   Lets you define the following basic SNMP properties:

   Bind Address

   Specifies an interface to which the agent binds and listens for incoming SNMP requests. Valid addresses are IPv4 or IPv6 address.

   Note: The corresponding default _port is specified during installation.

   Bind Port

   Specifies the trap port the agent binds to for sending SNMP traps. If no bind_address is specified the agent binds to all available UDP addresses.

   Default: Port selected by the system

   IP Family

   Specifies the agent communication method: IPv4 only, IPv6 only, or both. By default, the agent tries using IPv4 and then IPv6.

FIPS Mode

Specifies the agent to use FIPS-compliant encryption. Select Non-FIPS Mode to enable the CA eTrust Public Key Infrastructure libraries, and if this method fails, fall back to the internal minimum security solution. Select FIPS Co-existence Mode to enable FIPS-compliant encryption, and if this method fails, fall back to the CA eTrust Public Key Infrastructure Libraries. Select FIPS Only Mode to enable the RSA BSAFE Crypto-C Micro Edition FIPS-compliant libraries and perform no encryption if they fail.

Default: Non-FIPS Mode

Trap Source

Specifies the source address used to send traps. Valid addresses are IPv4, IPv6 address, or a host name.

Default: Host name of the agent

Security Settings

Lets you define the following security preferences:

Authentication Traps

Sends an authentication failure trap when the agent receives an SNMP message with a community name that the agent cannot recognize.

Default: Disabled

Process Sets

Permits access to processes and other software running on agent systems in the Process table and Running Software table. Allowing SNMP Sets on these tables can cause security issues.

Remote Shell Group

Permits management systems to remotely instruct the agent to run shell scripts and programs on the agent system through the Remote Shell group. The disclosure of this type of information can post a potential security risk.

Execution Action

Enables the execution of action commands with the monitoring tables when a threshold breach occurs. The capability to run action commands and scripts can be a security issue.

MIB Table Population

Populates the following tables in the Systems Management MIB:

- Process Table
- User Group Table
- Who Table
- Trap Community Table
- Monitor Mirror Table
- Aggregate Mirror Table
- Top Processes Table

Each table either contains sensitive information that you can expose in a MIB or non-essential information that you can disable to save disk space. The default settings enable population of all tables except for the process table.

Miscellaneous

Lets you define the following miscellaneous settings:

Allow agent to be Updated using SNMP

Permits agent updates using SNMP Sets (for example, removes write communities). If you permit SNMP Sets on the agent, any updates through this method cause a notification of an SNMP Set change and also an exception when viewing policy details for the system.

Notify Manager of Configuration Updates

Enables the agent to send a notification to the manager for any SNMP Set request that the agent processes.

Warm Start Discovery

Enables an agent rediscovery of all devices after every warm start configuration update. If you manage a system with many devices, a discovery after every warm start can consume too much time and too many resources.

Use Perl Compatible Regular Expressions

Perl Compatible Regular Expressions (PCRE) enables you to specify i18n compatible regular expressions while defining monitors that support regular expressions. The examples of regular expressions are log file, process, process group, Windows services and Windows events. You can also use this option to create more complex regular expressions. This option is provided in SystemEDGE agent 5.1.0 and above versions.

Automatically Resolve Index Conflicts

Enables you to resolve Index conflicts. When you apply the layered templates to all machines, indexes are assigned to the monitors added in the template. If the assigned indexes conflict with existing indexes either within the base policy or another template, this option reassigns unique index values.

Note: Indexes contained within the base policy are always maintained in the delivered configuration. If this option is disabled, you cannot resolve conflicting indexes. However, when you apply layered templates to the machines, the conflicting indexes are displayed as errors on the layered templates that caused the conflicting indexes.

Historical Performance Monitoring

Lets you define the following settings for the Performance Cube AIM, which collects history information into Systems Performance cubes for historical performance management:

Collection Interval

Specifies how often to collect information from the History table into performance cubes.

Index Range Start

Specifies the beginning of the reserved range of indices, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

Index Range End

Specifies the end of the reserved range of indices, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

UNIX Control Settings

Lets you define the following settings for agents running on UNIX systems:

Sub-program Group

Specifies a group name other than root under which to run subprograms.

Sub-program User

Specifies a user name other than root under which to run subprograms.

Linux Freemem Include

Specifies whether to include system buffers, disk cached memory, or both in free memory calculation.

Query System Devices

Lets you enable querying of the following system device metrics:

- Serial device status
- Floppy disk status
- Disk size, capacity, description, and other properties (Probe Disks)
- NFS file system status
- HP-UX graphics status

Querying these metrics can cause issues with potential agent blocking. The default settings enable querying of only serial device status and NFS file system status.

6. Click Plugins.

   The Plugins pane appears. This pane controls which AIMs to load with the agent.

7. Do one of the following:
   - Select 'Load all available plugins' to load all AIMs available on the agent system.
   - Select 'Load plugins selected in the table'.
   - Click + (New) on the External Plugins toolbar to add an AIM to the External Plugins table.

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

AIM loading is configured.

8. Click Aggregate Monitors.

   Configure aggregate monitors as described in Configure Object Aggregation (see page 313).

   The control settings are defined.

9. Click Save Policy.

   The policy is saved.

More Information:

Configure Object Aggregation (see page 313)

Define Traps and Communities

SNMP settings define the communities that the agent uses and the destinations to which it sends traps.

Follow these steps:

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

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Click the Traps and Communities tab.

   The Communities page appears.

4. Select one of the following and click Action, select Apply:

   - Include only Server specific SNMP settings
   - Include Server specific SNMP settings and all Default settings
   - Select 'Include Server specific SNMP settings and selected Default settings

The SNMP settings are updated and the community page in the Communities table displays the following:

Name

Specifies the name of the community string.

Port

Specifies the port of SNMP.

SNMP Version

Specifies the SNMP version that the community uses.

Access Rights

Specifies whether the community should have read write or read-only permissions.

Note: Add at least one read-only and one read write community.

Community/User

Specifies the community name.

Authentication Protocol

Specifies the protocol to authenticate SNMPv3 data.

Privacy Protocol

Specifies the protocol to authenticate SNMPv3 data.

Access Control List

Specifies a space separated list of IP addresses to restrict community usage to those addresses only. If you leave the list blank, the agent grants access to any system that uses the associated community name. Access lists are only for communities that use SNMPv1.

Note: For information about defining SNMPv2c and SNMPv3 access lists, see the *SystemEDGE User Guide*.

5. (Optional) Add, update, or delete other communities as necessary.
6. Click Save Policy.

   The policy is saved.
7. Click Trap Destinations.

   The Trap Destinations page appears.

8. Define a trap destination using the following controls and click Add:

Trap Type

Specifies the type of trap to send, depending on the SNMP version.

Destination

Specifies the IPv4 or IPv6 address to which to send traps.

Port

Specifies the UDP port to which to send traps.

Community

Specifies the community name sent with the traps.

Encoding

(Optional) Specifies how to include the source address you defined in the Trap Source field of the Control Settings pane in traps. This parameter is important if the trap source translates to an IPv6 address. Enter the encoding parameter in a three digit format XYZ, assuming leading zeros.

Default: 000

X

Controls extending the four byte IPv4 source address field (SNMPv1 traps only). Enter 0 to not extend the source address field to include the 16 byte IPv6 address, and enter 1 to extend the source address field.

Y,Z

Controls the inclusion of source information into the trap's varbind (Y) or UDP packet (Z; SNMPv1 traps only). Enter one of the following for these digits:

0: Do not modify the trap's varbind or the outer UDP packet.

1: Include the trap_source parameter as is in the varbind or packet (IPv4/IPv6 address or host name).

2: Include the trap_source parameter preferably as an IPv4 address (then IPv6 address, then host name).

3: Include the trap_source parameter preferably as an IPv6 address (then IPv4 address, then host name).

4: Include the trap_source parameter preferably as a host name (then IPv4, then IPv6).

5: Follow the preference for 2 and include the host name.

6: Follow the preference for 3 and include the host name.

7: Follow the preference for 1 and include the host name (if trap_source is an IPv6 address).

Trap Source

(Optional) Specifies the IPv4 or IPv6 address or the host name to use as trap source.

Default: Global Trap

The trap destination appears in the Defined Trap Destinations table.

9. (Optional) Add, update, or delete other trap destinations as necessary.
10. Click Save Policy.

    The policy is saved.

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

Define MIB Extensions

Defining MIB extensions provide functional benefits that are not available in local file manipulation. The policy configuration feature provides field names and the list of key properties such as object type.

When you configure a policy or a monitoring template, click the MIB Extensions tab to add the following objects:

- MIB Extensions
- Windows Performance
- Windows Registry

Note: MIB Extensions within templates are supported for the purposes of applying the MIB Extensions directly to monitored systems. MIB Extensions for use within policies should be created directly in the Policy itself.

How to Create SRM Policy

You create SRM policy to define tests to perform, thresholds to monitor, configuration preferences, and other settings that control how the agent runs and what it monitors. Once you create a policy, you can apply it to any number of systems running SystemEDGE agents with the SRM AIM in managed mode. Policy lets you perform all configuration operations that you can manually configure locally with the benefit of a consolidated interface, pick lists, and dynamic deployment to remote systems.

The following process describes how to create SRM policy:

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

   The Service Response pane appears.

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

   The New Service Response Monitoring Policy dialog appears.

3. Enter a name and description for the policy and whether to base it on an existing policy and click OK.

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

4. Define tests to include (see page 375).
5. Define test thresholds (see page 375).
6. Define control settings (see page 376).
7. Click Save Policy.

   The policy is saved.

Test Definition

When configuring SRM policy click the Tests tab to view a summary of all tests currently defined. From the Tests pane you can define new tests and modify, delete, or copy existing tests. Defining tests using configuration policy provides the following functional benefits that are not available through local file manipulation:

- Intuitive field names
- Pick lists for key properties
- A pick list to specify the test type, after which the pane dynamically updates with the properties relevant to the test you specify
- A table for adding threshold definitions to a test

For more information about defining test for configuration policy, see the *Online Help*.

Threshold Definition

You can define threshold monitors for any attribute collected by an SRM test and add the monitors to any existing test. These monitors create SystemEDGE self monitor entries and let you track status information for SRM tests. Defining monitor thresholds using configuration policy provides the following benefits over manual file manipulation:

- Intuitive field names
- A pick list for choosing the attribute to monitor
- Fields for warning, minor, major, critical, and fatal values that provide the flexibility to define as many severities as necessary
- The ability to dynamically associate a threshold with any SRM test in the Test tab

For more information about defining monitor thresholds for tests, see the *Online Help*.

Define SRM Control Settings

SRM control settings define various aspects of AIM behavior that you typically control in the svcrsp.cf file, including the following:

- Security settings
- Log level
- Index reservations

Control settings defined in SRM policy are applied to all machines to which you apply the policy.

Follow these steps:

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

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Click the Control Settings tab on the policy page.

   The Control pane appears.

4. Configure the following control settings:

   Maximum Number of Threads

   Specifies the number of threads the jcollector should use to perform tests

   Log Level

   Specifies the log level of the SRM AIM. Default is Warning.

   Allow External Scripts

   Specifies if execution of external scripts is allowed.

   Allow Execution of File I/O Tests

   Specifies if execution of file I/O tests is allowed.

   Allow Untrusted SSL Certificates

   Specifies if SSL tests with sites that do not have trusted SSL certificates is allowed.

   Java bin Location

   Defines the location of the Java executable.

   Note: Specify the complete path and binary on AIX.

Override CLASSPATH in Environment

Defines extra classes to load. Overrides CLASSPATH in environment if defined.

No Collector

Specifies if SystemEDGE should start jcollector.

Bypass JRE Internal Cache

Specifies whether to bypass JRE internal cache.

No TOS for IPv4 (HP-UX)

Specifies whether to disable TOS.

Shared Memory Name

Defines the ID for the shared memory.

Reserved Test Indices

Defines reserved range of test indexes.

The control settings are defined.

5. Click Save Policy.

   The policy is saved.

Apply Policy to Machines

After you create configuration policy, apply it to machines across the enterprise. When you apply configuration policy, CA Server Automation pushes a compiled configuration file containing all policy settings to all specified agent machines, and the new policy is implemented after an automatic agent warm start.

You can also reapply the policy to machines if one of the following occurs:

- You updated the policy
- You received a notification that the configuration on an agent machine has changed

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Select the policy you want to apply.

   Policy details appear in the right pane.

4. Click Action and select Apply.

   Tabs appear for selecting machines on which to apply the policy. The 'Update machines running this policy' tab lets you apply the policy to machines that are already running the policy. The 'Apply to Machines not running this Policy' tab lets you apply policy to machines without any policy or using a different policy.

5. (Optional) Do one of the following from the 'Update machines running this policy' tab:

   - Select 'Update all machines using this policy' to deploy the policy on all machines currently running it. This option is useful if you have made configuration policy changes that you want to apply globally.
   - Select 'Update selected groups of machines' to update only machines that meet any of the following criteria:
     - Running an out of date version of the policy
     - Policy exceptions have occurred
     - Running a current version of the policy

     Select any of these options. Policy exceptions occur when a user applies a point configuration change to an agent that is not represented in the applied policy.
   - Select Advanced and manually the machines in the Select Machines pane to which the policy is to be reapplied.

6. (Optional) Select machines from the 'Apply to Machines not running this Policy' tab to which to apply the policy.

7. Click Apply Policy.

   The policy application is initiated.

Review Policy Application Progress

You can review the progress of policy application operations in detail for each individual policy.

Follow these steps:

1. Select the Resources tab, open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Click the Managed Machines tab.

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

4. (Optional) Click View Exceptions.

   The Policy Exceptions pane appears and lets you view SNMP sets that have been applied to the system since the policy was last applied.

   Note: This screen is only available for SystemEDGE Policies.

5. (Optional) Click View Configuration.

   The Policy Configuration page appears and lets you view the configuration file delivered for the agent.

6. (Optional) Click View Errors.

   The Policy Errors pane appears. If the policy could not be successfully applied, you can view a list of errors returned by the agent when the policy was rejected.

Configure and View Applied Policies

The Policy feature lets you manage the policies and templates applied to an individual server, a server group, or a service. You can perform the following operations:

- Update Policies and Templates
- View exceptions. Exceptions are SNMP sets that have been applied to the system since the policy or template that was last applied.
- View policy configuration
- View policy errors
- Bulk update Policies
- Delete Templates

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Select a system or a service. Click the Resources page, and then Monitoring Software.

   The Machine Details page appears.

3. Click Policies.

   The Policies page displays a list of Policies of SystemEDGE and SRM, and SystemEDGE Templates.

   Note: The Filter can display a list of layered templates that are Pending, Delivered (successful), Configured, and Failed.

4. You can bulk update policies and templates. In the Policies and Templates table, select the policy or template you want to bulk update, click Actions, and select one of the following options:

   - Bulk update SystemEDGE Policies.
   - Bulk update Service Response Policies.
   - Bulk update SystemEDGE Templates.

   Note: If the policy is being applied to a single server, you are prompted for the policy name.

   Bulk update for Policies:

   If the selected policy is applied to a service group, a dialog provides an option to select the machine on which you want to apply the policy.

   Bulk update for Templates:

   A dialog provides an option to select the templates from the Available Templates. After you select the templates, click one of the following options:

   - Apply the selected templates as the set for all machines - Removes the existing templates that are applied to all machines and applies the selected templates to all machines.
   - Apply the selected templates to the existing set already applied to all machines - Adds the selected templates to all machines. If the templates selected include the existing templates that are applied to machines, the existing templates are reapplied.

5. Click Apply Policy to apply the policy or template to the machines.

   On the Policies page, you can view the progress of the policy or template being applied to the machines.

6. (Optional) Click View Configuration Icon.

   The Policy Configuration page appears. For a machine with a template, it displays the Policies and Templates, and SystemEDGE Configuration file. For a machine with a Service Response Monitor, it additionally displays the Service Response Monitor Configuration file.

7. Click Save Policy.

   The policy is saved.

Agent Policy Dashboard Views

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

Policy Status Summary

Displays a pie chart and list showing the number of policies. A system can be in five different states:

Unconfigured

SystemEDGE agent is installed but no policy is configured.

Agent Installed

SystemEDGE agent is installed.

Configured

SystemEDGE agent is installed and a policy is configured.

Configuration Error

SystemEDGE is installed and a policy is configured but the last configuration failed.

Installed but not managed

SystemEDGE is installed but running in a mode that cannot be managed by policy configuration.

Policy Breakdown

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

Machines with nonstandard Policies

Displays a table listing systems that contain nonstandard changes to an applied policy.

Example: How to Monitor User-specific Metrics (MIB Extensions)

This step-by-step example describes how to monitor a user-specific metric.

How to monitor user-specific metrics (MIB extensions)

1. Create a program that returns the data required. For example, a simple DOS batch script on the agent system to return some fixed data.

```
@echo off
echo 99
```

2. Open a text editor and store these two lines in data.bat on the C: drive.

3. Create an MIB extension that references this batch file.
   a. From the user interface, click *Policy*, open *Configuration* in the navigation pane, expand the Policy tree, and open a SystemEDGE policy.

      The policy details appear in the right pane.

   b. Click the MIB Extensions tab.

      The MIB Extensions pane opens.

   c. Add the following data into the fields:

      Index: 1 (if it is the first MIB extension)

      Type: integer

      Extension Command: C:\data.bat

      Access Rights: Read Only

   d. Click *Add*.

      The MIB Extension is added to the Policy.

   e. Click *Save Policy*.

      The policy is saved.

4. Create a threshold monitor to check the value of the new monitor.
   a. Click *Monitors* and then *Thresholds*.

      The Threshold Monitor Details Edit pane appears.

   b. Add the following data into the fields:

      Index: (automatically added)

      Platform: OS Independent

      Object Class: extensionGroup [Extended mib from adding new scalar variables]

      Object Attribute: 1

      Object Instance Name: MyData

      Interval: 60

      Severity: Major Alarm

      Operator: greater than or equal to

      Value: 50

      Scale: 1

      Sample Type: absolute value

c. Click *Save*.

The policy is saved. A ‘major’ alarm with threshold ‘50’ is added. This threshold will be breached immediately as the script created previously always returns the value ‘99’.

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

The Selected Machines pane appears.

e. Verify that the selected machines are correct and click *Apply*.

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

Click *Return to Policy*.

The policy details pane appears.

Once the agent is configured, you can view the state of this threshold monitor from the Resources tab. You can see that the “major” threshold has been breached.

Example: How to Monitor a Specific Windows Performance Registry Metric

This step-by-step example describes how to monitor a user-specific metric.

How to monitor user-specific metrics (MIB extensions)

1. Create a MIB extension for a Windows Performance Registry metric.

   a. From the user interface, click *Policy*, open *Configuration* in the navigation pane, expand the Policy tree, and open a SystemEDGE policy.

   The policy details appear in the right pane.

   b. Click the *MIB Extensions* tab.

   The MIB Extensions pane opens.

   c. Click *Windows Performance*.

   The Windows Performance Defined Extensions pane appears.

   d. Add the following data into the fields:

   Index: 1 (If the extension is the first one)

   Type: integer

   Object: System

   Counter: Processes (Provides the total number of running processes).

   The System metrics have no ‘instance’ so this field is left blank.

e. Click *Add*.

The MIB Extension is added to the Policy.

f. Click *Save Policy*.

The policy is saved.

2. Create a threshold monitor to check the value of the new monitor.

a. Click *Monitors* and then *Thresholds*.

The Threshold Monitor Details Edit pane appears.

b. Add the following data into the fields:

Index: (automatically added)

Platform: OS Independent

Object Class: ntRegPerf (NT Registry and Performance Group)

Object Attribute: 1

Object Instance Name: NumProcesses

Interval: 60

Severity: Warning Alarm

Operator: greater than or equal to

Value: 50

Scale: 1

Sample Type: absolute value

c. Click *Save.*

The policy is saved. A 'warning' alarm with threshold '50' is added.

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

The Selected Machines pane appears.

e. Verify that the selected machines are correct and click *Apply*.

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

Click *Return to Policy*.

The policy details pane appears.

Once the agent is configured, you can view the state of this threshold monitor from the Resources tab.

Chapter 7: Provisioning Resources

This section contains the following topics:

Imaging Services

CA Server Automation can provision new physical and virtual computers and also reimage existing resources. Physical computer imaging is available for servers that use Windows and Linux operating systems. Provisioning functions let you clone, migrate, configure, and change the properties of VMs.

Imaging services integrate with and use the following technologies to perform provisioning operations:

- VMware vCenter Server integration for VM provisioning
- Amazon Elastic Cloud (EC2) for image provisioning in the cloud
- CA Software Delivery OSIM for imaging Windows and Linux servers
- JumpStart servers for Solaris server imaging
- Rapid Server Imaging (RSI) for imaging dissimilar hardware

Events are generated for the following Imaging services actions:

- Submitting an imaging job to the imaging server
- Changes to imaging job status
- Discovering the target computer after the imaging job succeeds.

CA Software Delivery

The software delivery service of CA Server Automation manages the software delivery requests generated by CA Server Automation and sent to CA Software Delivery. The software delivery service manages communication for the delivery of operating system imaging for Windows and Linux, and application packages to UNIX, Linux, and Windows environments. The software delivery service handles all incoming operating system requests from the imaging service.

The imaging service uses the CA Software Delivery service (OSIM) to initiate an imaging process on remote Windows or Linux servers. The requests are sent to the software delivery service for imaging which then initiates the process on the target server in CA Software Delivery.

The software delivery service provides the integration to OSIM. The service retrieves OS image information from CA Software Delivery about the images that are available on the server. These images can be Windows or Linux operating systems. The Windows images can also be Ghost images. These images are then displayed and you can select from them and submit an imaging job on a client computer that you define.

This documentation assumes that you are familiar with the CA Software Delivery (OSIM) solution. All requirements and restrictions imposed by the solution are valid for CA Server Automation.

Understanding Packaging

Hundreds of packages are potentially registered in the CA Software Delivery server. CA Server Automation lets you select a subset of packages, so that you can manage the packages that apply to your data center. You can then deploy these packages from the CA Server Automation user interface to remote computers or schedule deployments as jobs.

When you schedule multiple packages for deployment, CA Software Delivery deploys jobs one at a time. You create both jobs in the CA Server Automation user interface. The first job is sent to the CA DSM interface and the other job is queued. If you view the jobs from the CA DSM interface, one job appears. When the first job completes, then the next job appears. If you view the jobs in the CA Server Automation interface, you can view multiple scheduled jobs at the same time.

Patch Management

CA Patch Manager manages software patches in heterogeneous environments. You can have patches delivered to managed servers in CA Server Automation when you have CA Patch Manager installed with CA DSM. Patches can be delivered to computers that are identified as meeting the prerequisites for patch delivery.

CA Server Automation checks the status of the patch to verify if it is approved or not. Only approved patches appear in the list of available packages when you are setting up package delivery in CA Server Automation, regardless of whether the patch is available in CA ITCM.

To deliver a patch to a server, the server must have a CA IT Asset Manager agent installed on it. This agent classifies the server into one of the following groups:

- Server needs the patch
- Server does not need the patch
- Server does not meet the prerequisites for patch delivery

CA Patch Manager creates these groups automatically and updates them, so that CA Server Automation knows which servers are available for patch deployment. If the agent is not installed and you attempt to deliver a patch, CA Server Automation returns an error message. The message notifies you that the state of the server is unknown and the patch cannot be delivered.

Using Generic Groups and Templates

When you discover or provision a new system, your next step is likely to involve deployment of software packages, monitoring certain metrics, and applying rules. These operations are typically performed individually for each system. The process is as follows:

- Deploy each software package one at a time. Know exactly which packages to deploy and the order of installation of each package.
- Go to the Metrics page and select the metrics you want to monitor. Know exactly which metrics to select.
- Manually define the exact rules, statements, and actions you want to apply to the server. Know how to create rules, how and which statements to create, and which action to take.

This process can be inefficient, especially if you are configuring multiple systems. CA Server Automation provides a generic way to group software packages in one group. Using these groups, you can apply them to a server or service at once. Additionally, you can link the different entities together to form a template so that the same software delivery packages can be applied to other systems.

When grouping software packages together, be sure to create the package groups using software packages for the same operating system. Package groups do not support the capability to deploy packages for different operating system types in one package group. For example, create one package group to deploy to Windows XP, another package group to deploy to Linux, and so on.

Note: For information about creating generic groups and templates, see the *Online Help*.

Software Delivery Configuration File

The casdaconf.cfg is shipped and configured with a default version to auto-deploy the Asset Management agent only. Additional agents can be deployed individually, as part of a CA Server Automation package group, as part of a CA Server Automation software template, or as part of a provisioning action when you use the Additional Mgmt Agents option in the user interface.

The casdaconf.cfg file is located in the *Install_Path*\CA*productname*\conf directory. Use a text editor to edit this file.

Configuration parameters include the following:

CONFIG_KEY_SDA_LoggerCategory=sdadapter

Defines the logger category for the software delivery adapter.

CONFIG_KEY_SDA_HTTP_Protocol

Defines the protocol for the software delivery adapter host.

Default: https

CONFIG_KEY_SDA_Port

Defines the listening port for the software delivery adapter host.

Default: 443

CONFIG_KEY_SDA_PackageList_Sync_Interval

Specifies the time interval for software delivery adapter synchronizing the package list from the software delivery server. The software delivery adapter polls the software delivery server package or procedure group catalog and synchronizes this list with a list maintained by CA Server Automation. This attribute sets the frequency in which the two lists are synchronized.

Default: 43200000

Limits: milliseconds

CONFIG_KEY_SDA_ImageList_Sync_Interval

Specifies the time interval for software delivery adapter synchronizing the OS image list from the software delivery server.

Default: 600000

Limits: milliseconds

CONFIG_KEY_SDA_Imaging_job_Sync_Interval

Specifies the time interval for software delivery adapter synchronizing the OSIM imaging job status from the software delivery server.

Default: 360000

Limits: milliseconds

CONFIG_KEY_SDA_Packaging_job_Sync_Interval

Specifies the time interval for software delivery adapter synchronizing the software package job status from the software delivery server.

Default: 30000

Limits: milliseconds

CONFIG_KEY_SDA_Sync_Interval

Specifies the time interval to prevent WS timeout.

Default: 300000

Limits: milliseconds

CONFIG_KEY_SDA_CCM_Run_System_Discovery_Profile_Initial_Delay

When the CCA agent is successfully installed through a package job submitted by the software delivery adapter, a discovery profile is run. This attribute sets the number of seconds that elapse before the discovery profile runs after the software delivery adapter detects that the CCA agent was successfully installed.

Default: 2000

Limits: milliseconds

CONFIG_KEY_SDA_CCM_Run_System_Discovery_Profile_Max_Retry_Times

Specifies the number of retry attempts to run a discovery profile when an error is detected contacting the CA Configuration Automation web service.

Default: 5

CONFIG_KEY_SDA_CCM_Run_System_Discovery_Profile_Retry_Time_Interval

Specifies the time interval between attempts to run the Run System Discovery Profile processing when an attempt fails.

Default: 60000

Limits: milliseconds

CONFIG_KEY_SDA_Stage_SD_Agent_Time_Out

Specifies the time-out period for staging the software delivery agent to a scalability server.

Default: 360000

Limits: milliseconds

CONFIG_KEY_SDA_New_Package_Entry_State

Specifies the default entry state for new software packages that are added to CA Server Automation after the software delivery adapter detects their presence. Valid values are unmanaged or managed.

Default: unmanaged

CONFIG_KEY_SDA_Agent_Check_Retry_Count

Specifies the number of retry attempts when using Common Application Framework (CAF) to verify if the software delivery agent is installed.

Default: 3

CONFIG_KEY_SDA_DSM_URL

Defines the URL that the packaging component uses to connect to the CA DSM web services.

Example: DSM_URL=http://localhost/UDSM_R11_WebService/mod_gsoap.dll

CONFIG_KEY_SDA_Max_Img_Jobs

Specifies the maximum number of imaging jobs permitted to run simultaneously.

Default: 20

CONFIG_KEY_SDA_Img_Job_Rrtry_Delay

Specifies the delay for retrying image job.

Default: 600000

Limits: milliseconds

CONFIG_KEY_SDA_Img_Job_Max_Retry_Count

Specifies the maximum number of times to retry a failed imaging job.

Default: 3

CONFIG_KEY_SDA_Img_Job_Queue_Sync_Interval

Specifies the time interval that the software delivery adapter updates the imaging job queue.

Default: 120000

Limits: milliseconds

CONFIG_KEY_SDA_CLI_TIMEOUT

Specifies the default timeout value for the CLI.

Default: 60

Limits: minutes

CONFIG_KEY_SDA_PROVISIONING_TIMEOUT

Specifies the default timeout value for CLI OSIM job.

Default: 120

Limits: minutes

CONFIG_KEY_SDA_SCREG_RETRY_MAX_COUNT

Specifies the maximum number of attempts to register the software delivery component to the CA Server Automation service controller.

Default: 360

CONFIG_KEY_SDA_Img_Job_Pending_Timeout_Value

Specifies the default timeout value when OSIM is in a pending state.

Default: 60

Limits: minutes

CONFIG_KEY_SDA_Img_Job_Pending_Timeout_Retry_Count

Specifies the maximum number of attempts when OSIM is in a pending state.

Default: 3

CONFIG_KEY_SDA_Img_Job_Pending_Timeout_Failout={Yes|No}

Specifies to stop the OSIM job after the OSIM pending timeout expires.

Default: No

CONFIG_KEY_SDA_Img_Job_Progress_Timeout_Value

Specifies the default timeout value (in minutes) for OSIM progress state.

Default: 120

CONFIG_KEY_SDA_Img_Job_Progress_Timeout_Retry_Count

Specifies the maximum number of retry attempts when OSIM is in progress and times out.

Default: 2

CONFIG_KEY_SDA_Img_Job_Progress_Timeout_Failout={Yes|No}

Specifies if the OSIM job stops after the OSIM progress timeout expires.

Default: Yes

CONFIG_KEY_SDA_Img_Cancel_Job_In_Progress={Yes|No}

Specifies if an OSIM job that is already in progress can be canceled.

Default: Yes

CONFIG_KEY_SDA_ITCM_SESSIONPOOL_SIZE

Maximum number of sessions in the CA ITCM session pool for each ITCM server.

Default: 10

CONFIG_KEY_SDA_ITCM_RENEW_SESSIONPOOL_INTERVAL

CA ITCM session pool renew interval (milliseconds)

Default: 600000

Changing Agent Versions

The versions of the agents being used in your environment may differ from the defaults provided because they are dependent on the version of CA IT Client Manager that you are using. If this situation occurs, change the versions of the agents when deploying additional management agents. This process provides steps to change the version of agents listed in the casdaconf.cfg file.

After you change and save the file, restart the Apache HTTP Server for the changes to take effect.

The casdaconf.cfg file contains attributes that tell which packages to install for a particular operating environment. These attributes are used for deploying Additional Mgmt Agents and have the following format:

```
AutoDeploy:<platform>:<SD packageinfo>[index]
```

Descriptions are provided for the casdaconf.cfg file entries that retrieve package information for deploying Additional Mgmt Agents.

platform

Valid values include the following CA IT Client Manager platforms:

AUTODEPLOY: WINDOWS_X86

AUTODEPLOY: LINUX_X86

AUTODEPLOY: HPUX_HP

AUTODEPLOY: AIX_AIX

AUTODEPLOY: SOLARIS_SPARC

AUTODEPLOY: SOLARIS x86

SD packageinfo

Identifies which package to run. When you deploy a package, the following entries are used to describe the package:

SDA_PKG_NAME

Defines the name of the package.

SDA_PKG_VERSION

Defines the version of the package.

SDA_PKG_PROCEDURE

Defines a procedure to run the installation package.

index

Determines which package is installed first and groups package and procedure information to identify what to install.

Default entries for agents that are supplied in the casdaconf.cfg file. Modify these entries so that they work with the version of packages that are used in your environment. If the index is correct for all entries, you can add and remove entries. The index must start at 1 and additional entries must be sequential; you cannot use 1 and then 3.

Example: Define Auto Deploy Profile to Deploy a CCA Agent and a Performance Agent Sequentially

This example shows you how to define the Auto Deploy profile to deploy the CCA Agent first, and the Performance Agent second on a Windows system:

- AUTODEPLOY: WINDOWS_X86:SD_PKG_NAME1=CCA Agent Win32
- AUTODEPLOY: WINDOWS_X86:SD_PKG_VERSION1= r5.0
- AUTODEPLOY: WINDOWS_X86:SD_PKG_PROC1= CA_ACM_Windows_Agent_Install_VM
- AUTODEPLOY: WINDOWS_X86:SD_PKG_VERSION2= 12.0
- AUTODEPLOY: WINDOWS_X86:SD_PKG_PROC2=Install

Additional Provisioning Information

You can find more information about application and operating system deployment on the CA Support Online website. The following information provides you with advanced CA Desktop and Server Management topics and system requirement information.

- *CA Desktop and Server Management Advanced Topics* guide:

  https://support.ca.com/phpdocs/0/common/impcd/r11/troubleshooting/doc/DSM_Adv_topics_r11.pdf

- *CA Desktop and Server Management Green Books*:

  https://support.ca.com/phpdocs/7/common/greenbooks/CA_Unicenter_Desktop_Server_Management_Green_Book_ENU.pdf

- The following information identifies the operating systems and versions that are supported for provisioning:

  Windows and Linux:

  CA Software Delivery compatibility matrix:

  https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/234/234_compmatrix.html

  CA IT Client Manager compatibility matrix:

  https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/8176/8176_compmatrix.html

  Note: For specific information about supported operating environments and hardware requirements, see the *Release Notes*.

VMware vCenter Provisioning

The imaging service communicates with the VC PMM, which integrates with VMware vCenter Server to provide VM provisioning for CA Server Automation. This integration permits the imaging service to perform operations on VMs such as powering on, powering off, setting shares for CPU or memory, migrating a VM, or creating and provisioning a new VM.

This documentation assumes that you are familiar with VMware vCenter Server. All requirements and restrictions imposed by VMware on vCenter Server are valid for CA Server Automation.

Note: For documentation about VMware vCenter Server and vSphere, see the website http://pubs.vmware.com.

VMware vCenter and vSphere User Permissions

All VMware vCenter operations that CA Server Automation supports require that you configure user permissions in the VMware product. The minimum requirements are described here. Some users may require additional permissions depending on their role. These permissions apply to VMware vCenter versions 2.5, 4.0, and 4.1.

Permissions are set in the Edit Role dialog under All Privileges, Virtual Machine.

Users are located on the Administration, Configuration, Provisioning: vCenter Server page. After configuring permissions, perform the following actions to ensure that VM provisioning does not fail:

- Add the users to the hosts and clusters level in VMware. See http://kb.vmware.com/kb/1008192 for details.

Note: For more information about setting permissions, see the VMware documentation.

Power Operation

To let users perform power operations (Adjust VC Power Action), select the following options under Interaction:

- Power on
- Power off
- Suspend
- Reset

Resource Adjustment

Resource adjustment permissions let users perform the following operations:

- Set CPU Limit
- Set CPU Reservation
- Set Memory Limit
- Set Memory Reservation

To configure resource adjustment permissions, select Change Resource under Configuration.

Share Adjustment

Share adjustment permissions let users perform the following operations:

- Set Memory
- Add Memory
- Subtract Memory
- Set CPU
- Add CPU
- Subtract CPU

To configure share adjustment permissions, select the following settings under Configuration:

- Change CPU Count
- Memory
- Settings

Delete VM

Delete VM permissions let users delete virtual machines (Delete VC Virtual Machine).

To configure delete VM permissions, select Remove under Inventory.

VM and Template Conversion

VM and template conversion permissions let users convert templates to VMs and VMs to templates.

To configure VM and template conversion permissions, select the following settings:

under Provisioning

- Mark As Template
- Mark as Virtual Machine

under Resource

- Assign Virtual Machine to Resource Pool

Cloning

Cloning permissions let users perform cloning operations including VM to template cloning or from VM to VM.

To configure vCenter cloning permissions, select the following settings:

under Inventory

- Create
- Remove
- Move

under Interaction

- Power On
- Power Off
- Suspend
- Reset
- Device connection
- Tools Install

under Configuration

- Add Existing Disk
- Add New Disk
- Change CPU Count
- Memory
- Modify Device Settings
- Settings
- Change Resource

under Provisioning

- Customize
- Clone
- Deploy Template
- Mark as Template
- Mark as Virtual Machine
- Read Customization Specifications
- Modify Customization Specification

under Resource

Assign Virtual Machine to Resource Pool

VMotion

VMotion permissions let users migrate VMs from one physical server to another.

To configure VMotion permissions, select the following settings under Resource:

- Migrate
- Relocate
- Query VMotion

VM Snapshot

VM snapshot permissions let users take snapshots.

To configure VM snapshot permissions, select the following settings under State:

- Create Snapshot
- Revert To Snapshot
- Remove Snapshot
- Rename Snapshot

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 Server Automation. 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 Server Automation does not support the custom specification network setting "Prompt User." Custom Specifications that use this setting will be filtered out and unavailable.

   Click Next.

9. Click Add Computer.

   A confirmation message appears at the top of the pane.

   Note: Imaging takes time, so you should expect a delay during operating system installation. For more efficient discovery, you can adjust the discovery retry time or the interval in the caimgconf.cfg file located at: *install_path*\CA*productname*\conf.

10. Click Refresh to see the new VM in the left pane.

    Your data center has a new cloned VM. You can view the events of the imaging process in the dashboard and you can generate an imaging job report.

Add a Virtual Machine (Hyper-V Server)

You can create a VM to your data center. You must use a predefined template to create a VM.

Note: The value for Hyper-V "Total Storage" includes the total space required to create the VM from the template. This value is a combination of several factors that include all virtual disks, RAM size for the VM, snapshots, and a buffer. Use this information as guidance for the maximum amount of storage required to create a VM based on the template selected.

To create a VM

1. Select Resources, Explore.

   The Explore pane appears.

2. Right-click a Hyper-V resource, and select Provisioning, Provision Hyper-V VM.

3. Specify the following details and click Next.

   - SCVMM server and the Hyper-V server.
   - Template name that you want to use to create a VM.
   - Destination path where you want to create the VM.
   - Name of the VM that you want to create.
   - Specify whether to start the VM after it is created.

   The Virtual Machine Memory page appears.

4. Specify the VM memory details and click Next.

   The Guest OS Customization page appears.

5. Specify the guest operating system details and click Next.

   The Network Management page appears.

6. Specify the network details of the VM and click Next.

   Note: If your custom specification specifies the use of DHCP, you will only be able to edit the network connection cell in the table. If your custom specification specifies the use of a static IP address, you will be able to edit all cells except the NIC cell. CA Server Automation does not support the custom specification network setting "Prompt User." Custom Specifications that use this setting will be filtered out and unavailable.

7. Click Add Computer.

   A confirmation message appears at the top of the pane.

   Note: Imaging takes time, so you should expect a delay during operating system installation. For more efficient discovery, you can adjust the discovery retry time or the interval in the caimgconf.cfg file located at: *install_path*\CA*productname*\conf..

Add a Solaris Zone

Solaris Zones servers use Zones to run applications in isolated environments to make it appear as if they are running on physically separate systems. Each Zone on a server takes its resources from a resource pool and includes virtual network interfaces, file systems, memory, and other dedicated units. When you create a Zone, you must supply all of this information. The Zone installs automatically after creation.

To add a Solaris Zone

1. Select the Resources tab, right-click the Zone Host in the Explore pane, and select Provisioning, Provision Zone.

   The Solaris Zone Provisioning wizard appears.

2. Complete the following fields on the Zone Identity and Type page and click Next:

   Host

   Defines the host on which to create the Zone.

   Name

   Defines the name of the Zone.

   Description

   (Optional) Defines a description of the Zone.

   Type

   Defines whether the Zone is Native or Branded. A Branded Zone is based on an existing Zone template.

   Template Name

   (Optional) Defines the template from which to create the Zone when you set Type to Branded.

   Install Archive Path

   Defines the directory path of the installation archive on the Zone. This field is only required if you set Type to Branded.

   The CPU, Memory, and Additional page appears.

3. Complete the following fields and click Finish:

   Type

   Defines the scheduler type. Select FSS to use the Fair Share Scheduling class to control CPU allocation based on the number of CPU shares assigned to tasks.

   Capacity

   Defines the amount of physical memory capacity to allocate to the Zone, in megabytes.

Swap Memory

Defines the amount of swap memory to allocate to the Zone, in megabytes. The swap memory must be at least 50 MB.

Lock Memory

Defines the amount of lock memory to allocate to the Zone, in megabytes. The lock memory must be less than the physical memory.

Zone Path

Defines the root directory path of the Zone.

NIC Type

(Optional) Defines the NIC type. Select a type from the drop-down list. If you do not select a NIC, the Zone is not assigned a NIC card or IP address.

IP Address

(Optional) Defines the IP address of the Zone.

Resource Pool

Defines the resource pool to use with the Zone. Select a pool from the drop-down list. If you want to use a new resource pool with the Zone, create the pool first. If you do not select a pool, the default is used.

Auto Reboot

Defines whether to reboot the Zone automatically when the global Zone is rebooted.

Add a Logical Partition for an IBM AIX Computer

You can use the Provisioning wizard to create logical partitions for an IBM AIX system.

To add a logical partition for an IBM AIX computer

1. Click Resources.
2. Right-click IBM PowerVM Server in the Explore pane, and select Provisioning, Provision LPAR.

   The Provisioning wizard appears with the Partition and Memory page.

3. Select the HMC/IVM server and managed system name. Specify the partition name and, if using an HMC server, the profile name. Specify the minimum, desired, and maximum memory for the partition. Click Next.

   The Processor page appears.

4. Specify the minimum, desired, and maximum number of processors, and indicate whether to allocate a shared or dedicated processor. Advanced settings are available for shared modes and virtual processors. Click Next.

   The I/O Components page appears.

5. Select the I/O devices to associate with the partition, and click Next.

   Note: For each I/O device, you can specify that the I/O device is required or optional for logical partition activation. If the I/O device is required, the partition cannot be activated if the I/O device is unavailable or in use by another logical partition. If the I/O device is optional, and if the desired I/O device is available when the partition is activated, the managed system commits the I/O device to the partition. If the optional I/O device is not available, the managed system skips the I/O device.

   The I/O Pools page appears.

6. (Optional) To create a new I/O pool, click the + (Add) on the I/O Pools table, enter a numerical value, and click Save.

   Note: When you add an I/O device to a partition, and the I/O device belongs to an I/O pool. When this partition is activated, the managed system automatically adds the I/O pools defined for the partition to the logical partition.

7. Click Next.

   If an HMC server was selected, the Virtual Serial page appears.

8. (Optional) Specify the maximum virtual adapters for the partition. To create a new virtual serial adapters, click + (Add) and specify the Adapter ID, Remote Partition, and Remote Slot Number. You can require that the virtual adapter must be allocated and the managed system must have enough memory to run the required virtual adapters for the partition profile, or the logical partition does not activate.

9. Click Next.

   The Virtual Ethernet page appears.

10. Specify the maximum virtual ethernet adapters for the partition. (Optional) You can add new virtual ethernet adapters by clicking + (Add) and selecting an Adapter ID, Virtual LAN ID, Access External Network, Trunk Priority, IEEE 802.1 Q Compatibility, additional Virtual LAN IDs, and whether the Ethernet adapter is required.

11. Click Next.

    The Virtual SCSI page appears.

12. Specify the virtual SCSI devices for the partition. (Optional) To add a new virtual SCSI adapter, click + (Add) on the Virtual SCSI Adapters table.

    Select an Adapter ID, specify whether the SCSI adapter is Required, and pick a Device name from the SCSI Devices table. If the desired device is on the SCSI Devices list, click OK, click Next in the Virtual SCSI panel, and skip to the last step. To add a new SCSI backing device, click + (New Backing Device) on the SCSI devices table.

    Note: If the selected device has a slot number, that is the slot number of the virtual SCSI server adapter defined to the Virtual I/O server partition.

    If the selected device doesn't have a slot number, that means it isn't associated with a virtual SCSI server adapter yet. When the job to create the partition takes place, the virtual SCSI server adapter will be created and assigned to the device.Click Next.

    The Summary page appears.

13. Verify the Summary and click Add Computer.

    The logical partition is created.

Additional Provisioning Information

You can find more information about application and operating system deployment on the CA Support Online website. The following information provides you with advanced CA Desktop and Server Management topics and system requirement information.

- *CA Desktop and Server Management Advanced Topics* guide:

  https://support.ca.com/phpdocs/0/common/impcd/r11/troubleshooting/doc/DSM_Adv_topics_r11.pdf

- *CA Desktop and Server Management Green Books*:

  https://support.ca.com/phpdocs/7/common/greenbooks/CA_Unicenter_Desktop_Server_Management_Green_Book_ENU.pdf

- The following information identifies the operating systems and versions that are supported for provisioning:

  Windows and Linux:

  CA Software Delivery compatibility matrix:

  https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/234/234_compmatrix.html

  CA IT Client Manager compatibility matrix:

  https://support.ca.com/irj/portal/anonymous/phpdocs?filePath=0/8176/8176_compmatrix.html

  Note: For specific information about supported operating environments and hardware requirements, see the *Release Notes*.

JumpStart

The JumpStart adapter integrates with CA Server Automation so that you can provision Solaris systems with images. The JumpStart adapter searches the system for information to determine what images are available on the server. You can select one of the available images and submit an imaging job against a specific target computer.

The CA Server Automation documentation assumes that you are familiar with the JumpStart solution. All requirements and restrictions imposed by Oracle on JumpStart server technology are valid for CA Server Automation.

JumpStart Prerequisites

The prerequisites for using the JumpStart solution include:

- The server must run on Solaris 10 SPARC or Solaris 10 x86.
- JumpStart boot servers must exist on the same subnet of each SPARC client to be imaged.
- All host names must be configured as static with the DNS server regardless of client architecture.
- Protocols used for configuration cannot be blocked by firewalls.
- The JumpStart solution requires RARP/BOOTP/TFTP protocols for SPARC systems.
- The JumpStart solution requires PXE/DHCP/TFTP for Solaris 10 x86 based systems.
- SPARC and Solaris 10 x86 systems use Network File System (NFS) to access remote JumpStart installation images.
- The full Solaris 10 media is required for server installation.

- Initial installation only is supported, but upgrades of Solaris images are not supported.
- Only one JumpStart configuration directory and one profile/rules directory per supported Solaris version is allowed.
- SPARC client computers must already be running Solaris to be reimaged for dynamic use.
- SPARC client computers must be preconfigured for Secure Socket Shell (SSH) root access to allow rebooting.
- Solaris 10 X86 client computers must have service processors that are Intelligent Platform Management Interface (IPMI) 1.5 or 2.0 compatible.
- Each service processor must have a static IP address configured.
- The IPMI feature for the service processor must be enabled and configured in the BIOS.
- Each service processor must be preconfigured so that it can be reached from the public network.
- The DHCP server must be configured for SSH access to allow the CA Server Automation JumpStart adapter to communicate with it during JumpStart x86 provisioning requests.

IBM AIX Provisioning with NIM

NIM provisioning uses the Network Installation Manager (NIM) to provide AIX OS imaging services for IBM PowerPC-based Logical Partitions and physical computers. The NIM adapter resides on the NIM master, where you set up and maintain your NIM environment. The OS version running on the NIM master must be the same OS version or higher than the versions that it deploys to NIM clients.

You can only provision IBM systems that are already defined in the NIM environment and configured with AIX. You cannot provision physical (bare metal) computers that do not have an OS installed.

The CA Server Automation user interface displays the status of IBM AIX provisioning jobs that the NIM adapter monitors. Configure NIM to monitor status and deploy images.

This documentation assumes that you are familiar with NIM. All requirements and restrictions imposed by NIM on IBM AIX are valid for CA Server Automation.

Note: For more information about NIM and IBM AIX, see the IBM Redbooks on the IBM website, http://www.redbooks.ibm.com/. *NIM From A to Z in AIX 5L* is a good resource on the IBM Redbooks website.

Prerequisites

The following list comprises NIM environment requirements and restrictions:

- Each NIM master server must have a NIM adapter installed.
- Alternate NIM masters are not supported.
- NIM master servers must be available for logon using ssh.
- NIM clients can only be registered to one NIM master.
- NIM clients must be registered with the NIM master.
- NIM clients must be configured with TCP/IP.
- NIM clients must initially be configured to permit rsh (unsecure) or nimsh (secure) communication with the NIM master.
- nimsh clients must be running on the client computer if the NIM machine resource defines the nimsh protocol.
- All imaging jobs are performed with the option "Remain NIM client after install" set, so that NIM clients can be reimaged without additional configuration.
- NIM clients must be on the same subnet as the NIM master, due to TFTP network restrictions.
- NIM clients must be capable of being imaged by the NIM master using the NIM command line interface.
- The NIM master must be configured to be able to qualify all of its NIM clients' host names.
- When provisioning to a new or selected LPAR, the NIM adapter updates the NIM system MAC address to the LPAR virtual Ethernet adapter MAC address.

Add an IBM AIX Client System Using a Resource Group

CA Server Automation integrates with NIM, so that you can image a client computer with an IBM AIX operating system using a resource group.

To add an AIX client system using a resource group

1. Right-click an IBM PowerVM resource, and select Provisioning, Provision NIM.

   The Provision AIX with NIM dialog appears.

2. Complete the following fields:

   NIM Master

   Defines the computer where you set up and maintain your NIM environment. The NIM environment is a logical group of computers. Multiple NIM environments can be on the same TCP/IP network, but there can only be one active NIM master per environment.

   Machine Resource Name

   Defines the name of the target computer for NIM installations and update operations.

3. Select Resource Group from the Resource Type menu, and select a resource group from the drop-down list.

   Note: If you select Resource Group, individual Resources fields are not displayed.

4. Complete the remaining fields, and click Add Computer.

   Resource Group

   Defines a logical grouping of resources used for assigning resources to a NIM client more quickly than an individual association.

   Username

   Defines the root user used for agent deployment.

   Password

   Defines the root user password used for agent deployment.

   Use Logical Partition

   Specifies an IBM LPAR as the system to image instead of a physical system, when selected. The system does not need to have an operating system running at the time of imaging.

   HMC/IVM Name

   Defines the HMC or IVM server that manages one or more managed servers.

   System Name

   The name of the managed system assigned in the HMC or IVM server. This system hosts LPARs and contains the physical hardware that is virtualized so that LPARs can use the resources, such as SCSI and Ethernet adapters.

   Partition Name

   Defines the existing LPAR that will be used as the target system to image.

   Profile Name

   (HMC only) Defines the existing partition profile containing information about how to initialize the selected LPAR.

Template

Lists the software package groups already created and available for use as templates.

Domain Manager

Defines the name of the ITCM Domain Manager where the software delivery adapter resides and where the operation should be performed. This name is optional when only one SD adapter or ITCM domain manager is configured. This parameter is valid only for CA Server Automation.

DSM Boot/Scalability Server

Serves as the primary interface for the agent and distributes the CA Unicenter DSM workload across multiple hosts. Software delivery adapter supports multiple scalability servers for software distribution.

The imaging process starts on the client computer and a confirmation message notifies you when the imaging job has completed successfully. When the imaging process is complete, the computer is properly discovered and classified.

Add an IBM AIX System Using an Individual Resource

CA Server Automation integrates with NIM, so that you can image a client computer with an IBM AIX operating system using an individual resource.

To create an AIX client computer using an individual resource

1. Right-click an IBM PowerVM resource, and select Provisioning, Provision NIM.
2. Complete the following fields:

   NIM Master

   Defines the computer where you set up and maintain your NIM environment. The NIM environment is a logical group of computers. Multiple NIM environments can be on the same TCP/IP network, but there can only be one active NIM master per environment.

   Machine Resource Name

   Defines the name of the target computer for NIM installations and update operations.

3. Select Individual Resources from the Resource Type menu, complete the remaining fields, and click Add Computer.

   MKSYSB Image

   (Required for MKSYSB installation type) Specifies MYSYSB image for cloning.

Base Operating System Instance

Specifies a file that contains information for the Base Operating System (BOS) installation.

Licensed Program Products

Defines the licensed program product files to use for an imaging request.

Shared Product Object Tree

Defines the shared product object tree to use for an imaging request.

Resolve Configuration

(Optional) Defines a file that contains valid */etc/resolv.conf* entries that define Domain Name Protocol name-server information for local resolver routines.

First Boot Script

(Optional) Defines the name of the file to use to configure devices when a NIM client is booting for the first time after the BOS installation process is completed.

Post Install Scripts

(Optional) Defines a list of scripts to run after installation.

Username

Defines the root user used for agent deployment.

Password

Defines the root user password used for agent deployment.

Use Logical Partition

Specifies an IBM LPAR as the system to image instead of a physical system, when selected. The system does not need to have an operating system running at the time of imaging.

HMC/IVM Name

Defines the HMC or IVM server that manages one or more managed servers.

System Name

The name of the managed system assigned in the HMC or IVM server. This system hosts LPARs and contains the physical hardware that is virtualized so that LPARs can use the resources, such as SCSI and Ethernet adapters.

Partition Name

Defines the existing LPAR that will be used as the target system to image.

Profile Name

(HMC only) Defines the existing partition profile containing information about how to initialize the selected LPAR.

Template

Lists the software package groups already created and available for use as templates.

Domain Manager

Defines the name of the ITCM Domain Manager where the software delivery adapter resides and where the operation should be performed. This is optional when only one SD adapter or ITCM domain manager is configured. This parameter is valid only for CA Server Automation.

DSM Boot/Scalability Server

Serves as the primary interface for the agent and distributes the CA Unicenter DSM workload across multiple hosts. Software delivery adapter supports multiple scalability servers for software distribution.

The imaging process starts on the client computer, and a confirmation message notifies you when the imaging job has completed successfully. When the imaging process is complete, the computer is discovered and classified.

LPAR Provisioning for IBM AIX

LPAR provisioning uses the Imaging Service to manage the logical partitions on an IBM PowerVM system. The Imaging Service sends requests to the LPAR PMM, which connects to the HMC or IVM server and issues commands to accomplish the requested actions.

The CA Server Automation user interface displays the IBM AIX provisioning job status that the LPAR adaptor monitors. Configure LPAR to monitor status and deploy images.

This documentation assumes that you are familiar with LPAR requirements. All requirements and restrictions imposed by NIM on IBM AIX also are valid for CA Server Automation.

Note: For more information about NIM and IBM AIX, see the IBM Redbooks on the IBM website, http://www.redbooks.ibm.com/. *NIM From A to Z in AIX 5L* is a good resource on the IBM Redbooks website.

iSCSI Storage Provisioning Prerequisites for LPAR

LPAR prerequisites for using iSCSI storage provisioning include:

- The default iSCSI Adapter initiator ID on IBM AIX does not conform to RFC 3720; it is in a generic format that is not unique, so the administrator must verify that the default iSCSI Adapter initiator ID is unique before provisioning storage. Do this on the Virtual I/O server as follows:

  ```
  chdev -dev iSCSIAdapterDeviceName -attr initiator_name=UNIQUE-ID
  ```

  iSCSIAdapterDeviceName is the adapter device name. For example, iscsi0.

  UNIQUE-ID must conform to RFC 3720 to help ensure uniqueness. For example, iqn.2011-06.com.domain:storage:myMachine.0xa1afc0f.800010EBCD900004

- Virtual I/O Servers must be available for login using ssh.

Amazon EC2 Provisioning

CA Server Automation lets you provision and manage Amazon Elastic Compute Cloud (EC2) instances in a public or private cloud.

Supported Features

CA Server Automation supports the following features in the Amazon EC2 operating environment:

Full Support

- On-Demand Instances
- Amazon Virtual Private Cloud
- Multiple Locations (Availability Zones and Regions)

  All Amazon regions are supported. When Amazon site adds new regions, the CA Server Automation user interface dynamically downloads them. You can provision Amazon Machine Instances (AMIs) and perform operations for each configured region.

Partial Support

- Amazon Elastic Block Storage (EBS)

  You can stop or restart EBS instance. When you stop an EBS instance, data and state are saved, and charges are avoided. When started, an EBS instance resumes the previous state and data. You can create an AMI from a running or stopped EBS instance.
- One Amazon Web Service (AWS) account in CA Server Automation
- Amazon EC2 instance types
  - t1.micro
  - m1.small
  - c1.medium
  - m1.large
  - m1.xlarge
  - c1.xlarge
  - m2.xlarge
  - m2.2xlarge
  - m2.4xlarge
  - cc1.4xlarge

Prerequisites

Review the following prerequisites before configuring Amazon EC2 for CA Server Automation:

- You are familiar with the Amazon EC2 operating environment
- The CA Server Automation Amazon EC2 Adapter is installed, and you can access the CA Server Automation user interface
- You have a basic understanding of the CA Server Automation user interface, and how to provision resources

How to Configure Amazon EC2 for Provisioning

CA Server Automation lets System Administrators provision AMIs from their Amazon EC2 operating environment. With help from a Network Administrator, a System Administrator can also connect users to compute resources from the Amazon Virtual Private Cloud (VPC).

The following diagram illustrates the process for configuring Amazon EC2 resources in CA Server Automation.

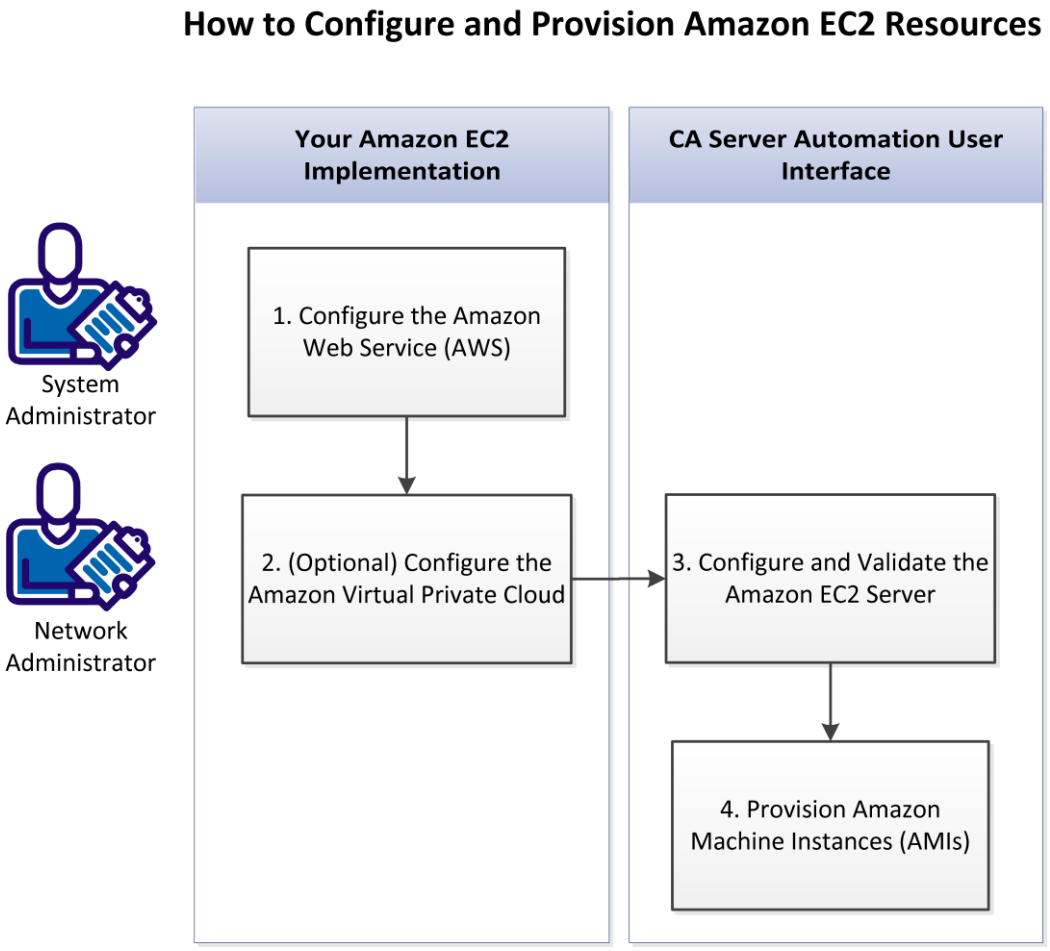

1. Configure the Amazon Web Service (AWS) (see page 418)
2. Configure the Amazon Virtual Private Cloud (see page 420)
3. Configure and Validate the Amazon EC2 Server (see page 421)
4. Provision Amazon Machine Instances (AMIs) (see page 422)

How to Configure the Amazon Web Service (AWS)

Use the following worksheet to verify configuration of your Amazon EC2 operating environment.

Tasks in Amazon Operating Environment	Notes	CA Server Automation User Interface Fields
Create an Amazon Web Services (AWS) account	Only one Amazon account is supported.	AWS Account ID:
Sign up for the Amazon EC2 service	This service includes the Amazon Simple Storage Service (public cloud) and the Amazon Virtual Private Cloud (VPC).	
Create the X.509 Certificate and Private Key	The certificate and private key are required credentials when configuring the Amazon EC2 server in CA Server Automation. Keep them in a safe place; if you lose them, generate a new certificate.	X.509 Certificate File Full Path: Private Key File Full Path:
(Optional) Collect proxy server information	A proxy server lets you monitor the Amazon EC2 instances in your network. This server is optional, but is often desirable for security reasons.	Proxy server name: Host Name: Port: User name: Password:

Create, launch, bundle, upload, and register the Amazon Machine Image (AMI)	The AMI Identifier is equivalent to an *image* in CA Server Automation. **Create security groups** Security groups are visible in the CA Server Automation interface during provisioning. **Create Key Pair** To launch an Amazon AMI in CA Server Automation, you need a named key pair. The name of the key pair is specified in the web service call that launches the instances. SSH uses the private key for authentication. To create a key pair, use the AWS Console or the ElasticFox Firefox plug-in. Save the returned private key in a safe place on your file system. **Create EBS volume** EBS volumes can be attached to any instance in the same Availability Zone. EBS volumes are displayed as a Device Type in the CA Server Automation user interface.	AMI ID: OS Type: Key Pair Name: (SSH Key Pair) Manifest: Instance Type: Security Group: Public Availability Zone: Device Type: EBS
	Determine Regions Regions (for example, East US Northern Virginia) are dispersed and located in separate geographic areas. Using dynamic updating, all Amazon regions are supported in CA Server Automation. Availability zones are distinct locations with a region for isolating failures. If you do not supply a zone for an instance, Amazon selects one.	Select Region

More information:

Add Proxy Servers (see page 40)

How to Configure the Amazon Virtual Private Cloud

CA Server Automation requires a VPN connection to access the Amazon Virtual Private Cloud (VPC). You can create AMIs within subnets for network isolation.

If you are implementing a VPC for CA Server Automation, a Network Administrator must complete the following tasks.

Tasks in Amazon Operating Environment	Notes	CA Server Automation User Interface Fields
Create subnets and IP address ranges	See the *Amazon Virtual Private Cloud Network Administration Guide.*	Subnet: Subnet IP Address: Domain Name:
Create the VPN connection	See the *Amazon Virtual Private Cloud Network Administration Guide.*	Private: Availability Zone:
Configure the gateway/router in CA Server Automation	Configure the gateway/router on the CA Server Automation Manager to access the AWS subnets.	

Configure and Validate the Amazon EC2 Server

The Amazon EC2 server must be configured and available in CA Server Automation. The connection from CA Server Automation to the Amazon public and private clouds must also be validated before you can provision instances.

Follow these steps

1. Copy the certificate and private keys to the CA Server Automation Manager.
2. On the server where the Amazon EC2 Adapter service is installed, start the CA Server Automation user interface.
3. Click the Administration tab.
4. Select the Configuration panel, Provisioning, EC2 Server.
5. Enter information for your implementation:
   - **AWS Account ID:** The AWS Account ID.
     Example: 538614568157
   - **X.509 Certificate Full Path:** The Amazon X.509 Certificate file (.pem).

     Example: C:\ec2_account\cert-XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX.pem
   - **Private Key File Full Path: T**he Amazon Private Key file (.pem).

     Example:C:\ec2_account\pk-YYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYY.pem
   - **Configured Regions**: The regions configured for this EC2 server.

     Example: us-east-1
   - **Default Region**: The default region to display for provisioning.

     Example: ap-northeast-1
   - **Proxy Server**: A server that acts as an intermediary for service requests.

     Example: yourproxy.com
6. On the CA Server Automation Manager, access the system clock application and synchronize the time with a Network Time Protocol (NTP) server.

   Note: For security reasons, the Amazon EC2 server invalidates connection attempts by an outside server, if the outside server time does not match its own time. Synchronizing time on the CA Server Automation Manager avoids a potential validation error in the next step.
7. From the Actions menu, select Validate.

   The connection between CA Server Automation and your Amazon public or private cloud is validated. If the Connection Status is green, you can start provisioning EC2 instances. If the Validation fails, verify the following before retrying the Validate action:
   - AWS Account ID is correct
   - Certificate and Private Key files are available and in the correct place

- You synchronized the system time on the CA Server Automation Manager
- Proxy credentials are correct (applicable only if you run CA Server Automation inside the firewall)

More information:

Add Proxy Servers (see page 40)

How to Provision Amazon Machine Image Instances (AMIs)

In CA Server Automation user interface, use the following wizard to provision AMIs: Provisioning, Add Resource, AMI. The following table provides additional help.

User Interface	Problem	Solution
Instance Type	**SOAP Error**: "The requested instance type's architecture (i386) does not match the architecture in the manifest for <*ami_name*> (x86_64)." You used a 64-bit image on a small instance type.	Create a large instance type for the 64-bit image.
	SOAP Error: "AMIs with an instance-store root device are not supported for the instance type 't1.micro'." You used a t1.micro instance type on a non-EBS AMI. Amazon does not support this configuration.	Create the t1.micro instance with an EBS AMI.

Provision AppLogic Applications

The Provision AppLogic Application wizard lets you provision an instance of an application template to an AppLogic grid.

Follow these steps:

1. Right-click a grid in the Explore pane, and select Management, Provision AppLogic Application.

   The Provision AppLogic Application wizard, General and Configure Resources pane opens.

2. Select the Grid Name to deploy to, the Template Name to create an instance of, and input an Application Name.

   The Constrain by Resources section populates with resource limits and default settings from the template.

3. Set the CPU, Memory, and Bandwidth values for the application, and click Next.

   The Config Boundary pane opens displaying the boundary properties from the template.

4. Set the value for each property and click Next.

   The Summary pane opens.

   Note: IP properties automatically populate based on grid properties.

5. Review the details for application provisioning and click Finish.

   The Provision AppLogic Application wizard submits a request to provision the application to the selected grid with the specified resource allocations.

Rapid Server Imaging

Rapid Server Imaging (RSI) provides cross-platform and heterogeneous hardware provisioning, physical and virtual server migration, disaster recovery, and image capture and deployment. Images can be deployed across dissimilar hardware with multiple operating environments provided the hardware belongs to the same processor family.

You can capture images from a managed server and deploy the captured images to a different managed server or to a bare metal system. You can capture and deploy images between physical-to-physical, physical-to-virtual, virtual-to-physical, and virtual-to-virtual systems. Except in cases of bare metal provisioning, the CA Server Automation managed servers must run the RSI agent.

For documentation on how to install and configure an RSI server, see DVD3 of the product media.

More information:

Best Practices (see page 424)
Deploy RSI Images to AppLogic (see page 425)

Best Practices

The following list comprises RSI best practices:

- Access the RSI Resources panel by selecting the Resources tab, right-click Data Center in the Explore tree, and select Management, RSI Resources.
- To deploy Windows images, capture the driver set and collect drivers from the vendor media using the Driver Set and Driver Collection tabs of the RSI Resources panel.

  When deploying all other operating environment images, the proper driver set is automatically captured in the image and assigned when the captured image is deployed.

  Note: For more information about capturing Windows driver sets, see the *Online Help* or the *Reference Guide*.
- Delete images and driver sets that you no longer need using the RSI Resources panel.
- For virtual systems, register hypervisors (VMware, Hyper-V) with your RSI servers. Registering virtual systems with a hypervisor enables the hypervisor platform power management to handle bare metal operations requiring an initial boot.
- When working with virtual machines in a Microsoft Hyper-V/RSI environment, with any supported OS, verify the following:
  - Each VM must have at least one Legacy Network Adapter that is connected to one of the DynaCenter boot networks. Having only one Legacy Network Adapter is the best. But, if you have more than one, verify that the first one is connected to a DynaCenter boot network.
  - Hyper-V Tools do not support the legacy network adapter on the 64-bit edition of Windows Server 2003; therefore, DynaCenter does not support this operating system. For a complete list of supported OS, see the support matrix in the installation guide.
  - Each VM must have one or more paravirtualized network adapters (used by OS when it is deployed to the VM).

Deploy RSI Images to AppLogic

Use CA Server Automation to deploy RSI images to an AppLogic grid.

Note: The AppLogic grid must have a deployed RSI server which shares the depot containing the images to deploy.

Follow these steps:

1. Right-click the AppLogic grid to deploy to, and select Provisioning, Deploy Image.
2. Select the RSI server in the grid, the shared depot, and the image to deploy. Set the boot network and domain according to the settings for the RSI server in the grid, and click Next.
3. Select the Grid Name to deploy to, the Template Name to create an instance of, and input an Application Name.

   The Constrain by Resources section populates with resource limits and default settings from the template.

   Important! The template must be an OS template matching the OS for the image.
4. Set the CPU, Memory, and Bandwidth values for the application, and click Next.
5. Review the details for image deployment to the select OS application and click Finish.

   The wizard submits a request to deploy the image to a new instance of the application template in the selected grid with the specified resource allocations.

Bare Metal Provisioning to a Cisco UCS Blade

You can use either the Rapid Server Imaging dpmrsi CLI or the Cisco UCS Provisioning dpmucs CLI for bare metal provisioning to a Cisco UCS blade. Consider the following information when deciding which method to use.

- Significant delays can occur during Cisco UCS blade provisioning, due to the blade power cycle and the time required to associate the blade with a specific service profile.
- For physical systems, RSI bare metal provisioning requires that both the network and the target system are configured for Wake-On-LAN to avoid manual intervention to boot the target system during deployment. Unless network routers are configured across subnets to allow directed broadcasts, Wake-On-LAN only works for systems that are in the same subnet as the CA Server Automation instance.
- Cisco UCS provisioning does not require Wake-On-LAN because it uses the Cisco UCS Manager for power control.

Remote Deployment

CA Server Automation provides a comprehensive solution for remotely deploying SystemEDGE and other agents to all managed systems. You can create deployment templates based on provided packages that contain customized installation parameters and simultaneously deploy these templates to numerous managed systems. This automated deployment solution provides one location from which to deploy and configure the agents throughout your enterprise.

Remote deployment provides the following features:

Deployment configuration

Allows creation, editing, and deletion of configurations which define how software packages are installed on target systems. These are referred to as Package Wrappers.

Deployment job management

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

Deployment job reporting

Provides the ability to query the status of deployment jobs.

Deployment events

Provide a source of deployment-related events that tracks the state of the managed nodes.

Remote Deployment Architecture

The overall architecture of the deployment solution is driven by the Domain Server and Distribution Server components. The following diagram represents an overview of the deployment-related components:

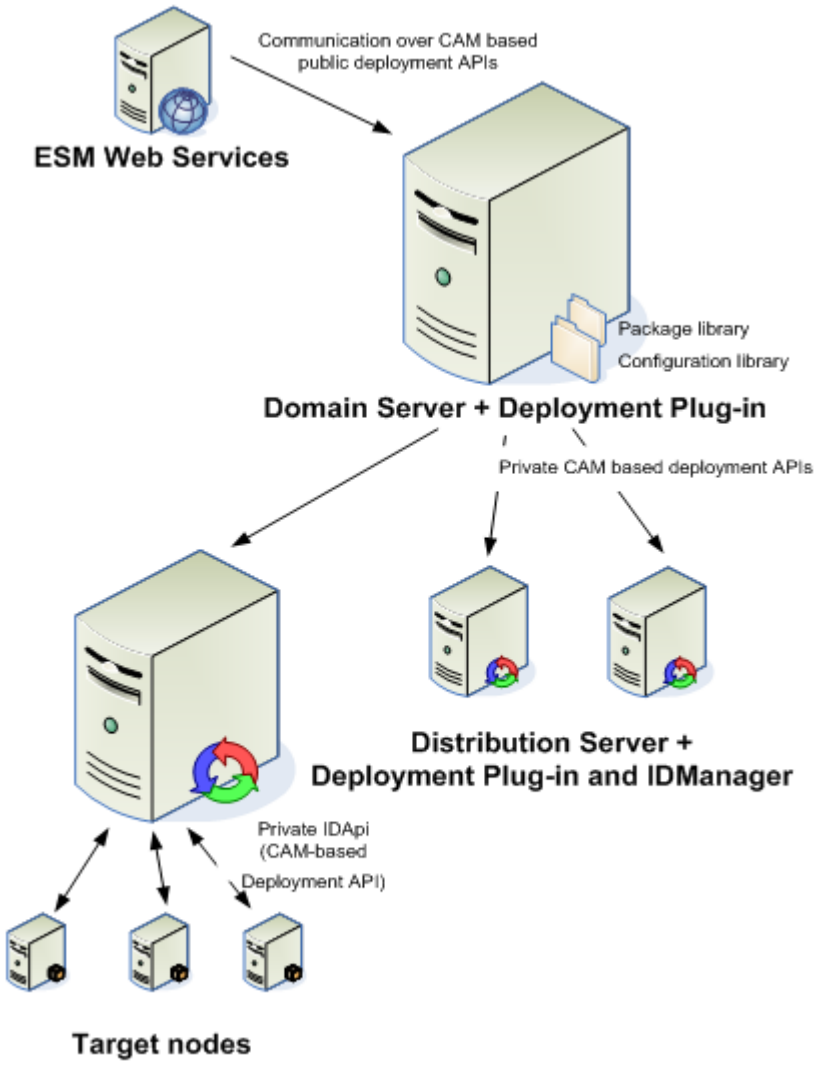

Deployment Components

This section lists and briefly describes the deployment key components:

Domain Server

The domain server is the repository for all configuration and control data. The server is responsible for managing configuration and software package data required for deployment operations and it manages all configuration operations. Detailed event data is passed between the domain and distribution servers during the deployment process. A single Domain Server is responsible for maintaining the status of all Distribution Server jobs.

Distribution Server

The distribution server controls the Infrastructure Deployment Manager (IDManager) server located on the same machine. The architecture allows for multiple distribution servers offering deployment services.

Infrastructure Deployment

Infrastructure Deployment initiates and manages deployment jobs. During the deployment process, the Infrastructure Deployment Manager (IDManager) provides access to remote systems and the Infrastructure Deployment Primer (ID Primer) provides the mechanism to remotely install agent software packages. The IDPrimer is used to transfer the deployment package data to the target computer and run the installation. All subsequent deployments to the same target computer can use the existing IDPrimer installation. The IDManager controls all the deployment operations and handles job status.

Using Multiple Distribution Servers

Although the remote deployment solution lets you have a single central server (manager machine) to manage all your deployments, it is recommended that you install a remote distribution server that points to the central domain server if you have any of the following requirements:

- You have two or more geographically remote locations where you need the agent software deployed to but need these managed centrally using a single manager machine.

  In this case we recommend that each location has at least one distribution server with the distribution server connects to the central domain server.

- You have a single location but have several hundreds of machines that you have the need to deploy to.

  In this case you can choose to install many distribution servers split logically across subnets and these distribution servers will connect to the central domain server.

Change the Domain Server a Distribution Server Connects To

In the situation where the network address of the domain server machine changes after the original installation, it is necessary to reconfigure the distribution server to connect to a new network address.

Prior to making the configuration change shown below it is important to verify that the new network address is connectable from the distribution server. If the distribution cannot make a connection to the domain server using the new address then deployment functionality will not work correctly.

To change the domain server a distribution server connects to

1. From the Start menu, open Administrative Tools, Services.

   The Services user interface appears, listing the installed services.

2. Right-click CA SM Distribution Server and select Properties.

   The Properties dialog appears.

3. Click Stop to stop the service.

4. Enter the following parameter into the Start parameters field:

   -m *domainserver*

   The *domainserver* parameter specifies the IP address or DNS name of the domain server.

5. Click Start

   The distribution server will now attempt to connect to the domainserver address entered.

Scalability

The deployment system provides a degree of scalability using multiple distribution servers as a scalability layer. Each distribution server communicates with one IDManager instance. The IDManager can manage multiple component deployments to multiple target computers. CA Server Automation supports many simultaneous deployments because of this federated model.

More Information

Remote Deployment and Policy Configuration Recommendations (see page 641)
Scalability Use Cases (see page 643)

Audit Trail

Jobs and tasks are the two fundamental concepts of the deployment system. A deployment job specifies one or more configurations to be installed on one or more target systems. Each individual installation of a software package on a target system is represented by a task. Deployment job reporting provides the ability to query the status of deployment jobs.

You can create, control and inquire the state of deployment jobs. When the job is started, its individual deployment tasks are delegated to available distribution servers which perform the actual deployment. You can track the progress of the job as it occurs, to verify that the deployment is going well and to quickly identify and correct any problems.

Remote deployment is able to provide the following information:

- Which deployment jobs are currently:
  - Inactive (not yet started)
  - Active
  - Completed, and of those, which were:
    - Successful
    - Partially successful
    - Unsuccessful
- Which deployment jobs are:
  - Associated with a specific target machine
  - Associated with a specific package/package group
- What packages have been deployed to a specific target machine
- Which user created/started deployment of a specific package
- Which machines are targeted in a specific deployment job
- Which machines are targeted by active deployment jobs

Note: Remote Deployment supports deploying software to UNIX/Linux systems with /tmp file system mounted with noexec flag.

The Jobs pane in the UI displays a table of all created deployment jobs that lists the job ID and description, the packages included, job status, and so on. From this table, you can drill down to view more details about a specific job, including why a job failed. Click the Jobs folder in the UI to access this view. For more information, see the *Online Help*.

Deployment Restrictions

Consider the following before performing a deployment:

- If you want to install an agent on the CA Server Automation Manager system, you must perform a manual standalone agent installation. Deployment of agents on the CA Server Automation Manager system is not supported.
- The deployment process is dependent on the availability of existing host operating system services in order to gain remote access to target systems. When these services are not available on the target nodes it will be necessary to install the IDPrimer client package and a corresponding key on a target system. For more information, see the section *Manual Installation of the Remote Deployment Primer Software*.
- Deployment is supported to most, but not all, supported agent platforms. For more information about deployment support, see the *Release Notes*.

Deployment Credential Restrictions

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

Remote Deployment to UNIX/Linux Using Non Privileged User Account

If you want to use a nonprivileged user account, consider the following requirements about the sudo configuration:

- Sudo must not enforce that the executed program has a valid pseudo terminal attached to it. To disable such validation for a particular user (if it is globally enabled), add the line “Defaults:$username !requiretty” to the /etc/sudoers file. Replace $username by the actual username used for Remote Deployment.

  The standard way to edit the file is using the visudo command. The visudo command invokes $EDITOR. When editing is finished, it verifies the syntax of the file. If the result is not valid, visudo blocks saving the file.
- Sudo must not ask the user for a password before running the elevated program. To achieve this behavior, the NOPASSWD: keyword must be present on the line giving privileges to the user.
- Sudo must be allowed to run the necessary commands or all. Configuration entries (lines in /etc/sudoers) satisfying the previous requirements are, for example:

  ```
  $username ALL=(ALL) NOPASSWD: ALL
  ```

  or

  ```
  $username ALL = NOPASSWD: /usr/bin/id,/bin/sh /tmp/idprimer/PifInst -r
  /tmp/idprimer/install.rsp
  ```

  Note: Replace $username by the actual username used for Remote Deployment.

On Solaris, consider the following requirements for pfexec:

- Any local user can be given profile “Primary Administrator” with the following command

  ```
  usermod -P “Primary Administrator” {user}
  ```

- Any nonlocal user can be given profile “Primary Administrator” by manually adding an entry in file /etc/user_attr:

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

Deployment Sizing Key Factors

A number of key factors having a considerable impact on the infrastructure sizing, and system performance, including the following:

- Size of software packages to deliver.
- Number of software packages to deliver.
- Frequency of software package delivery.
- Network latency between deployment components and target computers.
- Network bandwidth management.

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

Under ideal conditions, a single distribution server can install about 50-100 SystemEDGE packages per hour. However, if many tasks fail due to any of the deployment prerequisites not being met, the throughput of a distribution server will be reduced.

Recommendations:

- Verify that target servers typically meet the requirements for deploying software remotely.
- A single distribution server can handle 500 target servers comfortably.
- Install additional distribution servers local to the target location.
- Deploy using distribution servers that are local to targets if possible.
- Schedule deployments to start during periods of low network traffic if possible.

Deploying/Installing SystemEDGE Agents Using Custom Ports

Deployment of SystemEDGE agents to a non-standard port requires a number of settings to be configured. To ensure the manager can discover and manage the system once it has been deployed, perform the following actions:

1. Update the package wrapper:
   - If you are using the remote deployment solution, you must first configure the package wrapper to specify the port to be used. Navigate to Provisioning, Deployment in the user interface and change the Port field. The Write Community string can also be updated here.
2. Update SNMP Community strings in CA Server Automation:
   - For the Manager to successfully monitor and manage a machine using a non-standard port, it must know the appropriate Port / Write Community string combination to use for monitoring and management. This can be done either by creating a global SNMP entry that can be used to monitor and manage multiple systems, or by creating a server-specific SNMP entry:
     - To update global SNMP settings: Navigate to Administration, SNMP in the user interface and add a new entry with the appropriate SNMP community string / port combination.
     - To update server-specific SNMP settings: Navigate to Policy, Explore, *Machine_Name*, Metrics, SNMP Settings and add a new entry for the required port / write community string.

Once these settings have been updated, the agent can be deployed / installed in the usual way. The SystemEDGE Platform Management Module will then use the custom port / write community string combination to discover, monitor and manage the server.

Reconfigure Ports for SystemEDGE Agents

You can reconfigure the port for SystemEDGE agent by reinstalling the agent. After reinstalling the agent, the settings remain unchanged with a provision to edit the details of the port to be reconfigured.

Deployment Package Library

The package library contains a configurable set of installable software packages where you can control which products, versions and platforms are available for deployment. You can control the way these products are installed by creating standard package configurations that define the parameters required for an unattended installation of a configured software package.

Each package must have an associated package configuration file. The configuration file provides information describing both the package details and how the package installation can be configured. For more information, see the Deployment Package Configuration File (see page 437) section.

The package library is located in the following directory:

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

The directory tree layout is defined by the requirements of the IDManager component. The package library itself consists of a top-level packages directory which contains two sub-directories, Public and Private. The Public directory contains all the deployable software packages.

```
..
└──SM
    ├──CA_LiteAgent
    ├──CA_ProcProbe_Utility
    ├──CA_SystemEDGE_AdvancedEncryption
    ├──CA_SystemEDGE_Core
    │   └──5.6.0
    │       └──ENU
    │           ├──AIX_aix
    │           ├──HPUX_hp
    │           ├──HPUX_ia64
    │           ├──Linux_ia64
    │           ├──Linux_x86
    │           ├──Solaris_sparc
    │           ├──Solaris_x86
    │           ├──Windows_ia64
    │           ├──Windows_x64
    │           └──Windows_x86
    ├──CA_SystemEDGE_ESAD
    ├──CA_SystemEDGE_HyperV
    ├──CA_SystemEDGE_LPAR
    ├──CA_SystemEDGE_MSCS
    ├──CA_SystemEDGE_RM
    ├──CA_SystemEDGE_SRM
    ├──CA_SystemEDGE_UCS
    ├──CA_SystemEDGE_Zone
    └──CA_VMVCAIM
```

The top-level Public directory has five sub-directories:

Component Name

Must be the IDManager instance name, which for CA Server Automation is SM.

Software Package

Contains all versions, localizations and architectures of a single deployable package, for example CA_SystemEDGE_Core, CA_SystemEDGE_SRM

Version

The version of the software packages contained within, for example 5.0.1

Language

The installation package language, for example. ENU

Architecture

The architecture specific installation materials, for example Windows_ia64, Solaris_x86. The architecture directory name must be one of the platforms supported by IDManager.

When run within the distribution server machine, the IDManager component uses directories under the distribution server. This contains a temporary cache of encrypted packages for its internal use. These packages should be removed upon job completion.

The private IDPrimer installation materials are contained in a different directory. By default these are stored under the installation directory of the IDManager component itself, in the following directory:

<CA Shared Components>/IDMgrApi/packages/private/idprimer

This directory contains the IDPrimer installation materials for all platforms supported by the infrastructure deployment component.

Deployment Package Library Maintenance

You may want to perform maintenance of the deployment packages, for example:

- Add new deployment packages
- Copy existing packages to a new package to create a patched version of the same base package (master image patch).
- Delete packages which are no longer required.
- Copy CA Server Automation compatible packages from installation media, or another specified directory into the package library.
- Copy an entire registered package to a new name and generate a new UUID for that package.
- Perform consistency checking of the package library and the associated deployment configurations data.

Package Filter

If upgrades have been applied to the server, Remote Deployment can show an increasing number of package versions. The default behavior of this release is to show only the latest packages available. As a consequence, the data in the Packages - Details tab is also filtered to show only the latest versions of each package. Selecting a wrapper from this panel, expands the tree at the selected wrapper position.

If you want to see all packages, you can 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 the software package installation materials, each deployable package must be referenced by an additional package configuration file, pkginfo_PLATFORM.xml. The package configuration file describes what packages to install and how the installation is configured. The configuration files provide the following:

- A localizable description of the installation package
- Text for any End User License Agreement (EULA) that must be accepted in the UI
- A mechanism by which package dependencies may be encoded in a machine-readable format
- Documenting the publicly accessible installation parameter types
- Additional context to the parameter types so a level of validation may be performed within the UI
- A mapping between parameter names and the tokens used to represent those parameters in the packages installation program, in a platform independent manner
- Specifies how the installation materials should be executed on a target machine
- Mapping between the installer exit codes and those understood by the deployment system

Localized elements of the pkginfo.xml file may optionally be provided using either side by side locale-specific files or embedded within a single file. At load time any file whose name matches filename pkginfo_PLATFORM.xml may be loaded to obtain localized message data.

The deployment system requires the package configuration files to be located parallel to the platform-specific sub-directory in the packaging tree. A single package description file can describe the installation for multiple platforms. See the following directory:

%AllUsersProfile%\Application Data\CA\SM\domainserver\Deployment\Packages\SM\CA_SystemEDGE_Core\5.6.0\ENU

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

Deployment Packages

Deployment packages provide the materials necessary to deploy monitoring software to systems across your enterprise. The following deployment packages are available:

CCA Agent

Provides the CCA Agent.

CA LiteAgent

Provides the CA Performance Lite Agent.

SystemEDGE

Provides the core SystemEDGE agent.

SystemEDGE Advanced Encryption

Provides a FIPS 140 compliant encryption package for SystemEDGE.

SystemEDGE AIX LPAR

Provides the LPAR AIM.

SystemEDGE ESAD

Provides the Exchange Server and Active Directory (ESAD) AIM.

SystemEDGE Hyper-V

Provides the Hyper-V AIM.

SystemEDGE MSCS

Provides the Microsoft Cluster Support (MSCS) AIM.

SystemEDGE RM

Provides the Remote Monitoring AIM.

SystemEDGE Solaris Zone

Provides the Solaris Zone AIM.

SystemEDGE SRM

Provides the SRM AIM.

SystemEDGE UCS

Provides the Cisco UCS AIM.

SystemEDGE VC

Provides the VMware vCenter AIM.

The Solaris Zone, Remote Monitoring, vCenter, and SRM AIMs all depend on both SystemEDGE 5.6 and Advanced Encryption packages. Therefore, if you want to deploy any of these packages, SystemEDGE and Advanced Encryption must either already exist on the system, or you must include these in the deployment job.

Deployment packages are broken into platform-specific variants, package wrappers are available for all platforms that support deployment.

Default Package Wrappers

Default package wrappers are provided out of the box for the software packages that can be deployed using Remote Deployment. These package wrappers contain installer parameters with default values where possible for the chosen software package. Where a package needs a mandatory parameter which is not safe to default for security reasons (for example, SNMP write community string, user name, password), or cannot be defaulted because it is user environment specific (vCenter host name), the default package wrappers will not contain a pre-filled default. This means the default package wrappers containing mandatory parameters need to be edited first with the mandatory parameter values filled in and saved before they can be used in a deployment job.

The following package wrappers contain mandatory parameters that need to be filled in prior to deployment:

SystemEDGE RM

Specify a user name and a password to deploy this package.

SystemEDGE MSCS

Specify an MSCS hostname, MSCS user name, and a password to deploy this package.

SystemEDGE VC

Specify a vCenter host name, user name and password to deploy this package.

SystemEDGE LPAR AIM

Specify a host name, user name, and password to deploy this package.

SystemEDGE Solaris Zone

Specify a host name, user name, and password to deploy this package.

SystemEDGE UCS

Specify a host name, user name, password, protocol, and port to deploy this package.

SystemEDGE Hyper-V

Specify a host name, user name, and password to deploy this package.

The mandatory values for the default package wrapper will need editing only once before the first deployment. These need not be edited again unless there is a need to modify the installer parameter values for a package. If you proceed to deploy a package with a default package wrapper that does not have the mandatory parameters filled in, the deployment wizard will indicate that the package wrapper is not in a deployable state.

For more information, see the *Online Help.*

Agent Configuration Without Write Community

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

- The SystemEDGE agent can be discovered by the SystemEDGE PMM even if the agent is configured with only SNMP read community and no write community. However, point configuration changes cannot be made to the agent without the agent configured with SNMP write community.
- Full vCenter and Remote Monitoring functionality is only supported if the agent is configured with write community. AIM configuration and administration from the CA Server Automation 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 Server Automation UI using Policy Configuration. Policy Configuration also allows you to configure the agent to use SNMP v3, which is more secure than SNMP v1/2.

Create a New Package Wrapper

Package wrappers provide platform-specific instructions for the deployment mechanism to follow when deploying a specific package. Each package contains a default package wrapper for all platforms that support remote deployment. You can create new package wrappers if certain systems require different settings than the default.

Follow these steps:

1. Select Resources, Deploy.

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

2. Expand Packages.

   The list of available packages appears in the Deployment pane.

3. Right-click a package name in the Deployment pane and select Create New Wrapper. You can also click + (New) on the Available Wrappers toolbar.

   The New Wrapper dialog appears.

4. Enter a name and an optional description for the wrapper, specify the platform the wrapper should support, and click OK.

   The wrapper is created, and details appear in the right pane.

   Note: If you create a SystemEDGE package wrapper, consider the dependency between the Trap Port, Trap Destination, and Trap Community fields. The behavior is that either none of these fields or all must be set. In case of a partial setting, the installer displays an error message.

Modify a Package Wrapper

Package wrappers define a set of platform-specific installation settings for a deployment package, such as installation path, port, trap communities, and so on. You can edit a user created or default package wrapper to change this set of installation settings. The available properties vary by the package type.

To modify a package wrapper

1. Select Resources, Deployment.

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

2. Expand Packages, the specific package type, and the wrapper platform, and select the wrapper to modify.

   The wrapper details appear in the right pane.

3. Modify the package properties as necessary and click Save. The options that appear in the Properties pane depend on the package type that you select.

Package types:

CA AIX LPAR AIM

Hostname

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

Username

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

Password

Specifies the password to use for connecting to IBM LPAR server. You must specify a IBM LPAR password in order to deploy this package.

Serial Number

Specifies the serial number to use for connecting to IBM LPAR server. You must specify a serial number in order to deploy this package.

CA LiteAgent

Shared Path

Defines the root installation directory to use for CA Shared Components if not already set.

Install Path

Defines the root installation directory for the package.

Suppress Reboot check box

Specifies whether to suppress any automatic reboot that might otherwise occur at end of the installation.

CA Solaris Zone AIM

Zones Host

Specifies the hostname to use for connecting to Solaris Zone server. You must specify a Solaris Zone hostname to deploy this package.

Username

Specifies the username to use for connecting to Solaris Zone server. You must specify a Solaris Zone username to deploy this package.

Password

Specifies the password to use for connecting to Solaris Zone server. You must specify a Solaris Zone password to deploy this package.

SystemEDGE

Install Path

Defines the root installation directory for the package.

Data Path

Defines the data directory for the package.

Shared Path

Defines the root installation directory to use for CA Shared Components if not already set.

Port

Defines the SystemEDGE port number.

Default: 161

Description

Defines the SNMP system description.

Location

Defines the SNMP system location.

Contact

Defines the SNMP System contact.

Read Community

Defines the SNMP read-only community string.

Default: public

Read-Write Community

Defines the SNMP read-write community string

Trap Community

Defines the SNMP trap community string.

Trap Destination

Defines the SNMP trap destination host name.

Trap Port

Defines the SNMP trap port.

Default: 162

Privilege Separation User (UNIX/Linux)

Specifies the user name under which credentials the agent run during SNMP communication.

This entry instructs the agent to run SNMP communication under another user account. The agent also uses this user's default group as an effective group.

Default: The agent operates using root account.

Start Agent check box

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

Suppress Reboot check box

Specifies whether to suppress any automatic reboot that might otherwise occur at end of the installation.

Disable Native Agent check box

Specifies whether to replace the native SNMP agent.

Use native settings check box

Specifies whether to use native SNMP agent settings (if replacing a native SNMP agent).

Run in Managed Mode check box

Specifies whether to run SystemEDGE in managed mode.

Managed Policy Name drop-down list

Specifies a list of available SystemEDGE policies.

Note: When you upgrade SystemEDGE from Version 4.3 or Version 4.2 patch level 3, the installer uses the following parameters only:

CASE_PUBDATADIR
CASE_MANAGER_HOSTNAME
CASE_MANAGER_POLICY_NAME
CASE_START_AFTER_INSTALL
CASE_LEGACY_MODE
CASE_SNMP_PORT
CASE_INSTALL_DOCS
CASE_SNMP_TRAP_COMMUNITY (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 above parameters, see the *Installation and Deployment* chapter in the *SystemEDGE User Guide*.

SystemEDGE Advanced Encryption

Suppress Reboot check box

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

SystemEDGE Hyper-V

Suppress Reboot check box

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

SystemEDGE RM

Default WMI username

Defines the default username to use for connecting to remote machines. You must specify a username in order to deploy this package.

Default WMI password

Defines the default password to use for connecting to remote machines. You must specify a password in order to deploy this package.

Suppress Reboot check box

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

SystemEDGE SRM

Allow scripts check box

Specifies whether to allow running scripts as tests.

Allow File I/O tests check box

Specifies whether to allow running file I/O as tests.

Disable user TOS check box

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

Allow untrusted SSL check box

Specifies whether to allow accessing SSL site with unverified certificates.

Suppress Reboot check box

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

SystemEDGE UCS

UCS hostname

Specifies the hostname to use for connecting to UCS. You must specify a UCS host name to deploy this package.

UCS username

Specifies the username to use for connecting to UCS. You must specify a UCS user name to deploy this package.

UCS password

Specifies the password to use for connecting to UCS. You must specify a UCS password to deploy this package.

UCS Protocol

Specifies what protocol to use, HTTP or HTTPS.

Port

Defines the UCS port number.

Default: 80 for HTTP or 443 for HTTPS.

SystemEDGE VC

VirtualCenter hostname

Specifies the hostname to use for connecting to VirtualCenter. You must specify a VirtualCenter host name in order to deploy this package.

VirtualCenter username

Specifies the username to use for connecting to VirtualCenter. You must specify a VirtualCenter user name in order to deploy this package.

VirtualCenter password

Specifies the password to use for connecting to VirtualCenter. You must specify a VirtualCenter password in order to deploy this package.

VirtualCenter Port

Defines the VirtualCenter port number.

Default: 443

VirtualCenter Protocol

Specifies what protocol to use, HTTP or HTTPS.

Suppress Reboot check box

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

Deployment Jobs

To deploy agents to target systems, you must create a deployment job. Deployment jobs contain the details that are required for CA Server Automation to deliver the deployment packages to the appropriate systems at the appropriate time. When you create a deployment job, you specify the following information:

Job information

Includes the job name and whether to base the job on an existing template.

Deployment package

Includes the platform, the packages to deploy, and the specific wrappers for each package.

Machine information

Includes the systems to which to deploy the packages and system credentials required to establish a connection.

Deployment time

Specifies when to perform the deployment: immediately, staggered over a specific time period, or scheduled for a specific time in the future.

You can also save the job as a template after you create it. A template saves the package and machine selections so that you can easily reuse them for subsequent jobs.

For more information about creating a deployment job and running a deployment, see the *Online Help*.

Track Deployment Job Status

Once a job to deploy a set of agent packages to a set of computers has been started, you can track its progress and status. The Jobs folder displays a table of all created deployment jobs that lists the job name, packages included, job status, and so on. From this table, you can drill down to view more details about a specific job, including why a job failed.

To track a deployment job status

1. Select Resources, Deploy pane.

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

2. Click the Jobs folder.xxxxxx

   The Job Status pane appears.

3. Click the job you want to view.

   The Job Information page appears.

4. In the Task Status pane, filter out particular job tasks by using any of the filters available. Alternatively, use the paging interface to identify the desired task.
5. Click Status Extended to view extended information about a job.

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

   Information

   Displays general information about the job

   Message

   Displays a message about the job, for example, Package delivery failed

   Reason

   Displays the reason for the failure

   Action

   Displays what action to take to correct the problem

Common causes for deployment failure include the following:

- Lack of machine availability
- Invalid system credentials
- Inability to resolve the system host name
- Unfulfilled package dependency (for example, SystemEDGE is required for a successful SRM AIM deployment).

View Deployment History

Deployment history information is available from the following places:

Deployment Pane

Displays a count of completed, active, pending, and failed deployment tasks and a summary of successful deployment. Click the top-level Deployment folder to access this view.

Jobs pane

Displays a table of all created deployment jobs that list the job name, packages included, and job status. From this table, you can drill down to view more details about a specific job, including why a job failed. Common causes for deployment failure include the following:

- Invalid system credentials
- Inability to resolve the system host name
- Unfulfilled package dependency (for example, SystemEDGE existence is required for a successful SRM AIM deployment)

You can resubmit a job from this pane to correct the reason for failure and redeploy.

Click the Jobs folder to access this view.

Deployment Dashboard Views

The following views are available on the Dashboard for tracking deployment metrics:

Deployment Task Summary

Displays a pie chart and a list showing the number of completed, active, pending, and failed deployment tasks.

Unresolved Deployment Tasks

Displays a list of deployments that did not complete successfully. You can click the job ID to view details about why the task is unresolved.

Active Deployment Tasks

Displays a list of deployment tasks that are currently active. Details include the associated deployment job, target, package, and current state. You can click the task ID for details about the current status.

Deployment Package Summary

Displays a bar chart showing the number of deployment for each deployment package type.

Completed Deployment Jobs

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

Infrastructure Deployment Process

When executing a deployment, the primary steps of the process are as follows:

1. From the administrator computer, the infrastructure deployment client component issues a request to the IDManager to install an agent on a list of one or more target computers. The deployment manager may be running on a computer that is remote from the client. The list of targets can consist of explicit machine names or IPv4 addresses.

   Note: Only discovered resources can be deployed to.

   For deployment to succeed to each target computer, it is important to verify that its name, whether entered explicitly or obtained from a container, is suitable for resolving to the address of the target as seen on the deployment manager computer. If, for example, the list of targets retrieved from a directory is not fully qualified with network domain names, deployment may not be able to proceed in certain network configurations.

2. A check is made to see if the IDPrimer is already installed on the target computer. If not, IDPrimer will be installed first on the target computer. The IDManager tries to deliver the IDPrimer installation package. The delivery method used depends on the target operating environment and the security that has been enabled on it. After the IDPrimer image is copied across to the target computer, its installation is initiated.

   As some operating systems do not have a method for remote invocation of the IDPrimer installation, in which circumstances the IDPrimer installation may have to be performed manually.

3. The IDPrimer installer installs itself and the CA Messaging (CAM) component on the target computer. Once the IDPrimer is installed and IDManager has received the 'installation complete' signal from the target computer, package deployment can be initiated. An IDManager that has previously installed an IDPrimer and has authenticated with it can deploy packages without needing to resupply user names or passwords. On subsequent deployments, IDPrimer uses asymmetric cryptographic keys to authenticate and limit access to those managers from which we have already gained access.

For details on using the deployment manager through the CA Server Automation UI, see the *Online Help*.

More Information

Prerequisites for Automatically Deploying CA Server Automation Infrastructure (see page 452)
Notes on Infrastructure Deployment Using IPv6 Addresses (see page 454)
Protocols for Transferring Packages Employed by IDManager (see page 455)
Manual Installation of the Infrastructure Deployment Primer Software (see page 455)
Deployment Primer Installation on Windows (see page 455)
Deployment Primer Installation on Linux or UNIX (see page 456)
Provide the Deployment Management Certificate to a Primer Installation (see page 456)
Deployment Management Certificate on Windows (see page 456)
Deployment Management Certificate on Linux or UNIX (see page 457)
Compatibility Libraries for Linux (see page 457)

Prerequisites for Automatically Deploying CA Server Automation Infrastructure

The Infrastructure Deployment component lets you remotely install agent software to target computers. The installation can only be done using the facilities offered by the underlying operating systems on source and target computers, and is subject to any restrictions imposed by an enterprise network configuration.

The initial step when deploying software is to install a small primer application remotely, the IDPrimer, onto the target computer. This IDPrimer software is responsible for subsequent transfer of software component installation images, and the invocation of their installation. When delivering the IDPrimer to the target computers, the deployment manager must supply user credentials that are valid on the target.

The IDPrimer is transferred to the target system using one of the following mechanisms. If the target computer's operating system is known to the deployment manager, an appropriate transfer mechanism is selected. If the target operating system cannot be determined, each of the following mechanisms is attempted in turn.

- Opening a network share

  The deployment manager tries to connect to a Windows network share on the target system. By default, the share name used is ADMIN$, but this is controlled by an IDManager configuration option. This mechanism is available only from deployment managers running on a Windows-based environment and will only succeed on some Windows targets. Windows variants such as Windows XP Home do not support this deployment mechanism.

- Opening a network connection to the target computer using the SSH protocol, and transferring the primer installation package using SFTP

  This mechanism works on any computer running an SSH server, however, it is useful when targeting Linux or UNIX computers.

  Note: When deploying to Solaris systems, we recommend that you use either SunSSH v1.1 (or higher) or the latest version of OpenSSH. Refer to the following website for additional details about patches applicable for Solaris platforms and versions: http://opensolaris.org/os/community/security/projects/SSH.

  If you are running a firewall on the target computer, verify that the SSH port (22) is enabled to permit connection from the deployment manager. You should also check that the SSH server on the target computer is configured to use an RSA key with the 3DES cipher for encryption and the HMAC-SHA1 message authentication code (MAC). Most SSH servers support this configuration by default, but if they do not, consult your SSH server documentation for further instructions.

  To successfully deploy to a UNIX or Linux agent, configure the /etc/ssh/sshd_config configuration file of your recent SSH implementation as follows:

  - Set PasswordAuthentication to Yes
  - Set PermitRootLogin to Yes or configure sudo/pfexec as described in section Remote Deployment to UNIX/Linux Using Non Privileged User Account (see page 431)
  - Verify that SFTP subsystem is enabled

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

  When deploying to some IBM AIX systems that are running both an IPv4 and IPv6 stack, using an IPv6 address, the target computer SSH server may be listening only on port 22 for IPv4. This would cause the deployment to fail. To correct this, edit the sshd_config configuration file and set the ListenAddress to "::".

  Note: If you want the SSH communication between the deployment manager and the target computer to be FIPS-compliant, you must verify that the SSH server running on the target is also using FIPS-compliant cryptographic module, apart from setting FIPS-only mode on the deployment manager.

Important! Some modern operating systems do not encourage, and sometimes actively prohibit, the remote installation of software. If you try to deploy software to these systems, you will usually see the deployment fail with a status of No Primer Transport. In such cases, installation of software components may be performed in other ways, for example, installation off physical distribution media such as DVD.

Alternatively, you can pre-install or provision machines with the IDPrimer software. This will allow deployment without having to rely on facilities offered by the underlying operating systems. In cases where no authentication has been carried out, valid credentials would need to be supplied before deployments being authorized.

To determine whether automatic deployment is possible in your environment, you can perform some simple checks by running the following standard operating system operations:

- For delivery of the IDPrimer image using Windows shares, you must be able to map a share (default: ADMIN$) from your deployment manager host computer to each deployment target computer using the target user credentials supplied in the deployment request.
- For delivery of the IDPrimer image using SSH, you must be able to connect using SSH from the deployment manager to the deployment target computers.

More Information

Remote Deployment to UNIX/Linux Using Non Privileged User Account (see page 431)

Notes on Infrastructure Deployment Using IPv6 Addresses

If you are going to use CA Server Automation 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 Server Automation Infrastructure (see page 452).

Manual Installation of the Infrastructure Deployment Primer Software

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

As well as installing the primer software itself, you must install a security key that is generated by the deployment manager you want to use to deploy to your target computers.

Deployment Primer Installation on Windows

The installation of the deployment primer on a target computer running Windows requires the following actions:

- Make the CA Server Automation installation media (DVD) available on the target computer, or manually copy the primer setup file to the target computer. The primer setup file is stored on the installation media in the following directory:

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

- Run IDPrimer_Setup.exe on the target computer to install the primer.

Deployment Primer Installation on Linux or UNIX

The installation of the deployment primer on a Linux or UNIX target computer requires the following actions:

- Make the CA Server Automation installation media (DVD) available on the target computer, or manually copy the primer installation image to the target computer. The primer installation image is stored on the installation media in the following directory:

  %PROGRAMFILES%\SC\IDMgrApi\packages\private\idprimer\Linux_x86

- Change to the directory containing the primer installation image on the target computer and run the following installation command to install the primer:

  ```
  # sh installidp
  ```

Provide the Deployment Management Certificate to a Primer Installation

The deployment manager generates a certificate that needs to be transferred to the target computer before the primer on the target computer will accept deployment packages. The deployment certificate file is named dmkeydat.cer

The location of the certificate is configurable at installation time. You may configure a different file location if you want to store the certificate in a more secure area or in a location shared between two managers providing a failover solution. In the latter case, sharing the certificate enables deployment managers to communicate with IDPrimer components delivered from either manager without the need to resupply authentication credentials.

Deployment Management Certificate on Windows

On Windows, the deployment certificate is located in the following directory:

C:\Program Files\CA\SC\IDMgrApi\config\SM

The certificate file (with the suffix.PMR for example, MANAGER1 SM.PMR) must be copied to the primer installation folder on the target computer, which by default is the following:

\Program Files\CA\SC\IDPrimer

Deployment Management Certificate on Linux or UNIX

On Linux and UNIX, the deployment certificate must be copied to the primer installation folder on the target computer, which by default is the following:

/opt/CA/SharedComponents/ID/primer/bin

Compatibility Libraries for Linux

The IDPrimer installer assumes that certain dependent libraries are present. If these libraries are not present, installed components may not work properly.

The following table details the software library prerequisites. These libraries must be present on Linux hosts prior to installing IDPrimer.

Linux Distribution/Version	Required RPM Packages
Red Hat 5.0 Enterprise Linux	glibc 2.3.3-84 compat-libstdc++ 33-3.2.3-61 For 64-bit versions of the OS: ncurses 5.4-1.3 (i386) or ncurses-devel 5.4-13 (i386) zlib 1.1.4-8.1 (i386) or zlib-devel 1.2.1.2-1.2 (i386) libstdc++.so.6
Red Hat 4.0 Enterprise Linux (AS/ES/WS)	glibc 2.3.3-84 compat-libstdc++ 296-2.96-132.7.2 compat-libstdc++ 33-3.2.3-47.3 For 64-bit versions of the OS: ncurses 5.4-1.3 (i386) or ncurses-devel 5.4-13 (i386) zlib 1.1.4-8.1 (i386) or zlib-devel 1.2.1.2-1.2 (i386) libstdc++.so.6
SuSE Linux Enterprise Server 10	glibc 2.3.3-84 compat 2004.4.2-3 libstdc++ 3.2.2-38 libstdc++.so.6
SuSE Linux Enterprise Server 9 (SP3)	glibc 2.3.3-84 compat 2003.5.12-56 libstdc++ 3.2.2-38 libstdc++.so.6

SuSE Linux Enterprise Server 9 (x64)	glibc 2.3.3-84 compat 2003.5.12-56 libstdc++ 3.2.2-38 libstdc++.so.6
SuSE Linux 10.1	glibc 2.3.3-84 compat 2004.4.2-3 libstdc++ 3.2.2-38 libstdc++.so.6
SuSE Linux Professional 10	glibc 2.3.3-84 compat 2004.4.2-3 libstdc++ 3.2.2-38 libstdc++.so.6
SuSE Linux Professional 9.3	glibc 2.3.3-84 compat 2003.5.12-56 libstdc++ 3.2.2-38 libstdc++.so.6
SuSE Linux Professional 9.2	glibc 2.3.3-84 compat 2003.5.12-56 libstdc++ 3.2.2-38 libstdc++.so.6
SuSE Linux Professional 9.1	glibc 2.3.3-84 compat 2003.5.12-56 libstdc++ 3.2.2-38 libstdc++.so.6
SuSE Linux Desktop 9	glibc 2.3.3-84 compat 2003.5.12-56 libstdc++ 3.2.2-38 libstdc++.so.6

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

Deployment to Windows Vista, Windows 2008 and Windows XP Computers Running Firewall Software

To enable deployment of agents to computers running firewall software, consider the following:

- If the firewall of a target computer running Windows Vista or Windows 2008 operating system is *off* (disabled) and deployment to the computer fails, create or set the following registry variable so that it is a DWORD type with a value 0x1:

  `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\LocalAccountTokenFilterPolicy`

  This is required because User Account Control (UAC) in Windows Vista or Windows 2008 does not automatically grant administrative rights to local users. This occurs even though the local users are members of the Administrators group.

  Note: Setting this value will result in remote UAC access token filtering being disabled.

  Setting this value is only worthwhile if the user has a local administrator account on the computer running Windows Vista or Windows 2008. Domain administrators will not benefit from this change.

- If the firewall of a target computer running Windows Vista or Windows 2008 is *on* (enabled), the following ports should be opened in addition to file sharing ports, to enable deployment to that computer:

  UDP ports

  CAM: 4104

  File and printer sharing, and so on: 137, 138

  TCP ports

  IDManager: 135

  File and printer sharing, and so on: 139, 445

- If deployment still fails, the following Outbound Rules in the firewall for Windows Vista or Windows 2008 should be fully enabled:
  - Remote Assistance
  - Network Discovery
  - File and Printer Sharing
  - Core Networking

- To enable agent deployment to Windows XP computers that run firewall software you must perform the following actions manually:

  1. Change Security Policy Network Access: Sharing and security model for local accounts from 'Guest only - local users authenticate as Guest' to 'Classic - local users authenticate as themselves'.

     The Classic model allows fine control over access to resources and prevents network logons that use local accounts from being mapped to the Guest account, which typically has Read Only access to a given resource.

  2. Configure the following firewall settings:

     - Allow File and Printer Sharing
     - Open UDP Port 4104
     - Open TCP Port 135

Storage Provisioning for NetApp

The CA Server Automation Storage Provisioning Manager lets you provision new or additional NetApp storage for virtual and physical systems through integration with NetApp Provisioning Manager. The Provision Storage and Provision Datastore wizards are available at appropriate stages of the system provisioning workflow.

The following table identifies the Storage Provisioning Manager components, physical and virtual servers they apply to, associated access points, and available tasks.

Component	Servers	Access Points	Storage Tasks
Provision Datastore Wizard	VMware vCenter ESX servers	■ Provisioning Quick Start menu ■ During VMware vCenter server provisioning where Datastores are selected ■ VMware vCenter Add New Disk dialog	Provision

Component	Servers	Access Points	Storage Tasks
Provision Storage Wizard	Physical and virtual servers: ■ Windows ■ Red Hat Linux ■ SUSE Linux ■ Solaris (including Solaris Zones on a virtual server) ■ AIX (including LPAR on a virtual server) ■ HP-UX	Right-click Managed server	Provision
	Cisco UCS	During Cisco UCS provisioning, when you define vHBA as the boot policy and select remote storage	Provision
	AIX LPAR Virtual I/O Server	During AIX LPAR provisioning, when you add a new virtual iSCSI adapter followed by a new SCSI adapter	Provision
CA Process Automation Connectors	Physical and virtual servers: ■ Windows ■ Red Hat Linux ■ SUSE Linux ■ Solaris ■ AIX ■ HP-UX ■ VMware vCenter ESX	■ CA Process Automation Client > Custom Operators\CA Server Automation\Storage folder to create processes and start forms ■ CA Server Automation UI > Policy tab > Actions & Rules > Actions tab to add an Action **Note:** For more information about CA Process Automation Connectors, see the chapter "Configuring Resources" in this guide.	■ Provision ■ Discover ■ Resize ■ Move ■ Deprovision
cadpmspm cli commands	Physical and virtual servers: ■ Windows ■ Red Hat Linux ■ SUSE Linux ■ Solaris ■ AIX ■ AIX LPAR Virtual I/O Server ■ HP-UX ■ VMware vCenter ESX	CA Server Automation cli commands	■ Provision ■ Discover ■ Resize ■ Move ■ Deprovision

The Storage Provisioning Manager supports the following protocol types:

- SAN-based iSCSI (VMware vCenter ESX servers, AIX LPAR Virtual I/O Servers)
- SAN-based FCP (VMware vCenter ESX servers, Cisco UCS)
- NAS-based CIFS (physical and virtual servers)
- NAS-based NFS (physical, virtual, and VMware vCenter ESX servers)

You can provision storage using the following methods:

- Enhanced storage policies

  Enhanced storage policies are policies you define for use with the CA Server Automation Storage Provisioning Manager. Using enhanced storage policies lets you create a set of predefined parameters you can quickly apply during storage provisioning.

- Storage services

  Storage services are a combination of storage resource pools and NetApp provisioning policy definitions (including thick and thin provisioning selections). Storage Architects create storage services using NetApp Provisioning Manager.

- Provisioning policy and resource pools

  A provisioning policy is a NetApp Provisioning Manager policy definition defined by the Storage Architect. When you use a provisioning policy as the provisioning method, specify the resource pools to use with the policy.

Use Case: Storage Provisioning

The following scenario represents a common business issue and the solutions available to administrators using the CA Server Automation Storage Provisioning Manager.

Business Issue

An administrator is managing a VMware vCenter ESX environment using CA Server Automation. The administrator realizes additional storage is necessary based on the free space indicated for the datastores on multiple servers being monitored.

Storage Provisioning Manager Solutions

To save time, the administrator uses CA Server Automation to create an enhanced storage policy and associates it with the monitored servers. Enhanced storage policies also support provisioning across multiple, diverse platforms. The administrator chooses from the following options:

- Provisioning method (NetApp Storage Service or NetApp Provisioning Policy)
- Protocol (NFS, CIFS, iSCSI, FCP)
- Storage server (NetApp DataFabric Manager)
- Policy attributes (provisioning-specific settings; including protocol export settings)
- Servers for association with the policy

The administrator assigns a meaningful name to the enhanced storage policy; facilitating reuse of the predefined set of options during the storage provisioning process.

During storage provisioning, the administrator selects from available enhanced storage policies, storage services, or provisioning policies. After the administrator makes the required choices, CA Server Automation orchestrates the storage provisioning with NetApp Provisioning Manager and indicates that the operation is complete.

The administrator also has the option of using CA Process Automation connectors to automate the storage provisioning process. CA Process Automation provides a graphical user interface that lets administrators configure and manage processes that CA Server Automation can activate. The following examples provide specific scenarios for using the Storage Provisioning Manager with an enhanced storage policy or CA Process Automation connector.

Example: Provision a new NFS datastore to VMware vCenter ESX using an enhanced storage policy

The following steps show how the administrator can provision a new NFS datastore during VMware vCenter ESX provisioning using an enhanced storage policy.

1. Provisions a VMware vCenter ESX server using Resource, Provisioning, Provision Storage.
2. Verifies that the storage available for the datastore in the Free field does not meet the requirements.
3. Uses the Provision Datastore wizard to verify and change storage requirements:
   - Views the Free field and sees that storage does not meet requirements
   - Selects the NFS enhanced storage policy that meets requirements on the Select Enhanced Storage Policy pane
4. Verifies that the VC Datastore field is updated with the new datastore and continues VMware vCenter ESX server provisioning.

Chapter 8: Setting Up Reservation Manager

Reservation Manager provides the capability to reserve physical and virtual machines, create reservation templates, view inventory, and manage reservations.

This chapter describes post-installation tasks required for setting up the Reservation Manager for end users.

Note: The Reservation Manager online help describes best practices and how to use the Reservation Manager. The installation process is detailed in the *Installation Guide*.

This section contains the following topics:

Prerequisites

Before you start setting up the Reservation Manager for end users, perform the following tasks:

- Verify that the Reservation Manager is properly registered. To do so, check the Reservation Manager status on the CA Server Automation Administration, Configuration page.
- Prepare your environment for the Reservation Manager.
- Prepare CA Server Automation for the Reservation Manager.
- Verify that all required components and servers are installed correctly.

More information:

Prepare Your Environment for Reservation Manager (see page 466)
Prepare CA Server Automation for Reservation Manager (see page 466)
Check Required Components (see page 467)

Prepare Your Environment for Reservation Manager

Before you set up Reservation Manager, prepare your environment by collecting the following information:

1. Identify the systems to include in the Reservation Manager inventory for fulfilling reservation requests. Required system information includes the following:
   - MAC address
   - Default password for Windows, AIX, Solaris, or Linux
   - Network Interface Controller model
2. Verify that the required Network Interface Controller device driver is available in one or more of the OS images that will be made available to users.
3. Identify the operating systems and versions to deploy.
4. Confirm CA ITCM support for each Windows and Linux operating environment.
5. Confirm Solaris JumpStart server support for Solaris operating systems.
6. Confirm NIM server support for AIX operating systems.
7. Confirm your HMC/IVM server support for IBM PowerVM logical partitions.
8. Identify all users to grant access to Reservation Manager. Confirm the following:
   - Reservation Manager administrator users
   - Reservation Manager end users
9. If you are using native CA EEM security, verify that all users are defined in the CA EEM database.

Prepare CA Server Automation for Reservation Manager

Configure CA Server Automation to prepare the material that Reservation Manager uses to create its inventory and fulfill reservation requests. To prepare CA Server Automation for a Reservation Manager setup, do the following:

1. Select the software packages to make available for deployment.
2. Discover all systems that you used to test CA IT Client Manager (CA ITCM) Operating System Installation Management (OSIM) imaging. Select the Resources tab, right-click a system in the Explore pane, and select Provisioning to provision an image to systems for the Reservation Manager inventory.
3. Deploy the CCA Agent to all of the systems discovered in Step 2. Verify that the agent collected system detail information such as MAC Address, number of CPUs, available memory, and disk space.

4. Create as many services as necessary to provide the appropriate level of control. Add the systems that you want to make available to Reservation Manager to these services.

   The Reservation Manager imports these services to create resource pools of systems that users can reserve. The Reservation Manager controls access to systems at the service and resource pool level and other usage policies.

   Configuring services requires you to install the Automation Management Framework that is provided with CA Server Automation. The Automation Management Framework optionally runs CA Server Automation actions during reservation setup or tear-down processing.

Check Required Components

Verify that the following components are installed correctly.

Imaging

Depending on the operating environment, the Imaging component requires the following servers:

Amazon Elastic Compute Cloud

Provides imaging functionality for Windows and Linux provisioning to the Amazon Elastic Compute Cloud (EC2).

CA IT Client Manager Server - OSIM Imaging

Provisions computers running Windows or Linux operating systems using the CA ITCM OS installation technology (OSIM), which also contains Software Delivery.

Solaris JumpStart Server

Provides imaging functionality for computers running Solaris operating systems. This component is only required if you provision Solaris operating systems.

Hyper-V SCVMM

The System Center Virtual Machine Manager (SCVMM) provides functionality to deploy systems using templates and customization profiles for hardware and operating system options. Most sites have either Microsoft Hyper-V or VMware vCenter, but not both.

Note: Reservation Manager supports Hyper-V provisioning of Windows Server 2003 and Windows Server 2008 operating systems. Windows 7 and Windows Vista are not supported.

VMware vCenter Server

Provides functionality to deploy virtual machines using custom templates and specifications.

NIM Master Server

Provides imaging functionality for computers running IBM AIX operating systems. This component is required only if you provision IBM AIX operating systems. Multiple NIM Master Servers are supported. The NIM Master Server must be configured and system resources and resource groups must be created and configured on the NIM Master Server. The NIM Master Server must be registered with CA Server Automation.

AppLogic Grid Controller

Provides functionality to deploy and manage AppLogic applications and application templates.

IBM PowerVM (HMC/IVM)

Provides management and virtualization capabilities for IBM PowerVM logical partitions on AIX operating systems.

Software Delivery Adapter

The Packaging component that installs requested software to the servers allocated to fulfill a reservation request. Software Delivery is required to add software to existing servers.

Post-Installation Configuration

After the installation is complete, you can perform several operations to configure and customize your environment. The procedures in this section describe how to tailor the configuration of the Reservation Manager to your site.

Most settings are located on the Configuration Settings page of the user interface. This page organizes the settings into groups with names like Approvals and Notifications. Within each group are links to dialogs that contain detailed descriptions of the settings and let you change the values. You can expand or collapse the groups.

Use Help Desk for Reservation Approvals

Reservation Manager provides an option to use a help desk system for managing the reservation approval process.

To use a help desk system for reservation approvals

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Approvals area. Each one opens a dialog in which you can change the value and click OK.

   Open Help Desk Ticket

   Controls whether help desk tickets are opened when reservation requests that require approval are submitted. When set to true, tickets are opened.

   Default: false

   Help Desk Ticket Type

   Specifies the type of ticket that is opened.

   Limits: Request|Incident|Problem

   Default: Request

   Help Desk Ticket Template

   Specifies an optional ticket template that is used when opening tickets. Templates let you define default values for various ticket attributes such as priority.

   Default: None

   Automatically Close Help Desk Ticket Upon Approval

   Specifies whether help desk tickets are closed when a reservation is approved or rejected.

   Default: false (tickets are not closed)

   The configuration change takes effect when the next reservation is made.

Set Automatic Cancellation of Unapproved Reservations

If reservations are not explicitly approved in time, the Reservation Manager can cancel them automatically. If automatic cancellation is enabled, you can also specify how long after the reservation start time to wait before automatically canceling the reservation. By default, this option is disabled.

To set automatic cancellation of unapproved reservations

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Approvals area. Each one opens a dialog in which you can change the value and click OK.

   Automatically Cancel Unapproved Reservations

   Specifies whether reservations are canceled automatically if not approved by the time specified in the next setting, Automatically Cancel Unapproved Reservations Time Delay.

   Default: false (no automatic cancellation)

   Automatically Cancel Unapproved Reservations Time Delay

   Specifies the number of hours to wait after the scheduled reservation start time before automatically canceling unapproved reservations.

   Default: 2

   The configuration change takes effect when the next reservation is made.

Specify When to Send Pending Approval Request Notification

Reservation Manager can send an email alert to the administrator when the start time for a reservation is approaching and the reservation is not yet approved. You can configure how many hours before the start time to send this notification. By default, the notification is sent two hours before the reservation start time.

Note: This option applies only when Reservation Manager is configured for manual approval.

To specify when to send pending approval request notification

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Approvals area:

   Approval Required Notification Time

   Specifies how many hours before unapproved reservations start time to send the administrator an email. Enter -1 to send no emails. Enter 0 to send emails when reservations are submitted.

   Default: 2

5. Change the setting in the Value field and click OK.

   The configuration change takes effect when the next reservation is made.

Configure Chargeback

You can control whether chargeback is used, how often costs are calculated per day, and the currency to use.

To configure chargeback

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Chargeback area. Each one opens a dialog in which you can change the value and click OK.

Chargeback Calculation Currency

Indicates the currency to use for cost calculations, based on ISO 4217 standards. This setting determines the symbol shown for chargeback cost estimates and calculations.

The values have two extra decimal places. For example, U. S. Dollars have two decimal places, so the chargeback calculation would use four decimal places.

Values:

AUD (Australian dollar, $)

CAD (Canadian dollar, $)

CNY (Chinese yuan, ¥)

EUR (Euro, €)

GBP (Great Britain pound sterling, £)

INR (Indian rupee, ₹)

JPY (Japanese yen, ¥)

USD (US dollar, $)

Default: USD

Note: If you select the Indian rupee, check if you need a Windows update to support the new currency symbol. Go to support.microsoft.com, and search for the knowledge base article KB2496898.

Chargeback Calculation Frequency

Indicates how many times per day to calculate costs, starting at midnight. Modifications to this value take effect after the next scheduled calculation time or when the chargeback service is restarted. Changes to this setting take effect at the next calculation time or when the service is restarted. If you increase this value and restart the service, you do not have to wait overnight to run chargeback reports and see data.

Limits: 1, 2, 3, 4, 6, 8, 12, 24

Default: 1 (once a day at midnight)

Chargeback Retention

Indicates the number of days to retain cost calculation data. Chargeback records are purged at midnight, so changes are apparent the next day.

Default: 90 days

Chargeback Is Enabled

Indicates whether the chargeback feature is being used. If you change the setting to false, previous calculations are kept in the database for the number of days specified by the Chargeback Retention setting. If you change false to true, calculations start at the next interval specified by the Chargeback Frequency setting or immediately if the service is restarted.

Default: true

The configuration changes take effect when you log out and log back in.

Enter Home Page Welcome Text

Administrators can specify the welcome text that appears at the top of the home page, above the Tasks and Announcements.

To enter home page welcome text

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the General area:

   Home Welcome Text

   Specifies the text that appears at the top of the Reservation Manager home page in the welcome banner. The banner is displayed only if this setting is defined.

   Default: Welcome to the Reservation Manager. Use this site to automatically provision systems "On Demand" or reserve systems for a future time.

5. Change the setting in the Value field and click OK.

   The configuration change is made when you restart the browser.

Configure IBM PowerVM Logical Partitions

You can set default and maximum memory for logical partitions, disk size, starting slot number, and more.

To configure IBM PowerVM logical partitions

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the IBM PowerVM area. Each one opens a dialog in which you can change the value and click OK.

   Default Logical Partition Disk Size (GB)

   Specifies the default disk space in gigabytes for a new logical partition.

   Default: 10

   Default Logical Partition Memory Setting (MB)

   Specifies the default memory setting in megabytes for a new logical partition.

   Default: 1024

   Default Processor Units

   Specifies the default number of processor units for a new logical partition when using shared mode.

   Default: 0.1

   Default Processors

   Specifies the default number of processors for a new logical partition when using dedicated mode.

   Default: 1

   Default Uncapped Weight

   Specifies how any available unused capacity is distributed to contending logical partitions.

   Default: 128

Default Virtual Processors

Specifies the default number of virtual processors for a new logical partition.

Default: 1

Maximum Logical Partition Memory Setting (MB)

Specifies the maximum memory setting in megabytes for a new logical partition.

Default: 16384

Maximum Processor Units

Specifies the maximum number of processor units for a new logical partition when using shared mode.

Default: 0.1

Maximum Processors

Specifies the maximum number of processors for a new logical partition when using dedicated mode.

Default: 8

Maximum Virtual Adapters

Specifies the maximum number of virtual adapters for a new logical partition.

Default: 9

Maximum Virtual Processors

Specifies the maximum number of virtual processors for a new logical partition.

Default: 64

Starting Slot Number for Virtual Adapters

Specifies the starting slot number for virtual adapters.

Default: two serial adapters

Configure Announcements

Reservation Manager can display an optional Announcements pane on the end-user home page. Use announcements to communicate key operational information and news to users, such as planned and unplanned outages, new system or image support, and so on.

To configure announcements

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the General area:

   Announcements

   Specifies the location of the Announcements.html file, which contains announcement information. The default directory for the file is /Tomcat/webapps/ssrm/Announcements. You can specify a different directory here.

   Default: default

5. Open the directory where Announcements.html is located, and modify the text.

   Important! If you edit the Announcements file, save it with UTF-8 encoding. Non-English operating systems have multibyte characters that must be saved with UTF-8 encoding. Microsoft Windows Notepad can save with UTF-8 encoding.

6. Save and close the file.

   The announcement change is made when you restart the browser.

Configure Online Help

Reservation Manager displays a Help link at the top of the home page to open the online help. By default, the Help link points to the Reservation Manager help. Administrators can substitute a customized URL or remove the entry.

To configure online help

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the General area:

   End User Help

   Specifies the online location of Reservation Manager online help for users and administrators. You can enter a specific URL here. If you remove the entry, the Help link is not displayed.

   Default: default

5. Enter a value and click OK.

   The change takes effect when you restart the browser.

Configure the Contact Hyperlink

You can configure a hyperlink in the upper right corner of each Reservation Manager web page that lets an end-user request administrative or technical support. The link appears only if its label and URL are configured. By default, a contact hyperlink is not displayed.

To configure the contact hyperlink

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the General area. Each one opens a dialog in which you can change the value and click OK.

   Administrator Contact Name

   Specifies a label for the contact hyperlink. If no value is specified for this option, or for the Administrator Contact Url, Reservation Manager does not display the hyperlink.

   Default: Contact Us

Administrator Contact Url

Specifies an URL that opens a web page or an email message. If no URL is specified, or if Administrator Contact Name is blank, Reservation Manager does not display the hyperlink.

Default: None

Examples:

To open a web browser:

```
http://www.anycompany.com
```

To open an email message:

```
mailto:admin@anycompany.com?Subject=Reservation%20Manager
```

The email message contains the following values:

Email recipient: admin@anycompany.com

Subject: Reservation Manager

The configuration change is made after you restart the browser.

Specify a Timeout Value

Reservation Manager lets administrators specify a timeout value, which indicates how many minutes to wait for data requests to take effect.

To specify a timeout value

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the General area:

   Data Timeout

   Specifies how long, in minutes, to wait for data requests to return a value. Enter -1 to wait for an unlimited amount of time.

   Default: 4

5. Enter a value and click OK.

   The change takes effect when you restart the browser.

Configure Email Notifications

The Reservation Manager can send email notifications for key events to end users and administrators. An example of a situation that causes email notifications to administrators is when a datastore has insufficient space for a new virtual machine.

Perform post-installation configuration to activate email support. If CA EEM is configured to use an external directory, CA EEM automatically obtains the email address of the user. If you are using CA EEM without external directory support, the CA EEM administrator must specify the email address of the user.

To configure email notifications

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Notifications area. Each one opens a dialog in which you can change the value and click OK.

   Mail Notification Address

   Specifies the sender name that appears on emails. An example is <ssrm@somecompany.com>.

   Mail Server

   Specifies the name of the SMTP mail server that sends the emails. An example is mail.somecompany.com.

   Mail Server Port

   Specifies the port of the SMTP mail server that sends the emails. An example is 25.

   Mail Recipients

   Lists users and administrators who receive emails about important events. Enter their names separated by semicolons, for example, Brown, Jane <broja@companyA.com>;Smith, John <John.Smith@company.net>.

   The configuration change takes effect when the next reservation is made.

Customize Reservation Ready User Notification Email

The Reservation Manager can send email notifications to end users when the systems associated with a reservation are fully configured and ready for use. You can also specify whether the administrator or root passwords for reserved systems are sent in this email notification. The text of the email can display a custom message. For example, you can configure this message to indicate the credentials required to log in to the reserved systems.

To customize reservation ready user notification email

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Notifications area. Each one opens a dialog in which you can change the value and click OK.

   Reservation Ready Text

   Specifies text for the first paragraph in a reservation ready email to users. In addition to the text in this field, the email contains additional information about the reservation that you cannot modify. Enter the text on one line or on multiple lines if the last character is a backslash (\). The following escape sequences are supported: \t: tab, \n: line feed, \r: enter, \f: form feed.

   Default: Your reservation request has been fulfilled.

   Reservation Ready Text Contains Password

   Specifies whether to include in the email the password that accesses reserved systems.

   Default: true

   Note: When this option is set to true, the user name that is included in the email is the user name defined with the dpmutil –set –superuser command. This command allows only a single username per operating environment. Therefore, enable this option only if the environment is configured to use a common administrator user name. If a site is provisioning systems with different Windows administrator account names, we recommend that you disable this option. Use the Reservation Ready Text option to provide end users the site-specific instructions for accessing provisioned systems.

   The configuration change takes effect when the next reservation is made.

Specify the Time Period for User Expiration Notification

The Reservation Manager can send one email notification or repeated notifications at multiple time intervals to users when the expiration time for a reservation is approaching. You can configure when the notification is sent relative to the time of expiration. By default, the notification is sent 24 hours before the reservation expires.

To specify the time period for user expiration notification

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Notifications area.

   Reservation Expiration Warning Times

   Specifies how many hours before reservation expiration to send a warning email to users. Expiration notifications can be sent at multiple intervals by specifying more than one value separated by commas. The number of notification times you can specify is unlimited. To send no emails, set this parameter to -1.

   Default: 24

   Limits: 10000 hours maximum

   Example: 48,36,24,12,10,8,6,4,2,1

5. Change the value or values, and click OK.

   The configuration change takes effect when the next reservation is made.

Configure User Notification Email for Job Failures

The Reservation Manager can send email notifications to end users when the provisioning job associated with a reservation fails. Only administrators are notified of provisioning failures, by default.

To configure user notification email for job failures

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Notifications area:

   Notify End User On Job Failure

   Specifies whether end users receive an email when a provisioning job fails.

   Default: false

5. Change the value and click OK

   The configuration change takes effect when the next reservation is made.

Specify When to Send a Stalled Task Alert

Reservation Manager can send an email alert to the administrator when a task is taking longer than expected. An alert is sent if the elapsed time since the last status update for a task was received and the present time exceeds a defined time interval.

You can configure this time interval, which is two hours by default. This configuration setting applies for all types of tasks, including operating system imaging, software installation, and so on. Therefore, make the time interval long enough for all ordinary tasks to complete.

To specify when to send a stalled task alert

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Notifications area.

   Task Status Update Timeout

   Specifies the time interval, in minutes, during which the status of a provisioning task must be updated. If the status is not updated in this time period, an email is sent to the administrator. Specify 0 to disable sending emails.

   Default: 120

5. Specify a value and click OK.

   The configuration change takes effect when the next reservation is made.

Modify the Physical System Allocation Policy

After Reservation Manager determines that a reservation request can be met, it calculates the total cost of each system that meets the criteria for the request. You can modify the weight values for each property in the caaipconf.cfg file to customize the system allocation policy for your environment.

To modify the physical system allocation policy

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Physical Provisioning area. Each one opens a dialog in which you can change the value and click OK.

   Reservation Weight - Cpu Count

   Specifies the relative weight assigned to each CPU for calculating the total cost of a system.

   Default: 500

   Reservation Weight - Memory

   Specifies the relative weight assigned to each GB of memory for calculating the total cost of a system.

   Default: 50

   Reservation Weight - Disk Space

   Specifies the relative weight assigned to each GB of disk space for calculating the total cost of a system.

   Default: 2

   The configuration change takes effect when the next reservation is made.

More information:

How the Low Cost Algorithm Works (see page 484)

How the Low Cost Algorithm Works

The weighted algorithm is used with a preferred pool policy. The administrator defines the order of resource pools to search for resources that match a user request. If systems are available in the primary pool, the weighted algorithm is applied to determine which system is the best match. Secondary resource pools are only searched if a match cannot be found in the primary pool, even if the secondary pools have a closer matching system based on the weighting algorithm.

A weighted algorithm selects the physical system that most closely matches the user requirements. This algorithm weights the following system properties by default:

- Number of CPUs—weights each CPU at 500
- Available memory—weights each available gigabyte of memory at 50
- Hard disk space—weights each gigabyte of hard disk space at 2

Therefore, one CPU costs approximately the same as 10 GB of RAM and 250 GB of hard disk space using the default policy.

Customize the Home Page

By default, end users have access to the following links on the Reservation Manager home page:

- Reserve a system
- Create a virtual machine
- Create a virtual machine in the Amazon Cloud
- View reservation templates
- View system inventory
- View your reservations

Administrators can customize the options that users can see on this page by removing access to any of these links. For example, you want your users to create only reservation requests based on public templates that you define. In this case, you could eliminate the Reserve a system, Create a virtual machine, and View system inventory links from the home page.

Note: You cannot make the Administer your Reservation Manager link available to end users. This link only appears for administrator users.

To customize the home page

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Quick Start area. Each one opens a dialog in which you can change the value and click OK. Set any of these options to false to remove the corresponding link from the end-users home page.

   Reserve Machine Access

   Controls whether the Reserve a system link appears on the home page.

   Default: true

   Reserve VM Access

   Controls whether the Create a virtual machine link appears on the home page.

   Default: true

   Create AMI Access

   Controls whether the Create a virtual machine in the Amazon cloud link appears on the home page.

   Default: true

   View Template Access

   Controls whether the View reservation templates link appears on the home page.

   Default: true

   View Reservation Access

   Controls whether the View your reservations link appears on the home page.

   Default: true

   View Machine Access

   Controls whether the View system inventory link appears on the home page.

   Default: true

   The configuration change takes effect when you restart the browser.

Configure Short Descriptions on the Home Page

Administrators can configure whether task descriptions on the home page are short (one sentence).

To configure short descriptions on the home page

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Quick Start area.

   Use Short Descriptions

   Indicates whether the task descriptions on the home page are displayed in short form (single sentence).

   Default: false

5. Change the value and click OK.

   The configuration change takes effect when you restart the browser.

Configure Reservations

Reservation Manager lets the administrator configure the following features for reservations:

- Specify the default duration for a reservation. This value is used to calculate the initial end date that is displayed to users when they make a reservation.
- Make project ID mandatory for reservation requests. If this option is set to true, users must enter a value in the project ID field before they can submit reservation requests. This parameter can be used to help ensure that project information is available for charging costs back to projects or for reporting on usage by project.
- Give users the ability to override the minimum hardware specification in a template.
- Give users the ability to select a network.

To configure reservations

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Reservations area. Each one opens a dialog in which you can change the value and click OK.

   Default Reservation Length

   Specifies the default reservation duration in hours.

   Default: 24

   Network Selection Allowed

   Indicates whether users can select a specific network when making reservations.

   Default: false (the user cannot select networks)

Project Id Is Mandatory

Specifies whether the user must enter a project ID when submitting reservation requests.

Default: false (the user does not enter a project ID)

Specify Requirements Allowed

Indicates whether users are allowed to override minimum hardware specifications when making reservations using templates. If set to true, users can change the template settings for hardware requirements or the systems used to fulfill reservations.

Default: false (the user cannot override hardware specifications)

The configuration changes take effect when the next reservation is made.

More information:

Allow Alternate Selection (see page 488)
Configure Chargeback Display (see page 489)
Override Automatic Selection (see page 490)
Specify Memory and CPU Selections (see page 491)

Allow Alternate Selection

Configure the Select Systems page of the Reserve a Machine wizard to display a list of systems that support the selected system image, and that are available for the requested dates, but do not satisfy one or more of the following criteria:

- Number of CPUs
- Minimum memory
- Minimum disk space

This is useful when requests for a specific system cannot be fulfilled.

If an acceptable alternate system is listed, the user can select it from the displayed list.

If the user requested more than one system, and the request could only be partially fulfilled, the list displays the automatically selected systems at the top. The user can select additional systems from the list and submit the reservation or return to the date specification page and try again.

Users are not limited to selecting the number of systems requested on the Specify Requirements page or when the template was created. As long as the number of systems does not exceed the maximum number the user is allowed based on the resource pool policy, the selected systems are reserved.

To allow alternate selection

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Reservations area:

   Specify Alternate System Specifications Allowed

   Specifies whether the Reservation Manager displays alternate systems if not enough systems meeting the minimum hardware requirements are available when making reservations.

   Default: true (alternate systems are displayed)

5. Enter a value and click OK.

   The configuration change takes place when the next reservation is made.

Configure Chargeback Display

The chargeback feature provides a way to charge on an hourly basis for the use of virtual machines. The product comes with pricing models for VMware and Amazon Cloud. The default is 0 (no hourly charge), and the administrator can change the value.

If chargeback is not used, the administrator can suppress the display of the hourly rate and total cost so that end users do not see it.

To suppress the display of chargeback hourly rate and total cost

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Reservations area.

   Display Chargeback Cost Information

   Indicates whether users see the chargeback hourly rate and total cost.

   Default: true (users see chargeback information)

5. Enter a value and click OK.

   The change takes effect when you restart the browser.

More information:

Create a Template for Amazon Machine Images (see page 549)

Override Automatic Selection

You can configure the Select Systems page for the Reserve a Machine wizard to display a list of all the systems that meet the request. Checkmarks indicate the systems that were automatically selected to fulfill the reservation. If necessary, end users can override this preselection by choosing other systems from the list.

Users are not limited to selecting the number of systems requested on the Specify Requirements page or when the template was created. As long as the number of systems does not exceed the maximum number the user is allowed based on the resource pool policy, the selected systems are reserved.

To override automatic selection

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Reservations area:

   Machine Selection Allowed

   Indicates whether users can select specific systems when reserving physical machines.

   Default: true (users can select specific systems or override the preselection)

5. Enter a value and click OK.

   The configuration change takes place when the next reservation is made.

Specify Memory and CPU Selections

Administrators can control the amount of memory and the number of CPUs that users can choose when they make a reservation.

To specify memory and CPU selections

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Reservations area. Each one opens a dialog in which you can change the value and click OK.

   Reservation Memory Selections

   Lists the amount of memory in gigabytes that users can request when making a reservation. The values must comma-separated whole numbers and fractions. An example is "10, 20.5" to put 10 GB and 20.5 GB in the memory drop-down lists.

   Default: 1, 2, 4, 8, 12, 16, 24, 32, 40, 48, 56, 64

   Reservation CPU Selections

   Lists the number of CPUs that users can request when making a reservation. The values must comma-separated integers. An example is "1, 2" to put 1 and 2 CPU choices in the drop-down lists.

   Default: 1, 2, 3, 4, 5, 6, 7, 8

   The changes take effect when the next reservation is made.

Set Limits on Virtual Machine Resources

Administrators can set limits on the number of VMs deployed simultaneously, the number of CPUs, the amount of memory, and the number of disks a user can request when reserving virtual machines.

To set limits on virtual machine resources

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Virtual Machines area. Each one opens a dialog in which you can change the value and click OK.

   Maximum Virtual Center Workload

   Controls how many VMs can be deployed simultaneously within one VMware vCenter or vSphere. By default, no limit is set. Setting a limit may be desirable to increase overall throughput in your vCenter environment. A value of 0 or -1 indicates that no limit is set. A value greater than 0 limits the number of simultaneous deployments within a vCenter to the specified value. The Apache service must be recycled for the changes to take effect.

   Default: -1 (no limit)

   Note: This configuration setting does not apply to Hyper-V.

   Virtual CPU Limit

   Specifies the maximum number of CPUs that users can request when making virtual machine reservations. A value of 0 or -1 indicates no limit.

   Default: -1 (no limit)

   Virtual Memory Limit

   Specifies the maximum amount of memory (MB) that users can request when making virtual machine reservations. A value of 0 or -1 indicates that no limit is set.

   Default: -1 (no limit)

Virtual Disk Limit

Specifies the maximum number of disks that users can request when making virtual machine reservations. A value of 0 or -1 indicates there is no limit as long as enough space is available.

Default: -1 (no limit)

Virtual Disk Space Limit

Specifies the maximum amount of disk space (GB) that users can request when making virtual machine reservations. A value of 0 or -1 indicates no limit.

Default: -1 (no limit)

The configuration changes take effect when the next reservation is made.

Add New Virtual Machines to a Service

Reservation Manager lets administrators specify whether new virtual machines are added to a CA Server Automation service automatically. This capability makes it easy for sites to monitor performance and usage of the virtual machines that have been reserved. Sites can also evaluate whether additional resources must be made available to improve performance, such as adding a VMware ESX server to a cluster. Sites can also see if virtual machines are under used and possible candidates for return. If this option is enabled, new virtual machines are added to the service with the same name as the resource pool. If the service does not exist, it is created.

To add new virtual machines to a service

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Virtual Machines area.

   Add Virtual Machine To Service

   Specifies whether to add the newly provisioned virtual machine to the service with the same name as the resource pool in which the virtual machine is created.

   Default: true

   The configuration change takes effect when the next reservation is made.

Specify the Maximum Number of NICs per Virtual Machine

Administrators can specify the maximum number of network adapters (often named network interface cards or NICs) that can be requested for a VMware virtual machine. The default and maximum is 10.

To specify the maximum number of network adapters per virtual machine

1. Log in to the Reservation Manager using the CA Server Automation administrator credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Virtual Machines area, and click OK.

   Virtual Network Interface Limit

   Specifies the maximum number of network adapters that users can request when creating virtual machines.

   Default and maximum allowed: 10

Configure Services

Reservation Manager lets the administrator configure the following features for services:

- Specify lower and upper thresholds.
- Specify lag.
- Enter a priority.

To configure services

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Virtual Machines area. Each one opens a dialog in which you can change the value and click OK.

   Service Lower Threshold

   Specifies the lower threshold for new services.

   Default: 20

   Service Upper Threshold

   Specifies the upper threshold for new services.

   Default: 80

   Service Lag

   Specifies the lag value for new services.

   Default: 2

   Service Priority

   Specifies the priority value for new services.

   Default: 1

   The configuration changes take effect when the next service is made.

Configure Snapshots

Reservation Manager lets the administrator configure the following features for services:

- Check for free space before taking snapshots.
- Specify growth percentage.
- Specify percentage of space reserved for managing snapshots.

Note: Reservation Manager permits one snapshot at a time. When you take a new snapshot, any previous snapshot is deleted so that there is enough space for the new one. This includes any snapshots created directly from vCenter or vSphere.

To configure snapshots

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following links in the Virtual Machines area. Each one opens a dialog in which you can change the value and click OK.

   Check For Free Space on Snapshot

   Specifies whether to check for sufficient disk space before taking snapshots.

   Default: true (checking is done)

   Filesystem Growth Allowance for Snapshots

   Specifies the percentage of potential growth for snapshots. For example, if disk space on a VM is 5 GB, and the value for this option is 20 (the default), one snapshot is allotted 1 GB of disk space on the ESX server. This option is useful only when Check For Free Space on Snapshot is true (the default).

   Default: 20 (percent of allocated disk space for VMs)

   Manage Snapshot Reserve

   Specifies the percentage of space reserved for managing snapshots.

   Default: 10 (percent of total capacity per host)

   The configuration changes take effect when the next snapshot is taken.

Disable Software Deployment

Administrators can configure Reservation Manager so that users cannot deploy software to virtual machines. A configuration setting can hide all software deployment features in the user interface.

To disable software deployment features

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the ITCM Integration area, and click OK.

   Software Deployment Enabled

   Enables or disables software deployment features in the Reservation Manager UI.

   Default: true (software deployment is allowed)

   The change takes effect when Reservation Manager is restarted.

Allow Users to Select Storage Tiers

Storage tiers are classifications for the data stores associated with each disk. Tiers generally indicate different levels of performance of the data store on which a VM and its hard drives are created. Administrators can enable or disable storage tiers.

When storage tiers are enabled, users can select them when they make virtual machine reservations.

To allow users to select storage tiers

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Virtual Machines area:

   Allow Data Store Tier Selection

   Indicates whether users can select a storage tier when they reserve a virtual machine.

   Default: Not Allowed (tiers not enabled)

   Note: When storage tiers are not allowed, administrators cannot use them when defining resource pools and chargeback policies.

5. Enter a value and click OK.

   The configuration change takes effect when the browser is restarted.

More information

Configure Chargeback for Storage Tiers (see page 564)

User Access to Reserved Systems

Users that reserve systems through the Reservation Manager must know the user name and password to access the systems. By default, the Reservation Manager communicates the user name and password in the email that is sent when the systems are ready. Reservation Manager must be able to determine what user name and password to include in this email.

To deploy software to systems being prepared for users, the Reservation Manager also must know the user name and password. The administrator account name used to deploy software is a privileged user, defined to Reservation Manager as the superuser. The superuser account names and passwords can be defined during installation or later using the dpmutil –set –superuser command. A superuser account name must be defined for each operating system. However, only one superuser account name can be defined for a single operating environment (for example, Windows). This single account name is used when accessing a newly imaged system to deploy software. If, for any operating environment, systems are configured with different administrator account names, attempts to deploy software to systems that are not configured with the defined superuser account name known to Reservation Manager will fail. To avoid deployment problems, we recommend that you set up an account to support the account name defined as the superuser for that operating environment. Alternatively, do not allow your end users to install software when choosing that system image.

How the Reservation Manager determines the valid credentials is described in the following sections.

JumpStart Provisioning Password Configuration

Bare metal provisioning of Solaris operating systems is implemented using Solaris JumpStart technology. The root password with which target Solaris computers are configured is defined as a hashed password in the JumpStart sysidcfg configuration file. All Solaris OS images that Reservation Manager uses must be configured to use the same root password.

Hyper-V Windows Provisioning Password Configuration

In a Hyper-V environment, you can set up things so that all users have the same credentials to access their Windows systems or you can let users choose their own passwords when submitting reservation requests. You specify which policy to use at the resource pool level.

If users are not allowed to choose their own password, issue the *dpmutil set superuser* command. This command stores the Windows administrator password in the database for later retrieval. For example:

```
dpmutil -set -superuser
```

The command prompts for credentials with administrative privileges.

The Reservation Manager retrieves the password when emails are sent to users notifying them that their systems are ready. The password that you specify overrides any password specified in the Hyper-V OS profile used to provision the Windows system.

If the *Allow user to specify the Administrator password* option is set on the resource pool, the person creating a reservation can input an Administrator password that is configured in the VM.

Any password specified in the Hyper-V OS profile is overridden.

OSIM Provisioning Password Configuration

Bare metal provisioning of Linux and Windows operating systems is implemented through the CA IT Client Manager OS Installation Management technology (OSIM). When you add an OS image to the OSIM library, one of the parameters that must be defined is the root or administrator password for target systems. All Windows OS images must be configured to use the same administrator password and all Linux OS images must be configured to use the same root password.

NIM Provisioning Password Configuration

Bare metal provisioning of AIX operating systems is implemented using the IBM Network Installation Management technology (NIM). The root password with which target AIX machines are configured is defined as a hashed password in the ca_post_install.sh script that is installed with the CA Server Automation NIM Adapter. All AIX OS images must be configured to use the same root password.

VMware Windows Provisioning Password Configuration

In a VMware environment, you can set up things so that all users have the same credentials to access their Windows systems or you can let users choose their own passwords when submitting reservation requests. You specify which policy to use at the resource pool level.

If users are not allowed to choose their own password, issue the dpmutil set superuser command to store the Windows administrator password in the database for later retrieval. For example:

```
dpmutil -set -superuser
```

The command prompts for credentials with administrative privileges.

The Reservation Manager retrieves the password when emails are sent to users notifying them that their systems are ready. The password that you specify must match the password specified in the VMware customization specifications that are used when provisioning Windows systems.

Customization specifications are defined using the VMware Infrastructure Client. When using this approach, the administrator password for the virtual machine used to create the template must be set to a blank or empty value. When set to blank, the password defined in the customization specification is set during virtual machine provisioning.

When you add virtual machine templates to the Reservation Manager inventory, associate a saved customization specification with each template. All Windows customization specifications must be configured to use the same administrator password.

If users are allowed to choose their own password, the Windows administrator account is configured with the password the user has entered. The password is encrypted and stored with other reservation data. The user-specified password overrides the password defined in the customization specification. The requirement for setting up a VM template with a blank administrator password is also required to let end users specify their own password.

VMware Linux Provisioning Password Configuration

The root password for Linux VMs is the password defined on the virtual machine before it was converted to a template. All Linux virtual machine templates must be configured to use the same root password.

Issue the following command to store the Linux root password in the database for later retrieval. The command prompts for information.

```
dpmutil -set -superuser
```

The Reservation Manager retrieves the password when emails are sent to users notifying them that their systems are ready.

EC2 Windows Provisioning Password Configuration

The administrator password for Windows instances is generated when an AMI is launched. Each Windows instance is assigned a unique password. The private key that is used to start the instance is required to retrieve and decrypt the administrator password.

To support this requirement, the Reservation Manager requires that the private key is stored using the dpmutil set ec2-private-keypair command. When the Reservation Manager detects that a Windows instance is running, it initiates an action to retrieve the password using the stored private key. The Windows administrator password is included in the email notification sent to the end user that requested the system.

The private key that is used when starting an instance is determined by the EC2 resource pool with which the new instance is associated.

Note: The Windows administrator password defined with the dpmutil set superuser command is not used for accessing Windows instances.

EC2 Linux Provisioning Password Configuration

AMI instances that are running a Linux OS are accessed using SSH. When issuing the ssh command or using an SSH tool such as PuTTY, users must log in to the instance with the private key that was used to start the instance. The Reservation Manager does not make the private key available to end users. The administrator is responsible for providing the private key to end users. The EC2 resource pool determines the private key that is used to start an instance, so you can use multiple private keys at a site. Create separate EC2 resource pools and assign them different private keys to limit access. Supply members of each group with the private key file that has been assigned to the EC2 resource pool to which they have access.

Note: The Linux root password defined with the dpmutil set superuser command is not used for accessing Linux instances.

Access the Reservation Manager User Interface

Access the Reservation Manager user interface to configure organizational units, user access, system and virtual machine availability, and supported software, images, and templates for use in reservations.

To access the user interface

1. Select Start, Programs, CA, CA Server Automation, Launch CA Server Automation Reservation Manager from the Reservation Manager server.

   The Reservation Manager login page appears at the following URL:

   `https://servername:port/ssrm/`

   servername

   Specifies the name of the Reservation Manager server.

   port

   Specifies the port that the server is listening on.

   Default: 8443

2. Enter your admin login credentials and click Log In.

   The home page appears.

3. Click Administer your Reservation Manager. This link is only available to administrators.

   The Administration page appears. Perform all administrative tasks from this page.

The Start menu shortcut is only available on the Reservation Manager server. Users accessing the interface from a separate server must enter the URL in a web browser.

Filter Displayed Data

If the user interface displays data in table format, you can define filter conditions for each column of the table. The table then displays only the required amount of data.

To filter displayed data

1. Click Show Content at the Filter pane.

   The filter pane expands and the parameter fields appear that you can use to filter the displayed data.

2. Specify appropriate filter conditions and click Apply Filter.

   The table displays the data according to your filter condition.

Post-Installation Administrative Tasks

After installation, complete the following activities to prepare the Reservation Manager for end users:

- Set up organizational units to control user access to systems and system images.
- Define one or more resource pools and specify access policies.
- Add the systems that you want to make available to users to the inventory, classify them, and associate them with one or more resource pools.
- Identify operating system images (and virtual system templates) that you want to make available to users and specify access policies.
- Identify software packages that you want to make available to users and specify access policies.

There is no specific order in which you must complete these activities. The following instructions show how to prepare the Reservation Manager for end users so that they can reserve physical or virtual systems.

Organizational Units

An *organizational unit* (org unit) is a group of users. Org units provide security by giving users access to objects like resource pools, software groups, system images, and templates.

Consider the following information about org units:

- Users can belong to more than one org unit. They can switch to a different org unit by clicking the *Member of* link at the top left of the Reservation Manager window.
- Access to resources can be tailored for each org unit.

- Membership in org units can be based on the following properties of a CA EEM user. Reservation Manager administrators do not have to duplicate the CA EEM setup.
  - Global groups to which users belong
  - CA application groups to which users belong
  - Attributes like department, company, office, city, state, or country. So, all users in North American Sales Support could be members of an org unit named Sales.
- If the administrator removes a user from an org unit, that user can continue working in that org unit until one of the following situations occurs:
  - The user logs out and logs in again.
  - The user clicks the *Member of* link to change their org unit.
- CA EEM supports Native, Active Directory and LDAP.
- A default org unit, Public, is for users who do not belong to any other org unit. Do not explicitly add users to the Public org unit. Do add resource access to it, however, if you want to enable all CA EEM global users to make reservation requests.
- Use descriptive names for org units so that users can identify them easily.

Create an Organizational Unit

An organizational unit gives users access to Reservation Manager features like resource pools and templates.

To create an organizational unit

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click the Administer Your Reservation Manager link.

   The Administration page opens.

3. Click Manage your organizational units.

   The Organizational Units page opens.

4. Select Add from the Actions menu in the upper right corner of the list.

Add Users To An Organizational Unit

You can add CA EEM users to an existing org unit.

To add users to an organizational unit

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click the Administer Your Reservation Manager link.

   The Administration page opens.

3. Click Manage your organizational units.

   The Organizational Units page opens.

4. Select an org unit and then select Details from the Actions menu in the upper right corner of the list.

   The Organizational Unit Details page opens. This page contains several tabs on which you can modify org unit settings.

5. Click Members.

6. Select the type of property to search for (Users or User Attributes), the attribute to search for, and an operator. Enter a value and click Search.

   Reservation Manager queries CA EEM for all users or attributes that match the search criteria and displays the results in the Available Users/Attributes list.

7. Select users or attributes, click the right-arrow to move the names to the Select Users list, and then click OK.

   The Reservation Manager adds the selected users or attributes to the organizational unit and displays a confirmation message.

Let Users Perform Some Administrative Tasks

Administrators can give users permission to perform some administrative tasks on VMware virtual machines. Users can then perform the following actions without contacting an administrator:

- Change the administrative password when creating a virtual machine.
- Control the power status of a virtual machine. The administrator can let users turn on, turn off, and suspend virtual machines that are assigned to them.
- Take a snapshot of a virtual machine. A *snapshot* is a record of a virtual machine at a certain point in time.
- Change the configuration of a virtual machine after the original provisioning.

To let users perform some administrative tasks

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens, and lists existing pools.

4. Double-click a resource pool.

   The Resource Pool details page opens, with the Properties tab displayed.

5. Select or clear the following fields:

   Allow user to specify the Administrator password

   Lets users set the administrator password.

   Default: Off

   Note: If you enable this option, the password must be blank in templates that you create.

   Allow users to manage VM power state

   Lets users turn on, turn off, and suspend virtual machines that are assigned to them.

   Allow users to take a VM snapshot

   Lets users make a record of a virtual machine at a certain point in time. You can also specify the maximum number of snapshots per VM (up to 10), and indicate whether the VM slows down to conserve power when not in use.

   Allow users to change the VM configuration

   Lets users change such things as reservation start and end dates, the notification email address, and notes.

6. Click OK.

   The permissions are granted.

Set Over Commitment of Memory on ESX Server or Cluster

You can increase the amount of memory used on VMware servers or clusters when the actual memory usage permits it. Doing this lets you deploy more virtual machines.

Overcommitment is specified as a percentage. For example, if an ESX server has 30 GB of physical memory available for the VMs it hosts, an overcommitment of 50 percent would increase the memory to 45 GB.

To set overcommitment of memory on an ESX server or cluster

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens, and lists existing pools.

4. Double-click a resource pool.

   The Resource Pool details page opens, with the Properties tab displayed.

5. Select the following field and enter a percentage:

   Allow memory overcommitment. Percent:

   Enables over commitment of memory by a percentage that you specify.

   Default: Off, 0%

   Note: A negative value sets undercommitment of memory.

6. Click OK.

   Memory is overcommitted.

Multi-Tenancy Environment

A *tenant* is a collection of users and resources. The tenant users are a subset of all CA EEM global users. The tenant resources are a subset of all resources that CA Server Automation discovers and manages.

In environments with tenants, there can be two types of administrators:

- super administrator
- tenant administrator

The super administrator can perform all administrative tasks in Reservation Manager. The super administrator specifies the users in a tenant and designates one or more tenant administrators to manage each tenant. Thus, in a large environment, the super administrator can delegate some day-to-day activities to the tenant administrators.

The tenant administrator can administer only the resources belonging to a tenant. These administrators perform a subset of administrative tasks on behalf of the users in their tenant.

The end users belonging to a tenant are users of Reservation Manager only. They log in to Reservation Manager to make reservation requests.

Super Administrator

A super administrator is a member of the user group AIPAdmins, and therefore can perform all administrative tasks in both CA Server Automation and Reservation Manager.

The super administrator is responsible for creating the objects and configuration that define tenant users, resources, and administrators. The super administrator performs the following actions to define a new tenant:

- Create one or more services for the tenant in CA Server Automation. Add tenant physical computer systems to these services.
- Create one or more physical resource pools for the tenant in Reservation Manager by importing the tenant services from CA Server Automation.
- Create one or more virtual resource pools for the tenant in Reservation Manager.

  Note: When you create a pool of this type, Reservation Manager automatically creates a service with the same name in CA Server Automation. Later, when an end user creates a virtual machine in a virtual resource pool, Reservation Manager automatically adds the VM to the corresponding service.

- Create a tenant in Reservation Manager, specify access to resource pools, and specify administrators and users.
- Create system images, software groups, and reservation templates for the tenant in Reservation Manager. These objects should reference resources owned by the tenant or available to the members of the tenant. Members of the tenant use these resources when making reservation requests.

  Note: Tenant administrators can also create reservation templates. Templates that the super administrator creates can be assigned to one or more tenants. Templates that tenant administrators create can be used only by members of that tenant.
- Grant access to these objects to the tenant in Reservation Manager. This lets the tenant administrator give tenant users access to these objects.
- Specify which network definitions are available to each tenant.
- View or hide tenants, or filter the view by tenant name.

The maintenance of a tenant consists of adapting the previous configuration to changes to tenant users and resources:

- Maintain tenant administrators.
- Modify services and create new services for the tenant in CA Server Automation as needed.
- Modify or create system images, software groups, and reservation templates for the tenant in Reservation Manager as needed.
- Modify tenant access to resources as needed.
- Modify membership in the tenant as needed.

Note: The super administrator cannot create organizational units for tenants. The tenant administrator is solely responsible for creating organizational units for the members of the tenant.

Add a Tenant

The super administrator can create new tenants, add tenant administrators and end users, and give access to resource pools and other resources.

To add a tenant

1. Log in to the Reservation Manager using the super administrator credentials.

   The home page opens.
2. Click the Administer Your Reservation Manager link.

   The Administration page opens.
3. Click Manage your tenants.

   The Tenants list opens.

4. Select Add from the Actions menu in the upper right corner of the list.

   The Add Tenant wizard opens.

5. Add information to the following pages as instructed in the wizard.

   Define Tenant

   Enter a name for the tenant and optional description.

   Select Administrators

   Search for tenant administrators based on ID (name) or on a user attribute like group membership, department, city, and so on. Specify the attribute to search for, and a comparison operator (EQUAL or LIKE). Enter a value and click Search. (With the LIKE operator, you can enter the wildcard character * to represent a sequence of characters.)

   Reservation Manager queries CA EEM for all users that match the search criteria and displays the results in the Available Users/Attributes list. To add administrators to the tenant, select users in the Available Users list and click the right arrow to move them to the Selected Users/Attributes list.

   Note: A user can administer only one tenant. However, a tenant can have multiple administrators.

   Select Users/Attributes

   You can specify members of a tenant either explicitly by user ID or implicitly by user attribute. Implicit membership can be a real time-saver because it reduces the amount of work necessary to define and maintain membership. Implicit tenant membership is based on any of these user attributes: global group membership, department, office, city, state, country.

   To specify explicit membership, search for users based on ID (name). Specify identity type of user.

   To specify implicit membership, search for user attributes based on an attribute value like group membership, department, city, and so on. Specify Identity Type of User Attribute. Select the attribute to search for, and a comparison operator (EQUAL or LIKE). Enter a value and click Search. (With the LIKE operator, you can enter the wildcard character * to match any sequence of characters.)

   Reservation Manager queries CA EEM for all users or attributes that match the search criteria and displays the results in the Available Users/Attributes list. Select users or attributes and click the right arrow to move them to the Selected Users/Attributes list.

   Note: Users can belong to more than one tenant.

Specify Pool Access

To grant a tenant access to a resource pool, select the pool and then select the type of access from the drop-down list in the Access column. Shared access lets you grant multiple tenants access to the resource pool. Tenants with shared access to a pool are not permitted to change the pool settings. Exclusive access lets only one tenant use the resource pool. A tenant administrator with exclusive access can modify the settings for the pool.

If your choice of shared or exclusive access conflicts with another tenant, a message explains the conflict. Change your choice before completing the tenant definition.

Specify Image Access

To grant the tenant access to a system image, select the image in the Available System Images list. Click the right arrow to move it to the Selected System Images list.

Specify Software

To grant the tenant access to software, select software from the Available Software Group list. Click the right arrow to move it to the Selected Software Group list.

Manage Reservation Template Access

To grant the tenant access to a template, select it in the Available Reservation Templates list. Click the right arrow to move it to the Selected Reservation Templates list.

Edit a Tenant

The super administrator can edit tenant properties, administrators and end users, and access to resource pools.

To edit a tenant

1. Log in to the Reservation Manager using the CA Server Automation super administrator credentials.

   The home page opens.

2. Click the Administer Your Reservation Manager link.

   The Administration page opens.

3. Click Manage your tenants.

   The Tenants list opens.

4. Click the name of a tenant in the Tenant list.

   The Tenant details pane opens.

5. Update information about the following tabs.

Properties

Update the optional description.

Administrators

Add or remove tenant administrators.

To add a tenant administrator, select an attribute to search for, an operator, and a value. Click Search. Select a user ID in the Available Users list and click the right arrow to move it to the Selected Users list.

To remove a tenant administrator, select the user ID in the Selected Users list. Click the left arrow to move it to the Available Users list.

Members

Add or remove tenant members.

You can specify members of a tenant either explicitly by user ID or implicitly by user attribute. Implicit membership can be a real time-saver because it reduces the amount of work necessary to define and maintain membership. Implicit tenant membership is based on any of these user attributes: global group membership, department, office, city, state, country.

To add a member, select the type of identity to search for S(Users or User Attributes), an attribute, and an operator. Enter a value and click Search. Select a user ID or attribute in the Available Users/Attributes list, and click the right arrow to move it to the Selected Users/Attributes list.

To remove a member, select the user ID or attribute in the Selected Users/Attributes list. Click the left arrow to move it to the Available Users/Attributes list.

Pool Access

Select a pool, and select the type of access from the drop-down list in the Access column. Shared access lets multiple tenants use a resource pool. Exclusive access lets only one tenant use the resource pool.

Image Access

Select a system image in the Available System Images list, and move it to the Selected System Images list.

Reservation Template Access

Select a template in the Available Reservation Templates list, and move it to the Selected Reservation Templates list.

Network Access

Select a network in the Available Networks list, and move it to the Selected Networks list.

Configuration Properties

Select a Property name to view or change the value of a property. In the Edit dialog, use the Effective Setting drop-down to specify Global or Tenant.

Global

Displays the current value of the global setting. For this tenant, the value of this setting is same as the global setting. Whenever the global setting is changed, the value of the setting changes accordingly.

Tenant

Allows you to view and change the tenant-specific value of this setting. The value of this setting is independent of the value of the corresponding global setting.

Tenant Administrator

A *tenant administrator* has a restricted (that is, scoped) role that is limited to the users and resources belonging to a single tenant.

The super administrator defines the tenant, specifies the membership of the tenant, specifies the resources that the tenant can access, and specifies the users who administer the tenant.

The duties of the tenant administrator are:

- View and modify physical and virtual resource pools created by the super administrator and assigned to the tenant. The tenant administrator can perform the following actions on resource pools:
  - Set a limit on how many systems users can reserve
  - Set a limit on how long users can reserve machines
  - Specify VM naming rules
  - Specify whether reservations require manual approval
  - Specify whether users are allowed to set the Windows administrator password
  - Specify whether users are allowed to issue VM power operations

- Specify whether users are allowed to take snapshots of VMware virtual machines
- Specify whether VMs are deleted when reclaimed
- Specify the level of memory overcommitment

Note: When multiple tenants share a resource pool, tenant administrators cannot modify the pool settings. They can view the pool and its settings and assign the pool to tenant organizational units.

If a resource pool is assigned exclusively to a tenant, the tenant administrator can modify all pool settings with some exceptions. Tenant administrators cannot modify two for ITCM (ITCM Domain Manager and Scalability Server), one for vCenter (Folder placement), and two for Amazon Cloud (Keypair Name and Network Selection). Tenant administrators can view these settings, however.

Tenant administrators cannot create or delete resource pools.

- Create and manage organizational units, which define the access of tenant members to Reservation Manager resources. Types of resources include reservation templates, system images, resource pools, and software groups. Resources that the tenant can access are available for assignment to tenant organizational units.

  Tenant administrators can add tenant members and network definitions to organizational units.

- Create and manage reservation templates. The tenant administrator can create, modify, and delete templates that tenant members use when making reservations. Organizational units control member access to templates.

- Perform operations on reservations that are based on resources in designated services. The tenant administrator can perform the following actions on reservations:
  - Approve or reject reservation requests
  - Extend reservations
  - Cancel reservations to reclaim resources
  - Check reservation status
  - Restart or skip reservation tasks

- View system inventory and check system availability for systems that belong to resource pools derived from designated services.

When a tenant administrator clicks *Administer your Reservation Manager* on the home page, the following links appear. These links let the tenant administrator perform all the previously described tasks.

- View all reservations
- Manage your system inventory
- Manage your resource pools
- Manage your reservation templates
- Manage your organizational units

Tenant End User

The experience of users who belong to a tenant in Reservation Manager is identical to that of users who do not belong to a tenant. Membership in an organizational unit determines the reservation templates, system images, software groups, and physical systems available to a tenant user for making reservations.

VLAN Scoping

Reservation Manager administrators must scope which VLANs are accessible for end user selection. This is done by specifying the network address pools that are accessible to members of each organizational unit.

To give users access to VLANs

1. Select Administration, Manage your organizational units.
2. Add a new org unit or open an existing one.
3. Go to the Network Access tab, and select one or more VLANs.

Make Physical Systems Available to Users

Use the following process to make physical systems available for reservation:

1. Identify the CA Server Automation service that contains the physical systems that you want to share.
2. Import the service to create a Reservation Manager resource pool.
3. Define the operating system images that a user can install on the available systems.
4. Set up an access policy for both the resource pools and system images.

The following sections describe a quick way to set the Reservation Manager up to support reserving physical systems.

Prerequisites for Supporting Reservations of Physical Systems

Before you access the Reservation Manager for the first time, use the Add Resource function in CA Server Automation to perform a one-time image installation of each system that will be added to the Reservation Manager inventory. After all the systems are registered, use CA Server Automation to create a service and add the systems to the new service.

Note: The Add Resource function configures CA Server Automation with information that the Reservation Manager needs to reimage these systems successfully later.

We recommend that you install a CA Configuration Automation agent on all servers in the inventory to collect the system configuration information that the Reservation Manager requires. Specifically, the Reservation Manager requires information about the number of CPUs, the amount of memory, and the amount of disk space for each system.

Hardware Classes

The hardware class relations specify which operating systems can be installed on a system. You create these relations during adding systems and system images to your inventory. When you add a system to the inventory, then you associate that system with a single hardware class. However, when you add a system image to the inventory, then you can associate it with one or more hardware classes. The following example illustrates these relations:

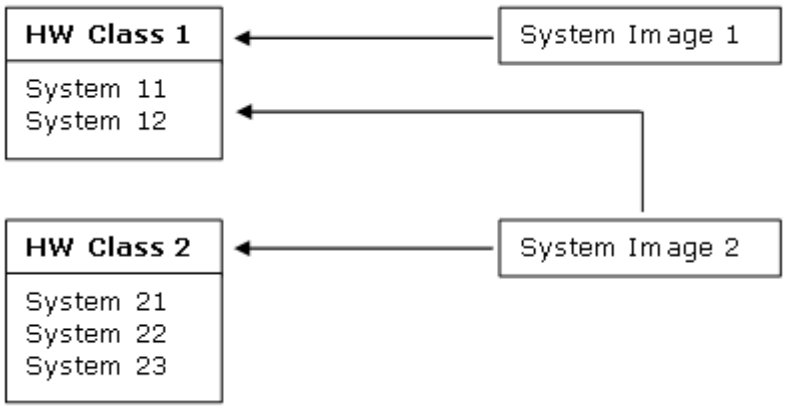

System Image 1 is associated with HW Class 1, so System Image 1 can be applied to System 11 and System 12.

System Image 2 is associated with HW Class 1 and HW Class 2, so System Image 2 can be applied to System 11, System 12, System 21, System 22, and System 23.

Create a Resource Pool Using a Service

You can set up a resource pool quickly by linking it to a CA Server Automation Service.

To create a resource pool using a CA Server Automation Service

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Select Import Service from the Actions drop-down menu in the upper right corner of the Resource Pools list.

   The Import Service wizard opens. Use the wizard to populate the Reservation Manager inventory with the physical systems that are members of the CA Server Automation service that you created.

5. Complete the following fields on the Select Service wizard page, and click Next.

   Service

   Specifies a CA Server Automation service to import.

   Description

   Defines a description for the resource pool.

   Maximum Days

   Defines the maximum number of days for which a user can reserve a system in this pool when submitting a single request.

   Maximum Systems

   Defines the maximum number of systems in this pool for which a user can simultaneously hold reservations.

   Scalability Server

   Defines the name of the server that will be used for deploying software to computers associated with this pool.

   Keep Synchronized with Service

   Specifies whether the resource pool remains synchronized with the CA Server Automation service. Synchronization updates this resource pool when systems are added or removed from the CA Server Automation service.

   The Specify Access Policy wizard page opens.

6. Select the organizational units whose members will be granted access to the systems in the resource pool.

   Note: Before you grant access to users, grant the rmadministrators resource group access so you can test creating reservations for systems in this pool.

   If an organizational unit that has been granted access to the resource pool is not listed or you are unsure which organization units to grant access to the resource pool, click Finish to skip this step. You can grant access to the resource pool later.

   After you have selected all required organizational units, click the right-arrow.

   The organizational units are moved to the Selected Organizational Units list.

7. Click Finish.

   The Reservation Manager creates a resource pool with the same name as the CA Server Automation Service and imports all systems in the CA Server Automation Service to its inventory. A confirmation message displays when the process completes successfully. The wizard closes and the Import Service page opens to display the Resource Pools list.

View Systems in the Reservation Manager Inventory

The following procedure explains how to view systems in the Reservation Manager inventory.

To view systems in the Reservation Manager inventory

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click the Manage your system inventory link.

   The System Inventory page opens. For each system in the Reservation Manager inventory, the System Inventory page lists the system name, the hardware class, model, processor, the number of CPUs, the amount of memory, location, and the amount of disk space each system has.

   Note: If a CA Configuration Automation agent has been installed on these systems, the CPUs, Memory (in megabytes), and Disk Space (in megabytes) values must already be set. On systems without a CA Configuration Automation agent, set these attributes before users can reserve the systems.

4. (Optional) If you need more details about the system, click the system name in the Name column.

   The System Details page opens for the selected system and displays system properties and associated resource pools. The Properties tab provides additional information like Location, Serial Number, IP Address, and MAC Address.

Modify System Attribute Values

Set attribute values for the systems in the Reservation Manager inventory before users start reserving the systems.

Note: If a CA Configuration Automation agent has been installed on these systems, most of the relevant attribute values must already be set.

To modify system attribute values

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your system inventory.

   The System Inventory page opens. For each system in the Reservation Manager inventory, the System Inventory page lists the system name, the hardware class, model, processor, the number of CPUs, the amount of memory, location, and the amount of disk space on the system.

4. Select the check box next to a system name for which to modify attribute values, then select Details from the Action menu.

   The System Details page opens for the selected system and displays the following tabs:

   - Properties
   - Resource Pools

5. Click Properties and modify the following settings as necessary:

   Hardware Class

   Specifies which operating systems can be installed on a system. The available Hardware Class values are site-specific. Select the value that seems most appropriate for the selected system or define a new class type by editing the field. When you add operating system images to the inventory, you associate each image with the hardware classes with which they are compatible.

   Note: You cannot delete a hardware class after it is created.

   Complete the remaining fields.

6. Click the Resource Pools tab and modify the list of selected resource pools for that system if necessary.

7. Click OK.

   Reservation Manager saves the changes, closes the System Details page, and opens the System Inventory page to display attributes you have set for the selected system.

More information:

Hardware Classes (see page 515)
Prerequisites for Supporting Reservations of Physical Systems (see page 515)

Define Your JumpStart Boot Servers

If users are allowed to request the installation of Solaris system images, administrators must configure Reservation Manager so that it can identify the required JumpStart boot servers.

The Reservation Manager uses Solaris JumpStart boot server technology for initiating the installation of Solaris operating systems on servers being set up for users. Because the Reservation Manager supports environments that have multiple JumpStart boot servers and installation servers, identify all the JumpStart boot servers required to support the systems that can be allocated to users.

To define your JumpStart boot servers

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer your Reservation Manager.

   The Administration page opens.

3. Click Manage your JumpStart boot servers.

   The JumpStart Boot Servers page opens. This page initially displays an empty table as long as no JumpStart boot servers are defined. If JumpStart boot servers have already been added, the page lists the available JumpStart boot servers and their associated IP masks, descriptions, and locations.

4. Select Add from the Actions menu to add new JumpStart boot servers.

   The Add JumpStart Boot Server wizard opens.

5. Complete the following fields to select and specify the JumpStart boot server to add to the inventory, then click Next.

   Name

   Specifies the JumpStart boot server. Select a boot server from the drop-down list. This list contains the JumpStart boot servers that were defined during the installation or afterward using the dpmutil command. The list also includes all the servers identified as JumpStart install servers.

   Description

   Lets you provide additional information about the JumpStart boot server.

   IP Mask

   Specifies an IP mask used to identify the computers that each JumpStart boot server will manage.

   IP masks can contain wildcards, IP ranges, or CIDR (slash) notation. Reservation Manager supports the ? and * as wildcard characters. Additionally you can use the following formats:

   - IP ranges: x.x.x.{x-x}
   - CIDR (slash) notation: x.x.x.x/x.

   Examples:

192.168.1.*	Wildcard notation
192.168.1.{1-127}	IP range notation
192.168.1.0/24	CIDR notation

   Location

   Identifies the location of the JumpStart boot server.

   The Select Install Servers page appears.

6. Select the JumpStart install servers that a boot server can use during the provisioning process, then click Finish.

   Note: If the boot server is also an install server, select it.

   The JumpStart Boot Servers page opens and lists the JumpStart boot server that has been successfully added.

Make Operating System Images Available to Users

One or more operating system images must be available for users to select.

To make operating system images available to users

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your system image inventory.

   The System Images page opens. This page lists the inventory of operating system images and virtual machine templates that are available for users to select when reserving systems. From this page, you can define which operating system images or virtual machine templates to make available to users.

4. Select the Add Image item from the Actions drop-down menu to add new operating system images to the inventory.

   The Add System Image wizard opens.

5. Complete the following fields to select the operating system image to add to the inventory, then click Next. Select the operating environment first and continue according to the list.

   Platform

   Specifies the operating system.

   Image Name

   Specifies the image for the selected operating system.

   Boot Image Name (Solaris only)

   Specifies the boot image associated with the selected operating system image.

   Note: The Image Name combo box contains a list of all installation images that reside on the JumpStart servers with which CA Server Automation is configured. If an installation image resides on multiple JumpStart install servers, the boot image associated with this installation image must have the same name on all boot servers.

Image Description

Specifies an image description displayed to users.

The Specify Hardware Classes page opens.

Note: If Reservation Manager cannot locate any system images, it displays a warning. If this situation occurs, confirm that the CA Server Automation has been configured to support imaging of one or more operating environment types. Another reason can be that you have already added all the images available for the selected operating environment type.

The Packaging component must be installed and properly configured to support Linux, Windows, or Solaris imaging. To support Solaris imaging, CA Server Automation must be configured to communicate with one or more JumpStart servers.

6. Select all the hardware classes in the Available Hardware Classes pane that support the operating system image you previously selected. Move them to the Selected Hardware Classes pane. Click Next.

   Note: Verify that you select the appropriate hardware classes to prevent provisioning job failures resulting from incompatibilities.

   The Specify Software page opens.

7. Move all software that you want to allow users to install on this image from the Available Software Group to the Selected Software Group.

   The Available Software Group list contains software that the administrator has identified as being compatible with the selected operating system. The list is filtered to display only software that you are authorized to access. Authorization is based on the organizational unit that your user name has been assigned to.

   The selected software will be installed on the system in the order it is listed in the Selected Software Group list. Use the  up and  down arrow buttons to modify the order of software in the list.

8. Click Next when finished. You do not have to select any software on this page.

   The Specify Access Policy page opens.

9. Select the organizational units whose members will be granted access to this operating system image, then click Finish.

   Note: Before you grant access to users, grant the rmadministrators resource group access so you can test creating reservations for systems using this image.

   If an organizational unit that has been granted access to the image is not listed or you are unsure which organizational units to grant access, click Finish to skip this step. You can grant access to the image later.

   The System Images page opens and lists the operating system image that has been successfully added to the inventory.

More information:

Hardware Classes (see page 515)

Reserve a System

To verify that the Reservation Manager is configured to reserve physical systems from the resource pool that you set up, click the Home breadcrumb to return to the home page.

You can reserve a physical system from an allocated resource pool for use at a specific date. You define the system requirements, such as operating system, software, and hardware. Reservation Manager verifies that sufficient resources exist and schedules the availability and provisioning of the resource.

To reserve a system

1. Click Reserve a system from the home page of the Reservation Manager user interface.

   The Reserve Systems wizard appears with the Specify System Image page displayed.

2. Select the operating system image and click Next.

   Note: You can filter the table on this page by image name, description, or operating environment.

   The Specify Software page appears.

3. Move all software that you want to install on the system from the Available Software Group to the Selected Software Group.

   The Available Software Group lists software that the administrator has identified as compatible with the selected operating system image. The list is filtered to display only software that you are authorized to access. Authorization is based on the organizational unit that your user name is assigned to.

   The selected software is installed on the system in the order it is listed in the Selected Software Group list. Use the  up and  down arrow buttons to modify the order of software in the list.

   Click Next when finished. You do not have to select any software on this page.

   The Specify Requirements page appears.

4. Specify the number of physical systems that you want to reserve and the minimum hardware requirements for the systems in the provided fields.

   Note: The fields prevent you from exceeding the resources available.

   Click Next when finished.

   The Specify Dates page appears.

5. Choose to reserve now for several days, or enter start and end dates and times. Both options also let you reserve a system with no expiration date and time.

   The times start and end on the hour. The end time must be at least an hour after the start time. If you specify an end time that exceeds the number of days for which you can reserve equipment, the request is rejected. Limits on reservation time are defined in policy at the resource pool level.

   Click Next when finished.

   Reservation Manager uses all defined information to determine whether the request can be fulfilled. One of the following occurs:

   - If the resources are available to fulfill the request, the Complete Request page appears.
   - If the requested resources are not available, the request is rejected. Modify the request to fit within the available resources.

6. Select a system and click Next.

   Note: This step may not be available if the administrator made selecting systems unavailable. See Override Automatic Selection (see page 490).

7. Enter information in the following optional fields as necessary:

   Project ID

   (Optional) Defines the ID that associates the request with a specific project. This field is useful for evaluating requests and for allocating or charging costs to a specific project.

   The Reservation Manager administrator can configure this field so that it is required. In this case, the Reservation Manager displays a corresponding note in the information bar. You can submit your reservation request only after you enter a project ID. The Reservation Manager does not validate the entered value.

   Notes

   (Optional) Defines additional information you want to provide about the request.

   Send Notifications to

   (Optional) Defines the email address to which Reservation Manager sends notifications regarding this request. You can override any default value with a different address. Reservation Manager sends email notifications when a reserved system is available for use and when the expiration of the reservation period is approaching.

Save as Template

(Optional) Defines a name for this request so that you can use it as a template for future requests.

Template Description

(Optional) Defines a description for the template.

Click Finish.

The request is submitted for approval, and a final check is performed to verify that the machines are still available. If the resources are available, one of the following events occurs:

- If Reservation Manager is configured for automatic request approval, the reservation immediately reflects that it has been approved.
- If Reservation Manager is configured for manual request approval, an email notification is sent to the administrator to approve the request.

8. Verify the reservation status on the View reservation details page or on the View all reservations page that is only available for administrators. The Job Status values in the table indicate the progress of the provisioning job. The options are as follows:

Not Started

Indicates that the job status remains in this state until the scheduled reservation start date.

Initiated

Indicates that the reservation start date has been detected and processing of the reservation request has started.

In Progress

Indicates that the provisioning of the first system has started.

Ready

Indicates that all systems requested have been set up and are available for use.

Failed

Indicates that there was a failure processing one or more tasks associated with preparing the requested systems.

Canceled

Indicates that the user or administrator requested cancellation for the reservation.

Expired

Indicates that the reservation end date has been reached.

9. (Optional) Select the reservation item in the table and select the Details item on the Action menu.

   Reservation Manager displays the details of this request.

   The following Job Status values in the administrator Details page appear as hyperlinks:

   - Not Started
   - Initiated
   - In Progress
   - Ready
   - Failed

10. (Optional) Click a link to go to a subordinate page that provides further details about the job status.

Specify a Folder for Virtual Machines

You can specify a folder for newly created VMware virtual machines. Folders make it easier to manage many VMs when using management tools like vSphere Client.

Note the following information about folders on the vCenter:

- Folders must be created from the “Virtual Machines & Templates” view in vCenter because they are a different type of folder from the ones created in the “Hosts and Clusters” view.
- Folders must exist on the vCenter before you create any reservations using this resource pool.

To specify a folder for newly created VMware virtual machines

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens, and lists existing pools.

4. Double-click a resource pool.

   The Resource Pool details page opens, with the Properties tab displayed.

5. Enter a name in the Folder field.

   Note: Folder names must be a nonempty string of less than 80 characters. The slash (/), backslash (\), and percent (%) are escaped using the URL convention (example: %2F). Folders in the same hierarchy cannot have the same name. You can specify folder/folder.

6. Click OK.

   New VMs are placed in the specified folder.

Public Templates for End Users

Administrators can create public templates for end users that determine how systems are allocated. The templates are used to create new virtual machines to fulfill reservation requests.

The Reservation Manager provides a wizard for creating templates.

More information:

Create a Template for the Image Type (see page 527)
Create a Template for VMware Virtual Machines (see page 536)
Create a Template for Hyper-V Virtual Machines (see page 545)
Create a Template for Amazon Machine Images (see page 549)
Create a Template for IBM PowerVM Logical Partitions (see page 557)

Create a Template for the Image Type

Administrators create one or more templates so that users can make reservations. This procedure shows how to create a template when the system image is the image type (as opposed to Amazon, Hyper-V, or VMware).

To create a template for the system image type image

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your reservation templates.

   The Resource Pools page opens.

4. Select Create from the Actions menu in the upper right corner.

   The Create Reservation Template wizard opens.

5. Follow the instructions in the wizard. The following pages are provided:

   Specify Name page

   Indicates a name and optional description.

   Specify System Image page

   Indicates the operating system to install on systems that are set up. Select a system image that has Image in the Image Type column.

   Specify Requirements page

   Indicates minimum system requirements and how to select systems.

   Specify Software page

   Select the software to install as part of the reservation.

   Specify Actions page

   (Optional) Select actions to run before provisioning a system, or after provisioning is completed or expired.

   Specify Access Policy page

   Select or remove the organizational units whose members are granted access to the template.

Create Group Templates

Administrators can bundle together individual reservation templates to create complete application configurations for systems. When using virtual resources, administrators can target specific virtual resource pools and determine whether all virtual machines are created in the same resource pool.

Note: Administrators can allow users to set the administrative password when creating a virtual machine (see Let Users Perform Some Administrative Tasks (see page 504)). When a reservation is made using a group template, users are not allowed to specify the password.

To create group templates

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your reservation templates.

   The Reservation Templates page opens. This page lists the inventory of public templates that are available to users for selection when reserving systems. From this page, you can create new public templates that you want to make available to users.

4. Select Create Group from the Actions menu to add new application server group templates to the inventory.

   The Create Application Group wizard opens.

5. Specify a name and description for the application server group template, then click Next.

   The Specify Processing Order page appears.

6. Specify the reservation templates that you want to group together as one template and reorder the templates as needed. If you want to create all virtual machines on the same VMware ESX server, select the check box “All virtual machines must be created in the same virtual resource pool”.

   Note: Reservation Manager determines the resource requirements for the full set of machines and only targets VMware ESX servers that have resources to accommodate creation of all the virtual machines. Click Next.

   The Specify Available Virtual Pools page appears.

7. Specify whether virtual machines are created in any available virtual resource pool or users are limited to a specific set of virtual resource pools. Click Next.

   Note: If you set the option to restrict all the virtual machines to a single resource pool, a page opens allowing you to specify any restrictions on the virtual resource pools for this template.

   The Specify Actions page appears.

8. Select the appropriate actions from the drop-down lists. You can specify CA Server Automation actions that are performed before or after the system imaging and software deployment operations are processed. You can also specify an action that is performed upon reservation end.

   Click Next.

   The Specify Access Policy page appears.

9. Select the organizational units whose members will be granted access to this public template, then click Finish.

   If an organizational unit that has been granted access to the template is not listed or you are unsure which organization units to grant access to this template, click Finish to skip this step. You can grant access to the template later.

   The Reservation Templates page opens and lists the template that has been successfully added to the inventory.

Make Virtual Machines Available to Users

Use the following process to make VMware virtual machines available for reservations:

1. Identify the virtual resource pool to which the VMs will be added when they are created.
2. Define the VM templates that a user can use to create VMs.
3. Set up access policy for both the resource pools and VM templates.

The following sections describe how to set up the Reservation Manager to support VM reservations.

Note: User permissions must be set up in the VMware product. See VMware vCenter and vSphere User Permissions (see page 396).

Prerequisites for Supporting Reservations of Virtual Machines

Before the Reservation Manager deploys virtual machines for usage, identify one or more VMware ESX servers or clusters of VMware ESX servers for use. Then, define which resource pools on each ESX server or cluster are targets for virtual machine creation. To determine whether an ESX server or cluster can create a virtual machine, the Reservation Manager calculates the amount of memory available for new virtual machines on the ESX servers.

If the resource pools on the ESX server or cluster have memory limits defined, the Reservation Manager requires exclusive access to the resource pool that is targeted to determine resource availability for the future.

In the absence of memory limits at the resource pool level, the Reservation Manager bases its calculations entirely on the amount of memory available to virtual machines at the ESX server level. The Reservation Manager requires exclusive use of the ESX Servers used for virtual machine creation to determine resource availability accurately in the future.

Note: The addition of resource pools and VM templates is performed in context to the vSphere Datacenter and the folder in which they are defined at the time they are added. If the VMware datacenter or folder is renamed or the VM templates are moved to a different folder, the resource pools and templates that were previously added to the Reservation Manager will no longer be usable. Thus, it is important to stabilize the vSphere structure before adding these items.

Create a Resource Pool for VMware Virtual Machines

The following procedure describes how to use a VMware vCenter or vSphere resource pool to create a virtual resource pool for the Reservation Manager.

Note: The addition of resource pools is performed in context to the vSphere Datacenter at the time they are added. If the VMware datacenter is renamed, the resource pools defined in the Reservation Manager will no longer be usable. Thus, it is important to stabilize the vSphere structure before adding resource pools.

To create a virtual resource pool for VMware

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Select Add Virtual Pool from the Actions menu in the upper right corner of the Resource Pools list.

   The Add Virtual Pool wizard opens to the Specify Pool page.

5. Follow the instructions in the wizard. The following pages are provided:

Specify Pool page

Specifies a name and optional description.

Note: Resource pool names are read-only after the pools are saved.

Specify Virtual Resource Pools page

Specifies the SCVMM server and the virtual resource pools to associate with the Reservation Manager pool. If storage tiers are enabled in the configuration settings for virtual machines, for each data store select an existing tier or define a new tier.

Configure Settings page

Specifies limits on the duration and number of logical partitions a user can request when submitting a reservation. The following fields require further explanation:

ITCM Domain Manager

If CA Server Automation is integrated with CA IT Client Manager (CA ITCM) for software delivery, enter its domain manager.

Scalability server

Select the name of the server used for deploying software to logical partitions associated with this pool.

Specify Access Policy page

Select or remove the organizational units whose members are granted access to the systems in the resource pool.

Note: You can grant access to the resource pool later by editing it.

Manage Snapshots in Resource Pools

Administrators can give users permission to take snapshots of virtual machines. They can also specify how many snapshots are allowed, and indicate whether to quiesce the file system (an explanation follows).

To manage snapshots in resource pools

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Click a resource pool.

   The Resource Pool Details pane opens to the Properties tab.

5. On the right side of the tab, perform the following actions:

   - Select **Allow users to take a VM snapshot**.
   - Select a number for **Maximum number of snapshots**.
   - (Optional) Select **Quiesce file system (Requires VMware tools installed)**.

     This option lets VMware Tools quiesce the file system in the virtual machine that is powered on when the snapshot is taken. Quiescing brings the on-disk data of a computer into a state that is suitable for backups. This process may include such operations as flushing dirty buffers from the operating system in-memory cache to disk, or other high-level application-specific tasks.

The snapshot permission takes effect for new reservations using this resource pool.

Stop VMs from being Provisioned from Resource Pools

Administrators can stop VMware virtual machines from being provisioned from specific resource pools.

To stop provisioning VMs from specific resource pools

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Double-click a resource pool.

   The Resource Pool Details pane opens.

5. Click the Resource Pool Details tab, and select from the drop-down list in the Operating Status column.

   Operating Status

   Sets the status to *In Service* or *Not Available*.

   Default: In Service

   When set to Not Available, the pool is not considered when checking for resource availability during creation of a VM. At provisioning time, if a reservation is allocated from this pool, the scheduler tries to move it to another pool. If no other pool can accommodate the request, the reservation fails.

Provision Storage for a VMware Resource Pool

Administrators can provision and attach new datastores to VMware hosts that have been added to VMware resource pools within Reservation Manager.

The Storage Provisioning Manager for NetApp must be configured before this feature can be used. See the section on Storage Provisioning Manager for NetApp.

To provision and attach storage to VMware hosts in a VMware Resource Pool

1. Log in to Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Select a VMware resource pool.

   The Resource Pool Details pane opens.

5. Click the Resource Pool Details tab.

   A page with a list of Datacenters opens.

6. Select a Datacenter.

   The Datacenter is highlighted.

7. Click the Actions drop-down and select Edit...

   The Edit Datastores page appears.

8. Click the Actions drop-down and select Provision Datastore...

9. Select the ESX Server, Enhanced Storage Policy from the drop-downs. Specify the amount of storage desired, and the name to attach to the datastore.

10. Press OK

    The datastore is added to the resource pool where the storage provisioning job was started.

To track the status of storage provisioning jobs, see the next section.

Provision Jobs for VMware Resource Pools

Administrators can view VMware Storage Provisioning jobs from within the Reservation Manager.

To view VMware storage provisioning jobs

1. Log in to Reservation Manager using the CA Server Automation administrator user credentials.
2. Click Administer Your Reservation Manager.
3. Click Manage your resource pools.

   The Resource Pools page appears.
4. Click the Actions drop-down and select View Storage Jobs...

   A page listing the storage jobs started from Reservation Manager appears.

Create a Template for VMware Virtual Machines

Administrators create one or more templates so that users can make reservations.

Note: The addition of VM templates is performed in context to the vSphere Datacenter and the folder in which they are defined at the time they are added. If the VMware datacenter or folder is renamed or the VM templates are moved to a different folder, the templates defined in the Reservation Manager will no longer be usable. Thus, it is important to stabilize the vSphere structure before adding templates.

To create a template for VMware virtual machines

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.
2. Click Administer Your Reservation Manager.

   The Administration page opens.
3. Click Manage your reservation templates.

   The Resource Pools page opens.
4. Select Create from the Actions menu in the upper right corner.

   The Create Reservation Template wizard opens.

5. Follow the instructions in the wizard. The following pages are provided:

Specify Name page

Indicates a name and optional description.

Specify System Image page

Indicates the operating system to install on systems that are set up.

Specify Custom Specification page

Lists the custom specification file to use when deploying a virtual machine using this template. Use the selected file or choose a different file.

Specify Requirements page

Indicates minimum system requirements. Enter the number of virtual machines to create, and enter a name or prefix for them. If you enter a prefix, the system appends a numeric suffix to create a unique name. Substitution variables are available instead of or in combination with a name or prefix. See Substitution Variables for Templates (see page 538).

Note: When you create VM names using either a prefix or a specified name, the name must contain alphanumeric characters (a - z, A - Z, 0 - 9) or a hyphen (-). If you use any other characters, they are automatically replaced with a hyphen.

Specify Software page

Select the software to install as part of the reservation.

Specify Actions page

(Optional) Select actions to run before provisioning a system, or after provisioning is completed or expired.

Specify Access Policy page

Select or remove the organizational units whose members are granted access to the template.

Specify Custom Specification page

Lists the custom specification file to use when deploying a virtual machine using this template. Use the selected file or choose a different file.

Substitution Variables for Templates

You can use any of the following substitution parameters instead of or in combination with a name or prefix for a Hyper-V or VMware virtual machine. There is no limit to the number of characters when you enter substitution parameters. These parameters are used on the Specify Requirements page of the Create Reservation Template wizard.

%DATACENTER%

Identifies the data center name on which the VM is created.

%HOSTSYSTEM%

Identifies the server name where the VM is created.

%ORGUNIT%

Identifies the organizational unit name of the user submitting the reservation.

%PROJECTID%

Identifies the project ID that the user enters when submitting a reservation.

%RESERVATIONID%

Identifies the numeric reservation ID assigned when the user submits a reservation.

%RESOURCEPOOL%

Identifies the Reservation Manager resource pool name.

%USERNAME%

Identifies the name of the user who logged in to Reservation Manager.

When using a prefix, a numeric suffix is appended to the name to help ensure a unique virtual machine name in Reservation Manager. You can use other methods to generate unique names, for example, by using the reservation ID in combination with the user name.

Example:

%USERNAME%-%RESERVATIONID% userkey01-62

%HOSTSYSTEM%-ServerA ESX1-ServerA

Use specified name

Specify a name.

Reserve a Virtual Machine

To verify that the Reservation Manager is configured to reserve virtual systems from the resource pool that was set up, click the Home breadcrumb to return to the home page.

You can create a reservation request for a virtual machine. Specify requirements such as the virtual machine template to use, software, and the number of machines to create. Reservation Manager uses these settings to create the virtual machine and reserve and provision it to fulfill your request.

The only supported method for virtual machine creation is deploying through templates. Select the template that you want to use for the virtual machine. The ensuing options are limited to virtual machines compatible with the selected template.

To reserve a virtual machine

1. Click Create a virtual machine from the home page of the Reservation Manager user interface.

   The Create Virtual Machine wizard appears with the Specify VM Template page displayed. This page displays a list of available virtual machine templates. The list is filtered to display only VMs that your organizational unit is authorized to use.

2. Select the virtual machine template that you want to use and click Next. The template determines the options presented throughout the rest of the wizard.

   Note: You can filter the table on this page by template name, description, operating environment, or location.

   The Specify Software page appears.

3. Move all software that you want installed on the virtual machine from the Available Software Group to the Selected Software Group.

   The Available Software Group list contains software that the administrator has identified as being compatible with the selected template. The list is filtered to display only software that you are authorized to access. Authorization is based on the organizational unit that your user name has been assigned to.

   The selected software is installed on the virtual machine in the order it is listed in the Selected Software Group list. Use the up and down arrow buttons to modify the order of software in the list. You do not have to select any software on this page.

   Click Next when finished.

   The Specify Requirements page appears.

4. Specify the number of virtual machines that you want to reserve, the required CPUs and required memory in the provided fields.

   (VMware only) If you require more disks, click Add Disk, select the disk size, then click Next.

   Note: The virtual machine is created with the additional disks, but the disks are not formatted automatically. Use Windows disk management to format them.

   The Specify Dates page appears.

5. Choose to reserve now for several days, or enter start and end dates and times. Both options also let you reserve a system with no expiration date and time.

   The times start and end on the hour. The end time must be at least one hour after the start time. If you specify an end time that exceeds the number of days for which you can reserve a virtual machine, the request is rejected. Limits on virtual machine reservation time are defined in policy at the resource pool level.

   Click Next when finished.

   Reservation Manager uses all defined information to determine whether the request can be fulfilled. One of the following actions occurs:

   - If the resources are available to fulfill the request, the Complete Request page appears.
   - If the requested resources are not available, the request is rejected. Modify the request to fit within the available resources.

6. Enter information in the following optional fields as necessary:

   Project ID

   (Optional) Defines the ID that associates the request with a specific project. This field is useful for evaluating requests and for allocating or charging costs to a specific project.

   The Reservation Manager administrator can configure this field so that it is required. In this case, the Reservation Manager displays a corresponding note in the information bar. You can submit your reservation request only after you enter a project ID. The Reservation Manager does not validate the entered value.

   Notes

   (Optional) Defines additional information you want to provide about the request.

   Send Notifications to

   (Optional) Defines the email address to which Reservation Manager sends notifications regarding this request. You can override any default value with a different address. Reservation Manager sends email notifications when a reserved system is available for use and when the expiration of the reservation period is approaching.

Save as Template

(Optional) Defines a name for this request so that you can use it as a template for future requests.

Template Description

(Optional) Defines a description for the template.

Click Finish.

The request is submitted for approval, and a final check is performed to verify that the machines are still available. If the resources are available, one of the following actions occurs:

- If Reservation Manager is configured for automatic request approval, the reservation immediately reflects that it has been approved.
- If Reservation Manager is configured for manual request approval, an email notification is sent to the administrator to approve the request.

7. Verify your reservation status on the View reservation details page or on the View all reservations page that is only available for administrators. The Job Status values in the table indicate the progress of the provisioning job:

Not Started

Indicates that the job status remains in this state until the scheduled reservation start date.

Initiated

Indicates that the reservation start date has been detected and processing of the reservation request has been initiated.

In Progress

Indicates that the provisioning of the first system has begun.

Ready

Indicates that all systems requested have been set up and are available for use.

Failed

Indicates that there was a failure processing one or more tasks associated with preparing the requested systems.

Canceled

Indicates that the user or administrator requested cancellation of the reservation.

Expired

Indicates that the reservation end date has been reached.

8. (Optional) Select the reservation item in the table and select Details from the Actions menu.

   Reservation Manager displays the details of this request.

   The following Job Status values in the administrator Details page appear as hyperlinks:

   - Not Started
   - Initiated
   - In Progress
   - Ready
   - Failed

9. (Optional) Click a link to go to a subordinate page that provides further details about the job status.

Specify a Prefix for Virtual Machine Names

You can specify a prefix for virtual machines at the resource pool level. A prefix provides a consistent naming convention. This procedure applies to AppLogic, Hyper-V and VMware virtual machines.

To specify a prefix for a virtual machine name

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens, and lists existing pools.

4. Double-click a resource pool.

   The Resource Pool details page opens, with the Properties tab displayed.

5. Select values for the following fields.

   VM Name Generation

   Lets you choose whether to use a prefix or a specified name. Reservation Manager appends the reservation ID to the name or prefix.

   Automatic using prefix

   Lets you enter a prefix for the VM name.

   Use specified name

   Lets you specify a name in the Name/Prefix field.

   Default: Automatic using prefix

   Name/Prefix

   Specifies a name or prefix for VM names.

   Default: none

   Limits: 10 alphanumeric characters (when Automatic using prefix is selected); 15 alphanumeric characters (when Use specified name is selected)

   Note: When you create VM names using either a prefix or a specified name, the name must contain alphanumeric characters (a - z, A - Z, 0 - 9) or a hyphen (-). If you use any other characters, they are automatically replaced with a hyphen.

6. Click OK.

   VM names will have a specified prefix or not.

Make Hyper-V Virtual Machines Available to Users

Use the following process to make Microsoft Hyper-V virtual machines available for reservations:

1. Identify the virtual resource pool to which the VMs will be added when they are created.
2. Define the VM templates that a user can use to create VMs.
3. Set up access policy for both the resource pools and VM templates.

Note: The following procedure for VMware virtual machines also applies to Hyper-V VMs. Other VMware procedures are similar.

Specify a Prefix for Virtual Machine Names (see page 542)

The following sections describe how to set up the Reservation Manager to support Hyper-V reservations.

Create a Resource Pool for Hyper-V Virtual Machines

The following procedure describes how to use a Hyper-V resource pool to create a virtual resource pool for the Reservation Manager.

To create a virtual resource pool for Hyper-V virtual machines

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Select Add Hyper-V Pool from the Actions menu in the upper right corner of the Resource Pools list.

   The Add Hyper-V Pool wizard opens to the Specify Pool page.

5. Follow the instructions in the wizard. The following pages are provided:

   Specify Pool page

   Specifies a name and optional description.

   Note: Resource pool names are read-only after the pools are saved.

   Specify Virtual Resource Pools page

   Specifies the SCVMM server and the virtual resource pools to associate with the Reservation Manager pool. If storage tiers are enabled in the configuration settings for virtual machines, for each data store select an existing tier or define a new tier.

Configure Settings page

Specifies limits on the duration and number of logical partitions a user can request when submitting a reservation. The following fields require further explanation:

ITCM Domain Manager

If CA Server Automation is integrated with CA IT Client Manager (CA ITCM) for software delivery, enter its domain manager.

Scalability server

Select the name of the server used for deploying software to logical partitions associated with this pool.

Specify Access Policy page

Select or remove the organizational units whose members are granted access to the systems in the resource pool.

Note: You can grant access to the resource pool later by editing it.

Create a Template for Hyper-V Virtual Machines

Administrators create one or more templates so that users can make reservations.

To create a template for Hyper-V virtual machines

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your reservation templates.

   The Resource Pools page opens.

4. Select Create from the Actions menu in the upper right corner.

   The Create Reservation Template wizard opens.

5. Follow the instructions in the wizard. The following pages are provided:

   Specify Name page

   Indicates a name and optional description.

   Specify System Image page

   Indicates the operating system to install on systems that are set up. Select a Hyper-V system image.

   Hyper-V Profiles

   Lists profiles for hardware and operating system. Select the profiles to use when deploying a virtual machine using this template.

   Specify Requirements page

   Indicates minimum system requirements. Enter the number of virtual machines to create, and enter a name or prefix for them. If you enter a prefix, the system appends a numeric suffix to create a unique name. Substitution variables are available instead of or in combination with a name or prefix. See Substitution Variables for Templates (see page 538).

   Note: When you create VM names using either a prefix or a specified name, the name must contain alphanumeric characters (a - z, A - Z, 0 - 9) or a hyphen (-). If you use any other characters, they are automatically replaced with a hyphen.

   Specify Software page

   Select the software to install as part of the reservation.

   Specify Actions page

   (Optional) Select actions to run before provisioning a system, or after provisioning is completed or expired.

   Specify Access Policy page

   Select or remove the organizational units whose members are granted access to the template.

Make Amazon Machine Images Available to Users

Reservation Manager lets you reserve Amazon Machine Images (AMIs) from the Amazon Elastic Compute Cloud (EC2). Use the following process to make AMIs available for reservation from EC2:

1. Add AMIs to the inventory.
2. Add an EC2 resource pool.

The following sections describe how to set up the Reservation Manager to support AMI reservations.

Add Amazon Machine Images to Inventory

Administrators can add AMIs to the system image inventory available for installation when users reserve AMI instances.

To add Amazon Machine Images to inventory

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your system image inventory.

   The System Images page opens. This page lists the inventory of AMI system images that are available to users for selection when reserving systems.

4. Select Add AMI Image from the Actions drop-down menu.

   The Add AMI Image wizard opens.

5. Select an operating environment for the AMI or select Unknown to view unclassified images. Select an image from the filtered list that is based on the operating environment you selected. Select an operating environment from Update Platform to classify an unknown image. Select the default instance type to use when launching instances of this AMI. Select the User Override Mode check box if you want to let users change the default instance type when submitting reservations. Enter any user data that you want to make available when this AMI is launched, and then click Next.

   The Specify Security Group page appears.

6. Select the security groups used to configure this AMI image when this instance is launched and move them to the Selected Security Groups area. Select the User Override Mode check box if you want to let users change the default security settings when submitting reservations, and then click Next.

   The Specify Access Policy page appears.

7. Select the organizational units whose members are allowed to install this image and move them to the Selected Organizational Units area, and then click Finish.

   The System Images page opens and lists the AMI that has been successfully added to the inventory.

Create a Resource Pool for Amazon EC2

The following procedure describes how to create an Amazon EC2 resource pool for use by the Reservation Manager.

To create an Amazon EC2 resource pool

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Click Actions and select Add EC2 Pool from the drop-down menu in the upper right corner of the Resource Pools list.

   The Create Resource Pool wizard opens the Specify Pool page.

5. Enter the pool name and an optional description for the Reservation Manager resource, and then click Next.

   The Specify Instance Parameters page opens.

6. Select information for the following fields, and then click Next.

   Key Pair Name

   Defines the key pair name used to log in to or access the AMI instance. Select one from the drop-down list.

   Note: Provide the private key file associated with the selected keypair to users who are reserving Linux AMI instances. After the instance is running, users must specify the path to the private key file to open an SSH connection to the running instance.

   Network Selection area:

   Public

   Indicates that the area in the cloud is a public area.

   Note: If you choose Public, you can select an availability zone, although Amazon recommends letting the availability zone be selected for you.

   Private

   Indicates that the connection to Amazon is through a Virtual Private Cloud (VPC). If you choose Private, select a Subnet.

   Availability Zone, Subnet IP Address, and Domain Name appear.

   The Configure Settings page opens.

7. Specify limits on the duration and number of Amazon EC2 instances a user can request when submitting a reservation for resources:

   Maximum Days

   Limits reservation length to the specified number of days.

   Maximum Systems

   Limits how many instances a user can reserve out of this pool at any single point in time.

   Click Next.

   The Specify Access Policy page appears.

8. Select the organizational units whose members will be able to launch instances in the context of the selected EC2 pool, and then click Finish.

   Note: Before you grant access to users, grant access to administrators so that you can test creating reservations for systems in this pool.

   If an organizational unit that has been granted access to the resource pool is not listed or you are unsure which organizational units to grant access, click Finish to skip this step. You can grant access to the resource pool later.

   The Reservation Manager creates a resource pool and displays a confirmation message when the process completes successfully. The wizard closes and the Create Resource Pool page displays the new Amazon EC2 resource pool in the Resource Pools list.

Create a Template for Amazon Machine Images

Administrators create one or more templates so that users can make reservations in the Amazon Elastic Compute Cloud (EC2).

To create a template for Amazon Machine Images

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your reservation templates.

   The Reservation Templates page opens.

4. Select Create from the Actions menu in the upper right corner.

   The Create Reservation Template wizard opens.

5. Follow the instructions in the wizard. The following pages are provided:

   Specify Name page

   Indicates a name and optional description.

   Specify System Image page

   Indicates the operating system to install on systems that are set up. Select an Amazon system image.

   Specify Requirements page

   Indicates the system requirements. Enter the number of virtual machines to create. (Optional) Select security groups, and select a different instance type (tier of service).

   Specify Actions page

   (Optional) Select actions to run before provisioning a system, or after provisioning is completed or expired.

   Specify Access Policy page

   Select or remove the organizational units whose members are granted access to the template.

Make AppLogic Applications Available to Users

Use the following process to make AppLogic Applications available for reservations:

1. Identify the virtual resource pool to add applications when they are created.
2. Define the application templates that a user can use to create AppLogic applications.
3. Set up access policy for both the resource pools and application templates.

The following sections describe how to set up the Reservation Manager to support AppLogic application reservations.

More information:

Specify a Prefix for Virtual Machine Names (see page 542)

Create a Resource Pool for AppLogic Applications

The following procedure describes how to use an AppLogic resource pool to create a virtual resource pool for the Reservation Manager.

To create a virtual resource pool for AppLogic Applications

1. Log in to the Reservation Manager using CA Server Automation administrator user credentials.
2. Click Administer Your Reservation Manager and Manage your resource pools.
3. Select Add AppLogic Pool from the Actions menu in the Resource Pools list.

   The Add AppLogic Pool wizard opens.
4. Follow the instructions on the following pages:

   Specify Pool

   Specifies a name and optional description.

   Note: Resource pool names are read-only after the pools are saved.

   Configure Settings

   Specifies limits on the duration and number of applications that a user can request when submitting a reservation. A default application name is provided in the resource pool.

   Specify Access Policy

   Selects or removes the organizational units whose members are granted access to the systems in the resource pool.

Note: You can grant access to the resource pool later by editing it.

Create a Template for AppLogic Applcations

AppLogic application reservations are created only through reservation templates. AppLogic configuration parameters provide information to the AppLogic grid so that the application can be set up and configured properly. IP addresses, system name, network name server addresses, application monitoring settings are examples of application configuration parameters.

Note: You do not need to create system images for AppLogic applications because they are created automatically when templates are created for AppLogic applications.

Set Up AppLogic Application Configuration Parameters

Reservation manager groups configuration parameters into two major categories:

Mandatory and Optional Parameters

- Mandatory parameters are required for the application to start up and run correctly. These parameters are marked appropriately in reservation manager and must have a value entered when creating the application or the application template.
- Optional parameters can be specified to provide more customization to the application but are not required for application operation.

Auto fill and User-Filled Parameters

- Reservation manager can extract information from the grid and automatically fill in mandatory parameters with correct values. For example, required parameters like IP address, DNS server address, Netmask parameters are auto-filled on behalf of the administrator. Auto filled parameters are not allowed to be overridden and are disabled in the reservation template wizard. The template becomes auto filled based on configuration values imported by reservation manager from the grid controller. Following are the configuration parameter names:
- **primary_ip**
  - in_ip
  - out_ip
  - usr_ip
  - mail_ip
  - netmask
  - gateway
  - dns1
  - dns2
  - dns3

To create a template for AppLogic Applications

1. Log in to the Reservation Manager using the CA Virtual Automation administrator user credentials.
2. Click Administer Your Reservation Manager and Manage your reservation templates.
3. Select Create from the Actions menu.

   The Create Reservation Template wizard opens. Follow the instructions in the wizard:

   Specify Name

   Indicates a name and optional description.

   Specify AppLogic Template

   Selects the Grid and Application template used when creating a reservation. The administrator determines the default CPU, Memory, and Bandwidth to provision with the application.

   Specify Requirements

   This page allows the administrator to set up the configuration parameters for the application. Default values can be specified (if appropriate). The administrator grants permissions to the end user to change default values by selecting appropriate checkboxes.

   Specify Actions

   (Optional) Selects actions to run before or after the system imaging operation is processed. You can also specify an end of reservation action.

   Specify Access Policy

   Selects or removes the access to the template for members of organizational units.

Configure Parameters for Email Notification

The administrator notifies the reservation requestor with the configured parameter values upon successful deployment of an application. The administrator can select the parameters for email notification when the application template is created. Select *Include in Email Notification* when the reservation template is created.

Logical Partitions

IBM PowerVM lets administrators create logical partitions that users can reserve. Administrators can also create and edit resource pools and public templates for logical partitions.

Logical partitions are small segments of a larger system. They have their own operating system and applications, and are independent from each other.

More information:

Create a Resource Pool for IBM PowerVM Logical Partitions (see page 554)
Edit a Resource Pool for IBM PowerVM Logical Partitions (see page 555)
Create a Template for IBM PowerVM Logical Partitions (see page 557)
Static IP Addresses for IBM PowerVM Logical Partions (see page 559)
Define Network Address Pools (see page 579)
Configure Chargeback by Tier for IBM PowerVM Logical Partitions (see page 566)
Select a Chargeback Tier for IBM PowerVM Logical Partitions (see page 567)

Create a Resource Pool for IBM PowerVM Logical Partitions

The following procedure describes how to create a resource pool for IBM PowerVM logical partitions in the Reservation Manager.

To create a resource pool for logical partitions

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Select Add IBM PowerVM Pool from the Actions menu in the upper right corner of the Resource Pools list.

   A wizard for creating a resource pool opens to the Specify Pool page. Enter information about the following pages as instructed by the wizard.

   Specify Pool page

   Specifies a name and optional description.

   Note: Resource pool names are read-only after the pools are saved.

Specify Logical Partition Resources page

Select the IBM PowerVM server (HMC/IVM) to associate with this resource pool.

You can then add a managed system and storage by selecting Add from the Actions menu. Note that tiers may not be enabled at all sites.

Configure Settings page

Specifies limits on the duration and number of logical partitions a user can request when submitting a reservation. The following fields require further explanation:

ITCM Domain Manager

If CA Server Automation is integrated with CA IT Client Manager (CA ITCM) for software delivery, enter its domain manager.

Scalability server

Select the name of the server used for deploying software to logical partitions associated with this pool.

Specify Access Policy page

Select or remove the organizational units whose members are granted access to the systems in the resource pool.

Note: You can grant access to the resource pool later by editing it.

The Reservation Manager creates a resource pool.

Edit a Resource Pool for IBM PowerVM Logical Partitions

The following procedure describes what you can edit in a resource pool for IBM PowerVM logical partitions in the Reservation Manager.

Follow these steps:

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens, showing existing resource pools in a table.

4. Select the pool that you want to edit, and then select Details from the Actions menu.

   The Resource Pool Details pane opens to the Properties tab. Enter or update information as described for the following tabs.

Properties tab

The Properties tab lets you update such things as maximum days users can reserve a system and the maximum systems user can reserve, whether to approve reservations automatically, and whether management agents are installed. The following fields require explanation:

Maximum Systems

Defines the maximum number of systems that a user can simultaneously hold reservations.

Default: 10

Maximum Days

Defines the maximum number of days a user can reserve a system in this pool when submitting a single request.

Default: 30

ITCM Domain Manager

If CA Server Automation is integrated with CA IT Client Manager (CA ITCM) for software delivery, enter its domain manager.

Scalability server

Select the name of the server used for deploying software to logical partitions associated with this pool.

Automatically approve reservation requests

Select Automatically approve reservation requests to enable users to self-approve the reservation requests without contacting administrator.

Allow users to manage logical partition power state

Select Automatically approve reservation requests to enable users to perform the power options.

Resource Pool Details tab

Add, edit, or delete the IBM PowerVM server (HMC/IVM), managed system, VIO server, and storage associated with this resource pool. Use the Actions menu.

Note: Tiers may not be enabled at all sites.

Access Policy tab

Select or remove the organizational units whose members are granted access to the systems in the resource pool.

Let Users Manage the Power Status of an IBM PowerVM Logical Partition

Administrators can give users permission to perform some administrative tasks on IBM logical partitions. Users can manage the power status of the logical partitions without contacting the administrator.

Follow these steps:

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens, and lists existing pools.

4. Double-click a resource pool.

   The Resource Pool details page opens, with the Properties tab displayed.

5. Select or clear the following fields:

   Allow users to manage logical partitions power state

   Lets users turn on, turn off, and reset on IBM logical partitions that are assigned to them.

6. Click OK.

   The permissions are granted.

Create a Template for IBM PowerVM Logical Partitions

Administrators can create public templates that define commonly used system configurations. You can define public templates for new IBM PowerVM logical partitions.

To create a public template for end users of IBM PowerVM logical partitions

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your reservation templates.

   The Reservation Templates page opens. This page lists the public templates that users can select when reserving systems. From this page, you can create new public templates.

4. Select Create from the Actions menu.

   The Create Reservation Template wizard opens. Enter information about the following pages as instructed by the wizard.

   Specify Name page

   Specify a name and optional description. Select the operating system AIX.

   Specify System Image page

   Select the Network Installation Manager (NIM) resource group to use with the reservation.

   Specify Requirements page

   Click Create new Logical Partition and indicate how many to create. Select other values as needed, or use the defaults.

   Specify Software page

   Select the software to install as part of the reservation.

   Specify Actions page

   (Optional) Select actions to run before provisioning a system, or after provisioning is completed or expired.

   Specify Access Policy page

   Select or remove the organizational units whose members are granted access to the template.

Static IP Addresses for IBM PowerVM Logical Partions

IBM PowerVM logical partitions require static IP addresses. These IP addresses must be DNS resolvable and correspond to a NIM Machine Resource configuration on the NIM Master. For this release, with the introduction of Dynamic NIM Machine resources, you can do the following:

- Create the NIM machine resource configurations as part of the reservation provisioning process automatically
- Remove on reservation expiration or cancelation.

Administrator does not need to precreate the configurations on the NIM Master.

Note: The pre-existing NIM Machine resource configurations are still used and will not be removed on reservation expiration or cancelation.

Important! Network address pools must be unique to IBM PowerVM. Do not add IP address ranges for virtual machines in the same network address pool as IBM PowerVM.

When a user creates a reservation, an IP address is assigned to it from a range of IP addresses the administrator defined. When a reservation is fulfilled or canceled, the IP address is released and made available for another reservation.

More information:

Define Network Address Pools (see page 579)

Approve or Reject Reservation Requests

If you have configured the Reservation Manager not to approve reservation requests automatically, you will receive an email notification from the Reservation Manager after a user has submitted a new request. The status of those requests is Pending Approval and you must approve or reject the request manually.

If you have configured the Reservation Manager to approve requests automatically, the Reservation Manager does not send an email and the reservation will be approved immediately if the requested resource is available.

To approve or reject reservation requests manually

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer your Reservation Manager.

   The Administration page opens.

3. Click View all reservations.

   The View Reservations page appears.

4. Select a request showing Pending Approval status, and select Details from the Actions menu.

   The Reservation Details page appears. If CA Service Desk Manager integration is set up, a hyperlink to the help desk ticket is included on this page. You can approve or reject the help desk ticket using either CA Service Desk Manager or from the Reservation Manager Actions menu.

5. Review the reservation request, click Actions and approve or reject the request.

   Reservation Manager returns to the View Reservations page and displays the updated status.

More information:

Specify When to Send Pending Approval Request Notification (see page 470)
Set Automatic Cancellation of Unapproved Reservations (see page 469)

Chargeback

The chargeback feature provides a way to charge hourly rates for the use of virtual machines. CA Server Automation supports two distinct methods of pricing. The method of chargeback for physical systems and virtual machines allows setting charges for individual CPU, memory, and disk resources reserved. The method of chargeback for Amazon EC2 instances, IBM PowerVM logical partitions, and virtual machines provisioned using AppLogic uses a tier-based pricing. Each supported platform has its own pricing model based on one of these methods. The default for all chargeback models is 0 (no hourly charge); the administrator can change the value. Reports that show usage charges are also available.

Note: Chargeback records are created at the end of the day after midnight. The option Within Past 24 Hours gets costs from the previous day.

When you make a reservation, the costs are shown. If you change reservation requirements, you can recalculate the amount before submitting the reservation request.

Note: You can control whether chargeback is used, how often costs are calculated per day, how many days to retain calculation data, and the currency to use. See Configure Chargeback (see page 471).

Chargeback by Resource

The chargeback policy for physical systems and virtual machines (Hyper-V and VMware) is based on resources that are reserved. It can be a flat hourly rate, but surcharges also can be applied for the following situations in any combination:

- Number of CPUs or CPUs over a base threshold
- Each GB of memory or GBs of memory over a threshold
- Each GB of disk space or GBs of disk space over a threshold

An example of chargeback policy is an hourly rate of 25 cents, with surcharges for using more than one CPU, 2 GB of memory, and 10 GB of space.

Note: When storage tiers are allowed, surcharges for disk space cannot be used. Instead, an hourly rate is assigned to each tier. Therefore, chargeback is by storage tiers or surcharge, but not both. For more information about tiers, see Allow Users to Select Storage Tiers (see page 497).

Chargeback by Tier for Amazon EC2

For Amazon EC2, the chargeback policy is a flat hourly rate with no surcharges available and is based on the following instance types:

- c1.medium
- c1.xlarge
- cc1.4xlarge
- m1.small
- m1.large
- m1.xlarge
- m2.xlarge
- m2.2xlarge
- m2.4xlarge
- t1.micro

Chargeback by Tier for IBM PowerVM Logical Partitions

For IBM PowerVM logical partitions, the chargeback policy is a flat hourly rate with no surcharges available. CA Server Automation comes with the following tiers:

- lpar.Large
- lpar.Medium
- lpar.Small

Administrators can add more, if needed, when creating or editing templates.

Chargeback by Tier for AppLogic Applications

For AppLogic applications, the chargeback policy is a flat hourly rate with no surcharges available. CA Server Automation comes with the following tiers:

- applogic.Large
- applogic.Medium
- applogic.Small

More information:

Configure Chargeback by Resource (see page 563)
Configure Chargeback for Storage Tiers (see page 564)
Configure Chargeback by Tier for Amazon EC2 (see page 565)
Configure Chargeback by Tier for IBM PowerVM Logical Partitions (see page 566)
Select a Chargeback Tier for IBM PowerVM Logical Partitions (see page 567)

Configure Chargeback by Resource

Chargeback by resource is the method used for charging for reserved resources on physical systems and virtual machines (Hyper-V and VMware).

When assigning chargeback rates, keep in mind that the charge is for 24 hours a day and 7 days a week during the reservation period. Enter realistic hourly rates and surcharge amounts based on factors like energy costs and hardware maintenance expense. Most rates will probably range from 10 cents to 50 cents an hour.

Note: Chargeback is optional. The default rates are 0 (no chargeback).

To configure chargeback by resource

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your chargeback models, select a model, and then select Details from the Actions menu.

   A pane with fields for resource chargeback appears.

4. Enter values in the following fields:

   - **Base Hourly Rate** (example, $0.15)
   - **Base Memory (GB)** (example, 1)
   - **Base Disk (GB)** (example, 10)
   - **Base CPU** (example, 1)

   Note: When storage tiers are allowed, Base Disk (GB) and Additional Disk in the next step are not available. Hourly rates are entered in the storage tiers instead. See Configure Chargeback for Storage Tiers (see page 564).

5. (Optional) Enter the following surcharges in any combination. The surcharges are charged when the base thresholds entered in the previous step are exceeded. Use the checkboxes to activate the surcharges.
   - **Additional Memory** (example $0.01)
   - **Additional Disk** (example $0.01)
   - **Additional CPU** (example $0.05)

     Note: When storage tiers are allowed, surcharges are not available. Instead, an hourly rate is assigned to each tier. Therefore, chargeback is by storage tiers or surcharge, but not both. For more information about tiers, see Allow Users to Select Storage Tiers (see page 497).
6. Click OK.

Configure Chargeback for Storage Tiers

Chargeback for storage tiers lets you charge different rates for each tier on physical systems and virtual machines (Hyper-V and VMware).

Storage tiers are classifications for the data stores associated with each disk. Tiers generally indicate different levels of performance of the data store on which a VM and its hard drives are created. Administrators can enable or disable storage tiers.

When the administrator enables storage tiers, users can select them when they make virtual machine reservations. See Allow Users to Select Storage Tiers (see page 497).

When assigning chargeback rates, keep in mind that the charge is for 24 hours a day and 7 days a week during the reservation period. Enter realistic hourly rates and surcharge amounts based on factors like energy costs and hardware maintenance expense. Most rates will probably range from 10 cents to 50 cents an hour.

Note: Chargeback is optional. The default rates are 0 (no chargeback).

To configure chargeback for storage tiers

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.
2. Click Administer Your Reservation Manager.

   The Administration page opens.
3. Click Manage your chargeback models, select a model, and then select Details from the Actions menu.

   A pane with fields for resource chargeback appears.

4. Enter values in the following fields:
   - **Base Hourly Rate** (example, $0.15)
   - **Base Memory (GB)** (example, 1)
   - **Base CPU** (example, 1)
5. (Optional) Enter the following hourly surcharges in any combination. The surcharges are charged when the base thresholds entered in the previous step are exceeded. Use the checkboxes to activate the surcharges.
   - **Additional Memory** (example $0.01)
   - **Additional CPU** (example $0.05)
6. Enter hourly rates for disk space on storage tiers.
7. Click OK.

More information:

Allow Users to Select Storage Tiers (see page 497)

Configure Chargeback by Tier for Amazon EC2

Chargeback by tier is the method used for charging for reservations on Amazon cloud instances.

When assigning chargeback rates, keep in mind that the charge is for 24 hours a day and 7 days a week during the reservation period. Enter realistic hourly rates based on factors like energy costs and hardware maintenance expense. Most rates will probably range from 10 cents to 50 cents an hour.

Note: Chargeback is optional. The default rates are 0 (no chargeback).

To configure chargeback by tier for Amazon EC2

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.
2. Click Administer Your Reservation Manager.

   The Administration page opens.
3. Click Manage your chargeback models, and Amazon Cloud Instances.

   A pane with fields for tier chargeback appears.

4. Select the instance types for chargeback, and enter hourly rates. The instance types are the following:
   - c1.medium
   - c1.xlarge
   - m1.large
   - m1.small
   - m1.xlarge
   - m2.xlarge
   - m2.2xlarge
   - m2.4xlarge
   - t1.micro
5. Click OK.

Configure Chargeback by Tier for IBM PowerVM Logical Partitions

Chargeback by tier is the method used for charging for reservations on IBM PowerVM logical partitions.

When assigning chargeback rates, keep in mind that the charge is for 24 hours a day and 7 days a week during the reservation period. Enter realistic hourly rates based on factors like energy costs and hardware maintenance expense. Most rates will probably range from 10 cents to 50 cents an hour.

Note: Chargeback is optional. The default rates are 0 (no chargeback).

To configure chargeback by tier for IBM PowerVM logical partitions

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.
2. Click Administer Your Reservation Manager.

   The Administration page opens.
3. Click Manage your chargeback models, and IBM Logical Partitions.

   A pane with fields for tier chargeback appears.

4. Select the tiers for chargeback, and enter hourly rates. The default tiers are the following, although the administrator may have defined more when creating or editing templates:
   - lpar.Large
   - lpar.Medium
   - lpar.Small
5. Click OK.

Select a Chargeback Tier for IBM PowerVM Logical Partitions

If chargeback is in use at your site, you can open a template for IBM PowerVM logical partitions and select a chargeback tier.

To select a chargeback tier for IBM PowerVM logical partitions

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.
2. Click Administer Your Reservation Manager.

   The Administration page opens.
3. Click Manage your reservation templates.

   The Reservation Templates page opens. This page lists the public templates that users can select when reserving systems.
4. Double-click a template for logical partitions.

   The Reservation Templates Details page opens.
5. Click the Allocation Policy tab, and select a tier from the Chargeback Tier Name drop-down list.

   Note: You can also create a new chargeback tier by entering a new name. Be sure to set an hourly rate for any new tiers. See Configure Chargeback by Tier for IBM PowerVM Logical Partitions (see page 566).

Configure Chargeback by Tier for AppLogic Applications

Chargeback by tier is the method used for charging for reservations on AppLogic applications.

When assigning chargeback rates, keep in mind that the charge is for 24 hours a day and 7 days a week during the reservation period. Enter realistic hourly rates based on factors like energy costs and hardware maintenance expense. Most rates range from 10 cents to 50 cents an hour.

Note: Chargeback is optional. The default rates are 0 (no chargeback).

To configure chargeback by tier for AppLogic applications

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your chargeback models, and AppLogic applications.

   A pane with fields for tier chargeback appears.

4. Select the tiers for chargeback, and enter hourly rates. The default tiers are the following, although the administrator can define more when creating or editing templates:

   - applogic.Large
   - applogic.Medium
   - applogic.Small

5. Click OK.

Select a Chargeback Tier for AppLogic Applications

If chargeback is in use at your site, you can open a template for AppLogic applications and select a chargeback tier.

To select a chargeback tier for AppLogic applications

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your reservation templates.

   The Reservation Templates page opens. This page lists the public templates that users can select when reserving systems.

4. Double-click a template for AppLogic applications.

   The Reservation Templates Details page opens.

5. Click the Allocation Policy tab, and select a tier from the Chargeback Tier Name drop-down list.

   Note: You also can create a chargeback tier by entering a new name. Set an hourly rate for any new tiers. See Configure Chargeback by Tier for AppLogic Applications (see page 567).

Email Customization

Reservation Manager provides templates for HTML notification emails. These templates let administrators customize email text and use substitution variables that include dynamic information like server name or the current time.

Success emails are color-coded with blue bars so that people can see at a glance that their job ran. Failure emails contain orange bars.

More information:

Directory Structure (see page 569)
Templates and Email Types (see page 570)
Substitution Variables (see page 570)
Email Types and Categories (see page 573)
Conditional Substitution (see page 575)
Configure Parameters for Email Notification (see page 553)

Directory Structure

Templates are located in the directory %INSTALL_ROOT%/MailTemplates, followed by a two-character code for the language (based on the ISO-639-1 standard), and then the template file name. For example, the path for English-language templates is:

%INSTALL_ROOT%/MailTemplates/en/*template_name*.html

Templates and Email Types

Template file names are based on the email types. The email types refer to situations that generate emails automatically. Some examples are the following:

RESERVED_SYSTEM_READY,

RESERVED_SYSTEM_TERMINATION_WARNING,

RESERVATION_READY,

RESERVATION_ABOUT_TO_EXPIRE,

RESERVATION_EXPIRED,

The template that corresponds to the RESERVED_SYSTEM_READY email type is RESERVED_SYSTEM_READY.html.

Template files are written using XHTML. One CSS file is provided for all the templates, and it is embedded in the HTML email messages rather than having an external style sheet. Customizing the CSS file (styles.css) is efficient because you edit only a single file.

Note: Do not delete styles.css.

Important! If you edit a template file, save it with UTF-8 encoding. Operating systems that are not English and have multibyte characters must be saved with UTF-8 encoding. Windows Notepad can save with UTF-8 encoding.

Substitution Variables

A template file can contain substitution variables, like the following:

System Name: %SYSTEM_NAME%

The substitution variables are associated with a category of email message. The categories are:

- Reservation
- System
- Task
- Detailed message
- Miscellaneous

Reservation

The following variables are associated with reservation messages.

%IMAGENAMES%	Image names used for the reservation. (comma-separated string).
%NUMSYSTEMS%	Number of systems reserved.
%ORGUNIT%	Organizational unit of the requestor.
%PROJECTID%	Project ID associated with the reservation.
%READYSYSTEMLIST%	List of system names or IP addresses for reserved systems that are ready (HTML list).
%READYSYSTEMTABLE%	Table of system names or IP addresses for reserved systems that are ready (HTML table).
%REQUESTEDSOFTWARE%	List of requested software (HTML list).
%RESERVATIONENDTIME%	Time when the reservation ends.
%RESERVATIONID%	Reservation ID.
%RESERVATIONNOTES%	Notes supplied for the reservation.
%RESERVATIONREADYTEXT%	The user-supplied text to include with all reservation ready e-mails.
%RESERVATIONSTARTTIME%	Time when the reservation starts.
%TEMPLATENAME%	The template used for this reservation.
%TICKETID%	The ticket ID associated with this reservation.
%TICKETURL%	The ticket URL associated with this reservation.
%USEREMAILADDRESS%	E-mail address of the requestor.
%USERNAME%	Username of the reservation requestor.
%VMNAMES%	List of reserved systems (HTML list).

System

The following variables are associated with system messages.

%DATACENTER%	The name of the data center.
%HOSTSYSTEM%	The name of the VM host system.
%IMAGENAME%	The name of the system image used to create the VM.

%IPADDRESSES%	A list of IP addresses associated with the system (comma-separated string).
%RESOURCEPOOL%	The resource pool name.
%SERVER%	The name of the server that is being reserved for a user. If a new VM is being created, this name is the same as the VM name.
%SYSTEMPASSWORD%	The system password (only included if "ReservedSystemReadyNotificationContainsPassword" is true). For Amazon EC2 systems, the variable is always included.
%SYSTEMUPDATEDTIME%	The time that the status of the system was last updated.
%SYSTEMUSERNAME%	User name associated with %SYSTEMPASSWORD% (only included if "ReservedSystemReadyNotificationContainsPassword" is true).
%VMCONSOLEURL%	URL for the VM Console.
%VMNAME%	The name of the virtual machine.

Task

The following variables are associated with task messages.

%TASKID%	The task ID.
%TASKDESCRIPTION%	The task description.
%TASKTYPE%	The task type.
%TASKTYPESHORT%	A shortened version of the task type.

Detailed message

The following variable is associated with detailed messages.

%DETAILEDMESSAGE%	The detailed message.

Miscellaneous

The following variables are associated only with an *Approval Required* message.

%AUTOCANCEL%	If the reservation is not approved in time (true or false), this message indicates whether the reservation is canceled automatically.
%AUTOCANCELMESSAGE%	If %AUTOCANCEL% is true, this message explains that the reservation is canceled automatically. Otherwise, no value.
%APPROVALDEADLINE%	If %AUTOCANCEL% is true, this message shows the time at which the reservation is canceled if it has not been approved.

The following variable is associated with any message.

%CURRENTTIME%	The current time.

More information:

Conditional Substitution (see page 575)

Email Types and Categories

The following table shows email types and the email categories associated with them. The categories were listed earlier in the Substitution Variables section (see page 570). The categories are the following:

- Reservation
- System
- Task
- Detailed message

Email Type	Reser-vation	System	Task	Detailed Message
APPROVAL_REQUIRED	Yes	Yes*	No	No
HELPDESK_TICKET_OPENED_FOR_RESERVATION	Yes	Yes	No	No
NOT_ENOUGH_SPACE_FOR_SNAPSHOT	No	Yes	No	Yes
RESERVATION_ABOUT_TO_EXPIRE	Yes	Yes	No	No

RESERVATION_APPROVED	Yes	Yes	No	No
RESERVATION_CANCELED	Yes	Yes*	No	No
RESERVATION_EXPIRED	Yes	Yes	No	No
RESERVATION_NOT_APPROVED_IN_TIME	Yes	Yes	No	No
RESERVATION_PROCESSING_RESUMED	Yes	Yes	No	Yes
RESERVATION_PROCESSING_SUSPENDED	Yes	Yes	No	Yes
RESERVATION_READY	Yes	Yes*	No	No
RESERVATION_REJECTED	Yes	Yes	No	No
RESERVATION_TASK_FAILED	Yes	Yes	Yes	Yes
RESERVATION_TASK_TAKES_TOO_LONG	Yes	Yes	Yes	Yes
RESERVED_SYSTEM_READY	Yes	Yes	No	No
RESERVED_SYSTEM_TERMINATION_WARNING	Yes	Yes	No	No
SCHEDULER_PROCESSING_RESUMED	No	No	No	Yes
SCHEDULER_PROCESSING_SUSPENDED	No	No	No	Yes
TERMINATION_TASK_FAILED	Yes	Yes	Yes	Yes
TEXT_SUPPLIED	Yes	Yes	No	No
VM_POWER_OPERATION_FAILED	Yes	Yes	No	Yes

* Yes applies only to reservations for a single system.

Conditional Substitution

In some cases, substitution variables do not apply to all situations. For example, %VMCONSOLEURL% applies only to VMware-based systems. If Reservation Manager sends an e-mail about a Hyper-V system, the e-mail would list the field for VM Console URL but with a blank value. If emails contain blank fields, they can be confusing and unattractive.

You can eliminate blank fields by surrounding variables with ampersands (@) in the template, for example:

@%VMCONSOLEURL%@

The following example shows how to code it in the template:

```
<table border=1>
<td>System Name</td> <td>%SERVER%</td>
<td>IP Address</td> <td>%IPADDRESSES%</td>
<td>VM Console URL</td> <td>@%VMCONSOLEURL%@</td>
</table>
```

Note: The variable must be on the same line as the text associated with it. The following example would not eliminate the bank field:

```
<td>VM Console URL</td>
<td>@%VMCONSOLEURL%@</td>
```

Suspend and Restart the Scheduling of Tasks

The Reservation Manager provides a scheduler feature that lets administrators suspend the start of scheduled reservation tasks to prevent jobs from running during system maintenance. When maintenance is finished, administrators can resume scheduler operations to start processing suspended tasks.

When the scheduler is suspended, all OS provisioning, software deployment and termination processing tasks that have not yet started are suspended. Additionally, all pre-, post-, and expiration actions specified in reservation templates that are to be executed as part of the reservation setup or expiration processing workflow are suspended.

Reservation tasks that are in progress at the time the scheduler is suspended are not stopped. Reservation tasks that are scheduled to run after the in progress tasks complete are started when the scheduler is restarted. For example, if a virtual machine is deployed when the scheduler is suspended, any subsequent task such as a post-action or software deployment task is not started while the scheduler is suspended.

If the maintenance activities to be performed after suspending the scheduler include stopping the Reservation Manager or CA Server Automation services, we recommend that time be allowed for all reservation tasks that are in progress to be completed before the services are stopped. This will help ensure that the Reservation Manager is able to monitor the success or failure of the in progress tasks.

The following operations are not affected by suspension of the scheduler:

- User requests to change the power state of a virtual machine, take or revert snapshots, and reconfigure virtual machines
- New reservations are accepted, but are not released for provisioning while the scheduler is suspended

Notify users in advance of planned outages to give them time to extend reservations that are set to expire during the outage. They may not be able to extend reservations during the outage. We recommend that they extend their reservations beyond the duration of the planned outage. Reservations that expire while the scheduler is suspended are immediately processed when the scheduler is restarted.

To suspend and restart *all* provisioning tasks

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   Note: You may see a message that Reservation Manager is offline for maintenance. Click Administrator Login at the bottom of the display.

   The home page opens.

2. Click the *Administer Your Reservation Manager* link.

   The Administration page opens.

3. Click the Maintenance link at the top right of the display.

   A dialog opens and provides the following options. Select or deselect one or both.

   Suspend pending reservation tasks

   Allows users to log in and enter reservations. Reservation jobs are suspended as described in the introduction to this topic.

   Block web access (maintenance mode)

   Does not allow users to log in. An error message is displayed if users try to log in.

   When maintenance is complete, click OK.

   Note: If maintenance lasts for more than one day, repeat these steps until maintenance is complete.

Suspension and Restart of Individual Tasks

The Reservation Manager automatically suspends a provisioning task when it detects that the task would fail under current conditions. In these cases, the administrator receives an email notification that a problem exists and must be resolved. For example, Reservation Manager checks whether sufficient disk space is available before provisioning a virtual machine. Insufficient disk space causes automatic suspension of the deployment task. The administrator must resolve the disk space issue before restarting the task that was suspended.

To restart *individual* provisioning tasks that the system suspended

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click the Administer Your Reservation Manager link.

   The Administration page opens.

3. Click View all reservations.

   The View Reservations page opens. The Job Status column shows that a job is suspended.

4. Click the check box next to the reservation you want to restart, click the Actions menu, and click Details.

5. Click the link in the Job Status column to resume the job.

6. Select the suspended job, click the Actions menu, and select Restart Selected Task.

   The job is restarted.

Run Frequently Used Reports

You can run all available reports from the main interface, but you can also run frequently used reports from the Reservation Manager administrator interface. Running reports from the administrative interface lets you launch reports in context of Reservation Manager resource pools.

To run frequently used reports

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your resource pools.

   The Resource Pools page opens.

4. Select one resource pool, click Actions, and select Reports.

   Note: If you select multiple resource pools, the Reports option is not available.

   The Create Report dialog opens.

5. Fill in the fields and press OK.

   The report opens in a separate window.

Resource Allocation Forecast

A Resource Allocation Forecast chart lets Reservation Manager administrators view the amount of resources that are reserved for a specified time period. The chart helps administrators evaluate whether there are sufficient resources to handle the expected demand.

Administrators can specify start and end times, and whether data is shown hourly, daily, or monthly. When data is displayed daily or monthly, the peak reserved resources for the time period are displayed.

To view the amount of resources that are reserved

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click the Administer Your Reservation Manager link.

   The Administration page opens.

3. Click a bar on the Resource Pool Utilization chart on the main Administration page.

   The Resource Allocation Forecast chart appears.

4. Select data for the following fields from the drop-down lists:

   Display by

   Provides data based on hours, days, or months.

   Start Time

   Specifies start date and time.

   End Time

   Specifies end date and time.

5. Click Refresh Chart.

   The chart displays resource allocation based on the values you chose.

Static IP Addresses

The Reservation Manager can provision virtual machines and perform other network tasks using static IP addresses. By default, Reservation Manager uses Dynamic Host Configuration Protocol (DHCP), but the administrator can configure the product to use static IP addresses instead.

When a user creates a reservation, an IP address is assigned to it from a range of IP addresses the administrator defined. When a reservation is fulfilled or canceled, the IP address is released and made available for another reservation.

Setting up Reservation Manager to use static IP addresses involves the following procedures:

Define Network Address Pools (see page 579)

Create VMware Customization Specification (see page 580)

Prepare for Reservation Creation (see page 581)

Define Network Address Pools

The first step in enabling static IP addresses at your site is defining network address pools in the main user interface of the product. Network address pools consist of a range of IP addresses.

Note: Reservation Manager does not permit editing or deleting a network address pool that has been used for reservations. If you must change something like the IP address range, make a new network address pool. Also, deleting a network associated with a VM or logical partition is not allowed.

To define network address pools

1. Enter the following URL in your web browser and log in using administrator credentials.

   `https://server:port`

   The main product interface appears.

2. Click Resources.

3. Click a data center in the left pane, and select Management, Manage Network Address Pools.

   The Network Address Pools dialog appears.

4. Click + (New), and enter information in the required fields, which are Network Address Pool Name, IP Address, Subnet Mask, and VLAN ID. Click Next.

   A new page for entering network information appears.

5. Enter information in the required fields, which are Default Gateway, Preferred DNS Server, and Domain Name. The unrequired fields are optional. Click Next.

   A dialog for IP address ranges appears.

6. Enter information in the required fields, which are Starting IP Address, Ending IP Address, and Type (choose Static). Click Add.

   Note: Each address in the pool must have a DNS name entry when the pool is created.

   Note: When defining network address pools, verify that the static IP address ranges are not also assigned to another type of data center.

   The Virtual Hosts pane appears.

7. Select one or more virtual hosts, click Add, and then click Finish.

   Note: Each virtual host must be associated with the network before IP addresses can be assigned to VMs created on the host.

   Note: This step is optional at pool creation time. However, it must be performed before reservations are made.

   The network address pool is created.

   Note: Network address pools are not accessible for use until you explicitly grant access to users. See VLAN Scoping (see page 514) for instructions.

Create VMware Customization Specifications and Templates

In VMware vCenter or vSphere, create customization specifications and templates that enable static IP addresses. These specifications and templates must be used to create reservations with static IP addresses.

Note: For more information about creating VMware customization specifications and templates, see the VMware documentation.

When creating these items, consider the following information:

- **Customization specifications**—The network interface (NIC) customization must specify an IP address. The value of the IP address in the customization specification does not matter. The virtual machine is configured with a free IP address from the network IP address pool selected for this virtual machine. The NIC customization must not specify “Use DHCP” for the NIC.

Prepare for Reservation Creation

The next step in enabling static IP addresses is preparing for reservation creation.

To prepare for reservation creation

1. Log in to the Reservation Manager using the CA Server Automation administrator user credentials.

   The home page opens.

2. Click Administer Your Reservation Manager.

   The Administration page opens.

3. Click Manage your configuration settings.

   The Configure Settings page opens.

4. Click the following link in the Reservations area, set the value to true, and click OK:

   Network Selection Allowed

   Indicates whether users can select a specific network when making reservations.

   Default: false (the user cannot select networks)

   The configuration change takes effect when the next reservation is made.

Considerations for Resource Pools and Templates

When you create resource pools and templates in Reservation Manager for virtual machines, be sure to enter information for static IP addresses.

Resource Pool

Specify the ESX server attached to the network address pool you created.

Template

Specify the VMware customization specification that contains the static IP setup.

More information:

Define Network Address Pools (see page 579)
Create VMware Customization Specifications and Templates (see page 580)

Chapter 9: Reporting

This section contains the following topics:

Run Reports

You can run reports based on the metrics that CA Server Automation collects about the servers and services in your data center.

Follow these steps:

1. Click Reporting.

   The Reports page appears. At the Reports level, a menu provides frequently requested selections.

2. Expand the folder for the type of report you are interested in.

   The specific reports that are available for the category appear.

3. Select the report that you want to run.

   A progress indicator appears in the right pane and an empty pane opens with the report description.

4. Select start and end times from the drop-down list and any other prompts requested, then click Create Report.

   A progress indicator appears in the right pane and the charted data appears in this pane.

Chapter 10: Managing Changes

This section contains the following topics:

CA Configuration Automation Overview

CA Configuration Automation is a CA Server Automation option that provides management of configuration changes and configuration change detection for the resources in your data center. These resources can be systems such as servers; applications such as software components; and networks, such as IP subnets. Events are generated to indicate changes to network and server level resources.

Discovery provides discovery and classification of all entities (new nodes and subnets) in your IP network. Only Discovery can add or delete systems from the Management Database. If you add systems from CA Configuration Automation directly, use the synchronize function to discover all new systems in CA Configuration Automation and add them to the Management Database. However, if you delete systems from CA Configuration Automation directly, CA Server Automation attempts to add them back in the next recording cycle.

You can have multiple CA Configuration Automation servers in your environment and CA Server Automation can connect to all of them, collect the data, and process it all in one location. You can decide which CA Configuration Automation server manages which individual server, but each system in CA Server Automation can only have one CA Configuration Automation source.

You can create and use up to three *standard snapshots* (gold, silver, and bronze) to compare servers. A standard snapshot has the correct configuration and settings for a particular environment, as determined by the administrator. After you designate standard snapshots, you can use them to run comparisons on servers and view the results. The comparison results show you the components that are installed on each server and the differences between the two servers.

For change detection, you can create *baseline snapshots* to determine changes on a server over time. A server can only have one baseline snapshot. After you have designated a baseline snapshot, you can run change detection to see what changes have occurred. Change detection and system comparisons also generate events.

View System Details

You can view details about each managed system from the CA Server Automation user interface.

Follow these steps:

1. In the Explore pane, expand the Managed node.
2. Select the server that you want to view.
3. Click the Details tab.

   The system details appear.

View Relationships

You can view relationships between a server and all other managed servers in the network. Discovering relationships can help you determine which servers could be grouped into a server automation service.

Follow these steps:

1. In the Explore pane, select the server for which you want to view its relationships.
2. In the right-click menu, select Verification, Relationship.

   The Relationship dialog opens. It lists all servers to which the selected server is related and the type of the relationship. For a graphical representation of relationships, click the Actions button and select the Launch Visualizer option.
3. (Optional) Run network and management profiles to discover more relationships of the server. Click the Actions button and select the Find more relationships option.

   Note: Running a CCA network profile helps discover more server relationships only if the following conditions are met:

   - The network profile uses a network scan policy with Softagent. For a qualified decision which network profile to select, check what network scan policy the network profile uses and read the policy description in CA Configuration Automation.
   - The server is in the inclusion list of the network profile and there are proper credentials in the selected credential vault profile.

Create a Baseline Snapshot

You can create a baseline snapshot and use CA Configuration Automation to detect changes of configuration and resources over time, or to make system or service comparisons. Both change detection and compare processes can use a baseline snapshot.

Note: A server or service can only have one baseline snapshot.

Follow these steps:

1. In the Explore pane, select the server or service for which you want to create a snapshot.

   Note: The system that you select (or is part of the selected service) must have the CCA Agent and components installed.
2. In the right-click menu click Verification, Snapshot.

   The Snapshot dialog, which displays snapshots for the selected system opens.

3. Click the Actions button and select the Create Snapshot option.

   The Take Snapshot dialog opens.

4. Type a snapshot name and select Baseline Snapshot as the snapshot type.

   A confirmation of the new baseline snapshot appears.

Create a Standard Snapshot

You can create a standard snapshot with CA Configuration Automation to use for system or service comparisons and control purposes. CA Configuration Automation provides Gold, Silver, and Bronze standard designations. A server or service can have one snapshot of each standard designation.

To create a standard snapshot

1. In the Explore pane, select the server or service for which you want to create a snapshot.

   Note: The system that you select (or is part of the selected service) must have the CCA Agent and components installed.

2. In the right-click menu click Verification, Snapshot.

   The Snapshot dialog, which displays snapshots for the selected system opens.

3. Click the Actions button and select the Create Snapshot option.

   The Take Snapshot dialog opens.

4. Type a snapshot name and select the snapshot type (Gold, Silver, or Bronze).

   A confirmation of the new standard snapshot appears.

Assign Profiles

You can assign a discovery, management, or access profile to a selected server. Assigning profiles to servers allows you to control operations, such as discovery of specific types of applications. You can apply profiles across multiple servers of the same type when you want consistent behavior.

Follow these steps:

1. In the Explore pane, select a system for which to assign a profile.
2. In the right-click menu select Verification, Assign Profile.

   The Assign Profile dialog opens.

3. Select a profile from the Profiles drop-down list. Options include the following:

   Management Profile

   Provides the operational rules for software component management.

   Access Profile

   Provides the rules for server access and CCA agent installation.

4. Click Assign.

   A message confirms that the profile for the system has been changed. The new profile is used for the next profile execution.

Run Change Detection

Use the Change Detection page to determine configuration changes for a server or service over time. Change detection can display All Differences or only Component Differences between the current configuration (Source Configuration) and its snapshots.

Note: A server or service must have a snapshot for change detection to run.

Follow these steps:

1. In the Explore pane, select the server or service on which to perform change detection.
2. In the right-click menu select Verification, Run Change Detection.

   The Change Detection dialog opens.

   The System Information section displays the server information and current activity. For a service, the Service Detail section lists all servers that are part of the service.

   The Change Detection section provides Source and Target Configurations for the change detection operation.

3. Select *one* of the following, and click the Perform Change Detection button:

   All Differences

   Detects all differences between the source and target configurations.

   Component Inventory Differences Only

   Detects only differences in components between the source and target configurations.

The comparison results are available on the Dashboard tab in the Jobs pane. Click the Actions button, and then the View Details option to display a table with columns for the source and target configurations based on the following criteria:

Source Only

Displays components in the source configuration (but not in the target configuration) in red.

Target Only

Displays components in the target configuration (but not in the source configuration) in green.

Differences

Displays components in the source and target configurations. Differences in configurations appear in blue.

4. (Optional) When changes are detected and you want additional details about the system changes, click the + icon. The row in the table expands to display change detection details.

   Note: Additional details are not available for services.

Compare Systems and Services

Use this functionality to determine differences between servers or services. The source server or service must have a snapshot created – it is used as a benchmark for comparison.

Follow these steps:

1. In the Explore pane, select the target server or service for the comparison.
2. In the right-click menu, select Verification, Compare Systems or Compare Services.

   The Compare System, respectively Compare Service dialog opens.
3. Select one of the following options:

   All Differences

   Detects all differences between the source and target configurations.

   Component Inventory Differences Only

   Detects only differences in components between the source and target configurations.
4. In the Source System section, select a server from the Host Name drop-down list or a service from the Service Name drop-down list.

   If you compare servers, you can select the Show All checkbox to display also servers in different services.

5. In the Target System section, select either the current configuration or a snapshot.
6. Click Compare.

   The comparison results are available on the Dashboard tab in the Jobs pane. Click the Actions button, and then the View Details option to display a table with columns for the source and target systems based on the following criteria:

   Source Only

   Displays components in the source configuration (but not in the target configuration) in red.

   Target Only

   Displays components in the target configuration (but not in the source configuration) in green.

   Differences

   Displays components in the source and target configurations. Differences in configurations appear in blue.

   If the comparison detects no changes, a message appears.
7. (Optional) When the system comparison detects differences, and you want additional details about the system changes, click the + icon. The row in the table expands to display change detection details.

   Note: Additional details are not available for services.

Run CA Configuration Automation Discovery

If you are using CA Configuration Automation, you can run defined discovery on selected servers or services. Discovery determines which components are installed on servers and which components have been removed.

Follow these steps:

1. In the Explore pane, select the server or service for which you want to run discovery.
2. In the right-click menu, select Verification, Run CA Configuration Automation Discovery.

   Completion of the discovery generates an event showing the number of components discovered. You can view the event on the Dashboard.

   Note: If the discovery process does not find components, verify that the CA Configuration Automation agent is installed and running.

Run Management Profiles

Verification lets you run management profiles on selected servers or services. Management profiles enable you to collect detailed blueprint component information.

Follow these steps:

1. In the Explore pane, select a server or a service for which you want to run a management profile.
2. In the right-click menu click Verification, Run Management Profile.

   An information message indicates that the management profile operation is in progress. The management profile generates an event showing the number of components discovered. You can view the event on the Dashboard.
3. In the right-click menu for the server or service click Verification, View Components to see the results of the discovery.

   Note: If the management profile does not find components, verify that the CCA Agent is installed and running.

Add a CCA Server

You can add more CA Configuration Automation servers to your environment, as needed.

To add a CCA server

1. Click Administration, then Configuration.

   The Configuration page appears.
2. Select Verification from the list of components.

   The CCA server data is displayed.
3. Enter the CCA server name, credentials, and port number, and then select Save As Default CCA Server from the Actions drop-down menu.

   The CCA server is added to the server list.

Test CA Configuration Automation Agent

You can test a CA Configuration Automation agent to find out whether the agent is installed and running.

Follow these steps:

1. Open the Explore pane.

   Available groups, services, and systems appear.

2. Expand the Data Center folder and the CA Server Automation Services folder.

   The Data Center discovered and managed resources appear.

3. Right-click a server in the Explore pane.

   A menu appears.

4. Select Verification, Test CCA Agent.

   A message verifies whether the agent is running.

Delete a CCA Server

You can delete CCA servers from your environment, as needed. When multiple CCA servers are registered, the default server is not deleted.

To delete a CCA server

1. Click Administration, then Configuration.

   The Configuration page appears.

2. Select Verification from the list of components.

   The CCA servers data is displayed.

3. Select the check box of the server that you want to delete, and select Delete from the Actions drop-down menu.

   A confirmation message notifies you if the server was deleted.

   Note: If multiple CCA servers are registered, the default server is not deleted.

Merge CA Configuration Automation Servers

You can merge the data from your CCA servers into CA Server Automation. Use this function when you have an existing CCA server that contains multiple systems and you want to install CA Server Automation and preserve your data. Merging servers puts those systems into CA Server Automation. You must do this manually, the CA Configuration Automation service must be running, and you must have the correct credentials for the CCA server.

To merge CA Configuration Automation servers

1. Click Administration, then click Configuration.

   The Configuration page appears.

2. In the Management section, click Verification.

   The CCA servers data appears.

3. Click the Merge Servers link next to the CCA server that you want to merge with CA Server Automation.

   A confirmation message notifies you when the merge is successful.

Chapter 11: Managing Traps

This section contains the following topics:

SNMP Trap Receiver

The SNMP Trap Receiver component is the integration point for CA Server Automation to receive incoming SNMP traps. The SNMP Trap Receiver component listens on a User Datagram Protocol (UDP) port (default 162) that catches the following types of traps:

- Incoming SNMP v1 traps
- Incoming SNMP v2c traps
- Incoming SNMP v3 traps
- v2c notification traps
- v3 notification traps

SNMP Trap Receiver handles SNMP v1, v2c, and v3 traps in the same manner and in some cases, there is no noticeable difference.

All traps that are not viewed as duplicates by the trap storm window are forwarded to ActiveMQ, where other components listen for specific trap IDs.

For SNMPv3 traps, you can configure SNMPv3 trap user authorization using the sysedgev3.cf file located in the *Install_Path*\CA*productname*\conf directory.

Note: For more information about configuring for SNMPv3 in sysedgeV3.cf, see the *SystemEDGE User Guide*.

Specify TrapReceiver SNMP Port Setting

TrapReceiver lets you specify the port which it listens for traps on. You can modify the setting for this option by editing the caaipconf.cfg file.

To specify the SNMP Port setting

1. Open the caaipconf.cfg file located in the CA Server Automation *Install_Path*\conf directory with a text editor, and edit the following setting:

   CONFIG_KEY_TRAPRECEIVER_SNMP_PORT

   Specifies the TrapReceiver listening port for any incoming SNMP traps.

   Default: 162

2. Save and close the file.

   The configuration change is enabled.

Specify TrapReceiver Storm Window Setting

TrapReceiver lets you specify the interval at which duplicate traps are ignored. You can modify the setting for this option by editing the caaipconf.cfg file.

To specify the TrapReceiver Storm Window setting

1. Open the caaipconf.cfg file located in the CA Server Automation *Install_Path*\conf directory with a text editor, and edit the following settings:

   CONFIG_KEY_TRAPRECEIVER_STORM_WINDOW

   Specifies the interval window in seconds when TrapReceiver ignores duplicate traps. Default is 30 seconds.

2. Save and close the file.

   The configuration change is enabled.

SNMP Trap Receiver Configuration File

The trapreceiver.conf configuration file is an XML-like file that contains all of the SNMP Trap Receiver settings. This file is provided so you can change these settings as needed. You can configure settings to enable trap blocking, format traps, send traps to CA Server Automation Event Manager, or convert traps into CIM_SNMPTrapIndications. These traps can be sent to the Product Management Modules (PMMs). You can configure the file dynamically without restarting Apache for all settings except for SNMPPort which requires you to restart Apache. All other changes take effect 30 seconds after you save them.

Locate the trapreceiver.conf file in the Install_Path\CA*productname*\config directory and open with a text editor.

Configuration parameters include the following:

SNMPPort

Defines the UDP port to which Trap Receiver connects and listens for incoming SNMP traps. If another application is connected to the port specified by this parameter, Trap Receiver does not start. Instead, it logs a message to the Event Manager and attempts to reconnect to that port every 30 seconds. If this parameter is modified with a new port number, Trap Receiver attempts to connect to the new port.

Default: 162

Note: If you change the SNMPPort setting after Trap Receiver is connected to the SNMP port, restart Apache for the new port to take effect.

TrapStormTimer

Prevents duplicate trap handling. A configurable value prevents identical traps from coming in more often than the TrapStormTimer value setting. If the value is set to 0, then trap storm is disabled.

Default: 30

PreFilter

Filters all incoming SNMP traps before being evaluated by the Converter or Subscriber components in the configuration file.

Default: All Traps

FilterType

Applies only to the PreFilter. Options include the following:

positive

Determines which traps are permitted to pass for handling.

negative

Determines which traps are blocked.

Default: positive if <FilterType>negative</FilterType> or <FilterType>positive</FilterType> is not included.

Note: If you specify the asterisk * for all fields in negative prefilters, then all prefilters are disabled and traps are blocked.

ConvertToDCAEvent (Including Filter and Format)

Specifies a descriptive name for the converted trap and can be duplicated when specifying other trap names.

Example: <Converter Name="eHealthTrap25">

Format

Specifies a string containing %s to represent the values that are defined by the list of varbind values below the format string.

Varbind

Specifies the variables that are used to replace the %s in the format string.

Valid values include the following:

SourceIP

EnterpriseOID

GenericTrapID

SpecificTrapID

Note: The SpecificTrapID filter entry is only evaluated if generic type from the incoming trap is 6 (enterpriseSpecific). If GenericTrapID is * and the generic type from the incoming trap is 0-5, then both GenericTrapID and SpecificTrapID are matches.

Varbind1, Varbind2 - up to and including the number of expected varbinds in the trap. The number of variables below the Format must match the number of %s values in the Format.

EventSource

(Optional) Substitutes the SourceIP of the trap to another value, such as a specific varbind of the incoming trap. Some incoming traps originate from the manager and are not from the node were the trap originated.

Default: SourceIP

EventTarget

(Optional) Substitutes the localhost node name with a value from the incoming trap, such as a specific varbind.

Default: localhost

SubscribedIndicationHandlers (Including Filter and Topic)

Specifies a descriptive name for the subscriber and can be duplicated when specifying other trap names.

Example: <Subscriber Name="SUNZoneAim_PMM">

Topic

Specifies the ActiveMQ topic name where messages are sent for CA Server Automation.

Filter

These settings apply to Filter and PreFilter. They allow incoming traps to be evaluated to determine if each Converter or Subscriber uses them. If the PreFilter blocks traps, then the other Filters do not evaluate them.

SourceIP

Specifies any regular expression string variable that is a subset of the range of IP addresses.

Examples:

111.*.*.*

112.223.*.*

112.223.334.5

*

Limits: IPv4 addresses only.

EnterpriseOID

Specifies any regular expression string variable that is a subset of the range of enterprise OIDs. Use backslash (\) before any period (.) that appears in the regular expressions. The period (.) is a special character in regular expression syntax.

GenericTrapID

Specifies an integer-based regular expression that is used to compare the trap Protocol Data Unit (PDU) TrapType field. Valid values are *, a single integer 0 through 6, or a range (0-1) or (0-6).

Valid SNMP TrapType values include the following:

coldStart(0)

warmStart(1)

linkDown(2)

linkup(3)

authenticationFailure(4)

egpNeighborloss(5)

enterpriseSpecific(6)

SpecificTrapID

Specifies an integer-based regular expression that is used to compare the trap PDU's TrapType field.

Valid values can be *one* of the following:

*

A single integer

A range from (1200-2225) or (165201-165299)

Comments

Specifies comments which can be on the same line or separate lines.

<!--

Defines the beginning of a comment.

-->

Defines the end of a comment.

Example:

<!-- If the port is already connected, it does not change dynamically. Restart Apache. -->

SystemEDGE Trap Forwarding

The SNMP Trap Receiver can receive any trap from SystemEDGE when configured to forward traps to the CA Server Automation manager. The incoming SNMP traps are evaluated using a filter and can then be sent to the PMMs or converted to CA Server Automation events.

Note: For more information about forwarding SNMP traps using the sysedge.cf and sysedgeV3.cf files, see the *SystemEDGE User Guide* in this bookshelf.

Chapter 12: Remote Monitoring

This section contains the following topics:

Overview

Remote Monitoring (RM) lets you monitor the health state of agent-less systems. RM provides the flexibility of monitoring systems without the need to install the monitoring agents (such as SystemEDGE) on the remote systems.

RM employs a mid-level manager named RM AIM to monitor the remote systems. RM AIM collects the metrics information using WMI queries on the remote Windows systems.

Advantages of Remote Monitoring

Remote Monitoring involves agent-less rather than agent-based technology and there are advantages to both strategies. Use this information when deciding whether to use RM or deployed agents.

RM offers the following benefits:

- Costs less to set up, configure, and deploy
- Simplifies software upgrades and maintenance
- Deploys quickly and is less intrusive on the monitored environment
- Utilizes fewer resources on the managed server

A deployed agent offers the following benefits:

- Provides more detailed data and higher levels of functionality for the monitored servers and applications
- Requires less network bandwidth to operate
- Provides a higher degree of scalability, scaling to thousands of servers
- Continues to monitor server health and conduct data gathering when network connections are unavailable (as agent can work autonomously)
- Provides stronger command and control functions over the managed servers

Features and Benefits

Remote Monitoring provides *seamless* integration of monitoring from an end-user perspective (that is, equal look-and-feel of management interfaces for the monitored systems whether by agent or RM).

RM includes features that let you manage systems by monitoring health states and key performance indicator (KPI) metrics. RM provides reports on system status and utilization metrics. RM includes benefits such as resilience, scalability, integration, and automation. The primary features and benefits are described in the sections that follow.

Agent-less Monitored Systems

Remote Monitoring enables seamless health monitoring for systems managed with agent-based and agent-less technologies.

The RM manager component (RM PMM) creates CIM system objects representing the RM systems and their health state.

This information is presented in the Dashboard and the Resources Panel.

Follow these steps:

1. Open Resources, Explore and expand the Remote Monitoring folder.

   The discovered systems appear in the components tree.

2. Select a system.

   The page of that system appears in the right pane.

3. Open the Remote Monitoring tab.

   The agent-less gathered data appears.

Key Performance Indicator Metrics

Remote Monitoring collects and provide Windows metric information by performing WMI queries on the monitored RM systems. A rich set of information is available in various Win32 CIM classes, made available through the RM AIM.

Visualization

The RM UI lets you configure the following information:

- What systems are remotely monitored
- What metrics are collected for those systems
- If and how those metrics are monitored (including severity and threshold)

Configuration

Remote Monitoring monitors KPIs out-of-the-box on a remote system when it is selected for monitoring without requiring configuration of the monitoring being performed. You can adjust the out-of-the-box monitoring thresholds to suit your needs.

You can also define and store configuration settings in a configuration set, which can then be assigned to one or more RM systems.

Access Control

When a user logs in to the UI as admin or as a nonadmin user, security mechanisms provide authentication and authorization functionality. Remote Monitoring allows or disallows certain actions (such as configuring an RM system) based on whether the user is an admin or nonadmin user.

The RM AIM accesses the RM system through a WMI connection to the root\CIMV2 namespace (using DCOM). The local RM system administrator user and password credentials are required for access. These credentials (provided by the user when an RM system is to be monitored) are stored in a file using password encryption.

Resilience

The RM AIM is a separate process from SystemEDGE; an error in the RM AIM does not cause SystemEDGE to crash. If the RM AIM crashes or no longer responds to SystemEDGE requests, the *RM AIM alive* check in SystemEDGE restarts it.

Scalability

There is one RM AIM per SystemEDGE and each RM AIM can monitor approximately 200 RM systems. There is a single RM PMM per manager and each RM PMM can manage approximately 20 RM AIMs. The default configuration set contains ten monitored metrics with two monitors for each metric.

In terms of SystemEDGE scalability, this results in the following:

- 10 * 2 * 200 = 4000 monitorTable entries
- 10 *200 = 2000 aggregateTable entries

Integration

RM monitor information is exposed in an SNMP MIB to enable easy access for Spectrum, eHealth, and CA NSM managers.

Automation

The RM AIM includes a command line utility (rmonwatch), which allows remote configuration of RM systems and their credentials using a script.

Architecture

The following diagram provides an overview of the main RM components.

One or more RM AIMs perform monitoring for Windows servers through WMI over DCOM/RPC. Within a particular site or subnet, direct TCP connectivity from the AIM to the monitored Windows servers is required. The AIM is deployed through the deployment component.

A Platform Management Module (RM PMM) provides the interface to the manager infrastructure and creates managed objects in the CIM object model. The PMM communicates with the RM AIM using SNMP.

Interaction Between Remote Monitoring Components

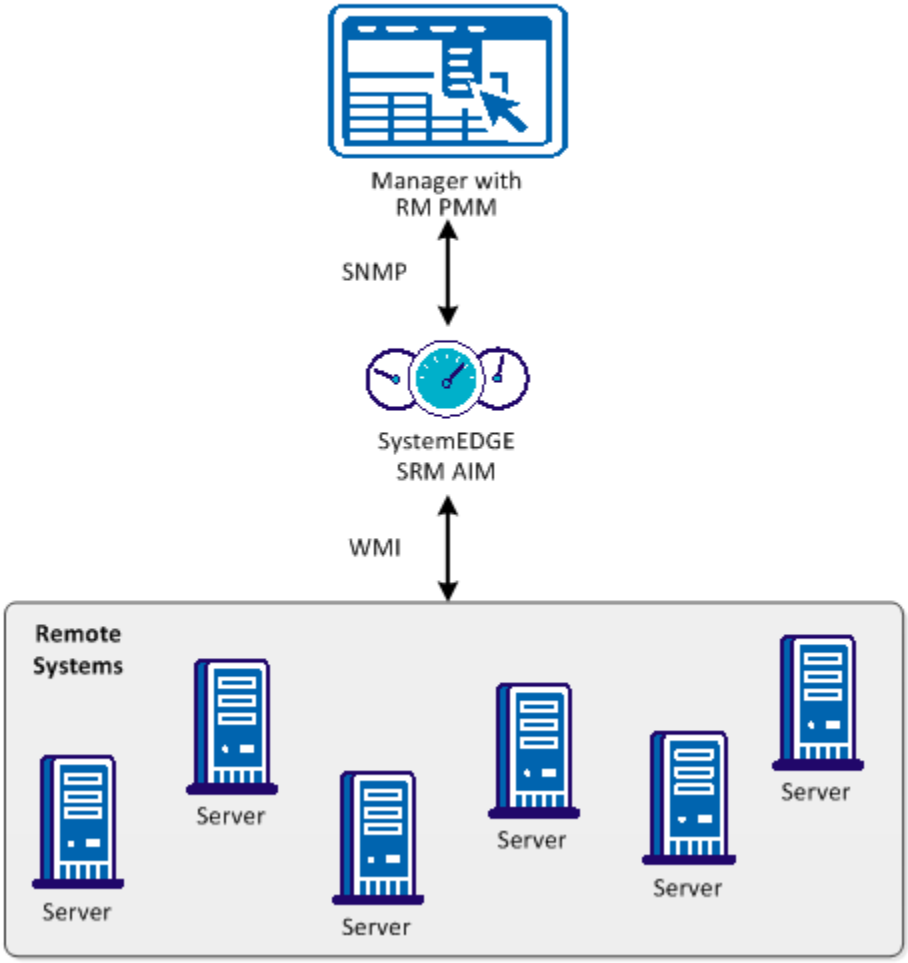

Use Case Scenario

Consider the following use case scenario for Remote Monitoring. An enterprise offers a web book store consisting of services to order books and services to browse for books.

- The order services are available from two servers in Munich and one server in London
- The browse services are available from two servers in London and one server in Munich

The servers are GERSERV1 and GERSERV2 in Munich and UKSERV1 and UKSERV2 in London. They are configured for load balancing and failover.

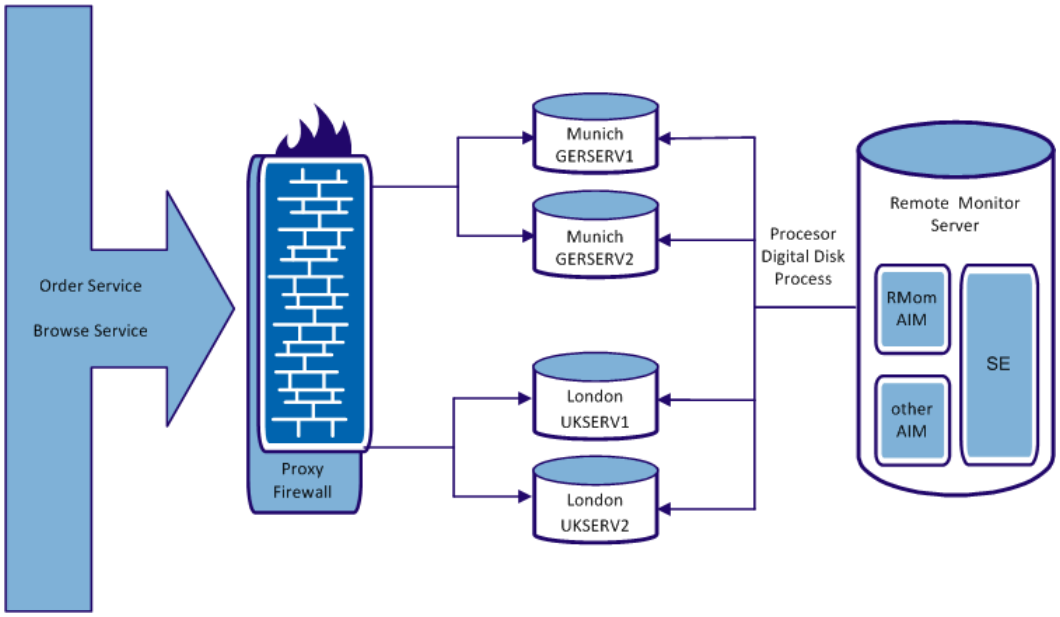

Monitoring of the services depends whether it is an order service or a browse service. In this example, two configuration sets (one for each service type) are defined. They comprise queries and monitors for the following information:

- CPU

  The percentage of the total CPU idle time.

- FSys

  The free space of the logical disk important for the respective service type (C: for an order service and D: for a browse service).

- Proc

  The working set size of the order process (single process order.exe) or the sum of the working set size of all browse processes (group of processes browse).

For each monitored system, depending on the service role (order, browse, or both), the following configuration sets are assigned:

SystemName	ConfigSet
GERSERV1	order browse
GERSERV2	order
UKSERV1	order browse
UKSERV2	browse

Configuration Prerequisites

Before you configure remote agents, verify that the following prerequisites are met.

- Firewall and port requirements for RM systems

  The RM AIM system accesses the RM system through a WMI connection. WMI uses DCOM communication which uses an End Point Mapper (EPMAP) Port TCP (135) and a DCOM TCP port which EPMAP dynamically identifies.

  To simplify configuration, the RM AIM must be located within the same firewall boundaries as the RM systems.

  Note: For more information about using a fixed port search for the article "Setting Up a Fixed Port for WMI" on the Microsoft MSDN website.

- Firewall and port requirements for the manager system

  The RM AIM utilizes the SNMP infrastructure provided by SystemEDGE; it does not require additional ports.

  RM configurations are performed using SNMP. Because the configuration data includes passwords, RM uses password encryption.

- SystemEDGE ports and SNMP

  The manager system accesses the SystemEDGE system through SNMP, which requires that the SNMP port (UDP 161 incoming) is open on the SystemEDGE system. The SystemEDGE system sends SNMP traps (UDP 162 outgoing).

- SystemEDGE ports and policy-based configuration

  The manager system accesses the SystemEDGE system through CAM, which requires that the UDP (4104) or TCP (4105) port are opened on the SystemEDGE AIM system. The SystemEDGE AIM system uses CAM to send messages to the manager system.

- WMI access best practices

  The RM AIM connects to RM Systems using WMI and requires credentials. As a best practice the RM systems must be a member of an AD Domain (for example, RIVER). This membership lets you use a domain account and avoids the need to define local user accounts on each RM System. Create a CARMuser domain account that is a member of the Domain Admins group of the AD Domain.

  When user credential settings are prompted for during RM installation, provide the domain account (for example, RIVER\CARMuser) with the password. For any system member of this domain, no additional configuration is required.

  Note: If necessary, you can restrict the CARMuser access rights so the user is not a member of the Domain Admins group. In this case, configure WMI Namespace access and DCOM access. For more information about defining WMI Namespace access and DCOM access, see the Microsoft website.

Configuring Remote Monitor Systems

A configuration set is the entity assigned to an RM System; it defines what metrics (WQL queries) are collected and how those metrics are monitored.

A configuration set consists of several configuration items. The configuration items consist of a metric definition (WQL query) and a monitoring definition (threshold, severity, and so on).

RM provides the following configuration sets with out-of-the box metric and monitoring definitions:

- default
- extended
- metricDisk
- metricFS
- metricNet

If different metrics must be monitored for an RM system with different threshold and severity settings, clone the out-of-the-box configuration set and adjust the cloned set to your system-specific monitoring needs.

The following table lists the RM metrics and the config sets to which they belong.

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 *Reference Guide*.

Collection Engine and Report Support for Remote Monitoring Metrics

The Collection Engine and Reports support a fixed set of RM metrics in the default config set.

As a result, assign the default config set (or a config set or group of sets containing those metrics) to all systems for which you want to use reports.

The supported default config set metrics are as follows:

- CPU_PercentIdle
- Disk_PercentIdle
- Event_SystemErrors
- Mem_PercentUsed
- FSys_FreeMB
- Fsys_FreeMBDecrease
- Net_QueueLength
- Proc_PercentCPU
- Proc_PercentMemory
- 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

Note: For more information about using the UI RM sub-tabs, see the *User Help*.

For the Dashboard, the following RM modules are available:

- CA SystemEDGE Machines Status
- CA SystemEDGE Objects Status

Chapter 13: Scenarios and Best Practices

This section contains the following topics:

How to Configure the vCenter Server Management Components

The following diagram provides an overview about the required actions. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the vCenter Server Management Components

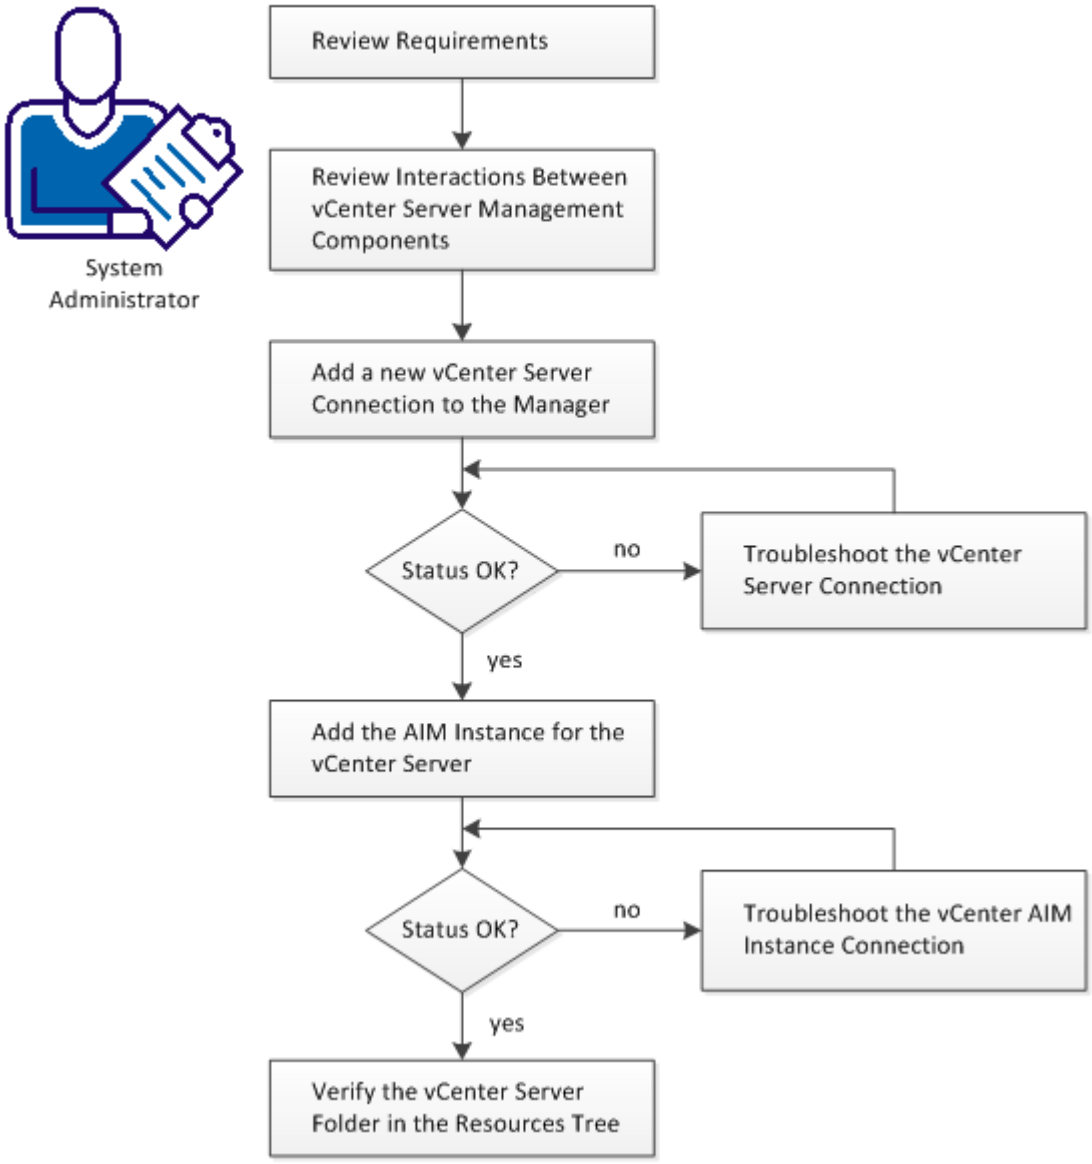

Follow these steps:

Review Requirements

Review the following requirements before you start configuring the vCenter Server management components of CA Server Automation for Infrastructure Managers:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You have a basic understanding of CA Server Automation, CA SystemEDGE, and VMware vSphere.
- You can access a CA Server Automation 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 Server Automation user interface.
- You have valid credentials available (user name and password) to access the vCenter Server of the new vSphere environment that you want to manage.
- You have found out which protocol (HTTP or HTTPS) and port to use for accessing the vCenter Server of the vSphere environment through web services. Default: HTTPS, Port 443
- You have verified that the new vSphere environment and its vCenter Server are running properly.
- If the vCenter PMM and vCenter AIM are installed on different systems, you have verified that the SNMP settings on these systems are consistent. Read and write community strings and SNMP port number must be identical.
- You have verified that the CA Server Automation manager has discovered any remote vCenter AIM Servers that you want to use.

More information:

Review Interactions Between vCenter Server Management Components (see page 618)
Add a New vCenter Server Connection to the Manager (see page 620)
Add the AIM Instance for the vCenter Server (see page 624)
Verify the vCenter Server Folder Appearance in the Resources Tree (see page 630)

Review Interactions Between vCenter Server Management Components

As a System Administrator, you want to manage a new VMware vSphere environment with CA Server Automation. CA Server Automation 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 Server Automation 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 Server Automation vCenter Server management components, that is, vCenter PMM and vCenter AIM.

The vCenter AIMs is a SystemEDGE agent plug-in which extends the functional scope of SystemEDGE. The vCenter AIM enables SystemEDGE to monitor the performance of multiple vSphere environments and to evaluate the states of monitored vSphere resources. Typical monitored resources are virtual CPUs, virtual memory, virtual switches, virtual disks, resource pools, vApps, and so on. Based on thresholds, SystemEDGE and the vCenter AIM determine the status of a monitored resource and propagates this information to the CA Server Automation manager using SNMP.

The vCenter PMM is a component of the CA Server Automation 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 Server Automation Management Database. Typical operations include but are not limited to: Creating, starting, stopping, or cloning a VM, adding, or removing CPU shares, adding memory to the VM while the VM is running, and so on.

Because the vCenter PMM and the AIM interact with each other, CA Server Automation can dynamically manage multiple vSphere environments. CA Server Automation can run operations automatically controlled by thresholds, status, and values gathered by the AIM. For example, CA Server Automation can add or remove CPU shares dynamically according to the workload of a VM.

The following diagram shows the interaction of the affected components in an example environment of four vSphere environments represented by four vCenter Servers. In general, the vCenter PMM and each vCenter AIM with its multi-instance support can connect to multiple vCenter Servers. The number of connections shown in the diagram do not specify any limitations. The required network connections are based on TCP/IP, SNMP, and web services.

Interaction Between vCenter Server Management Components

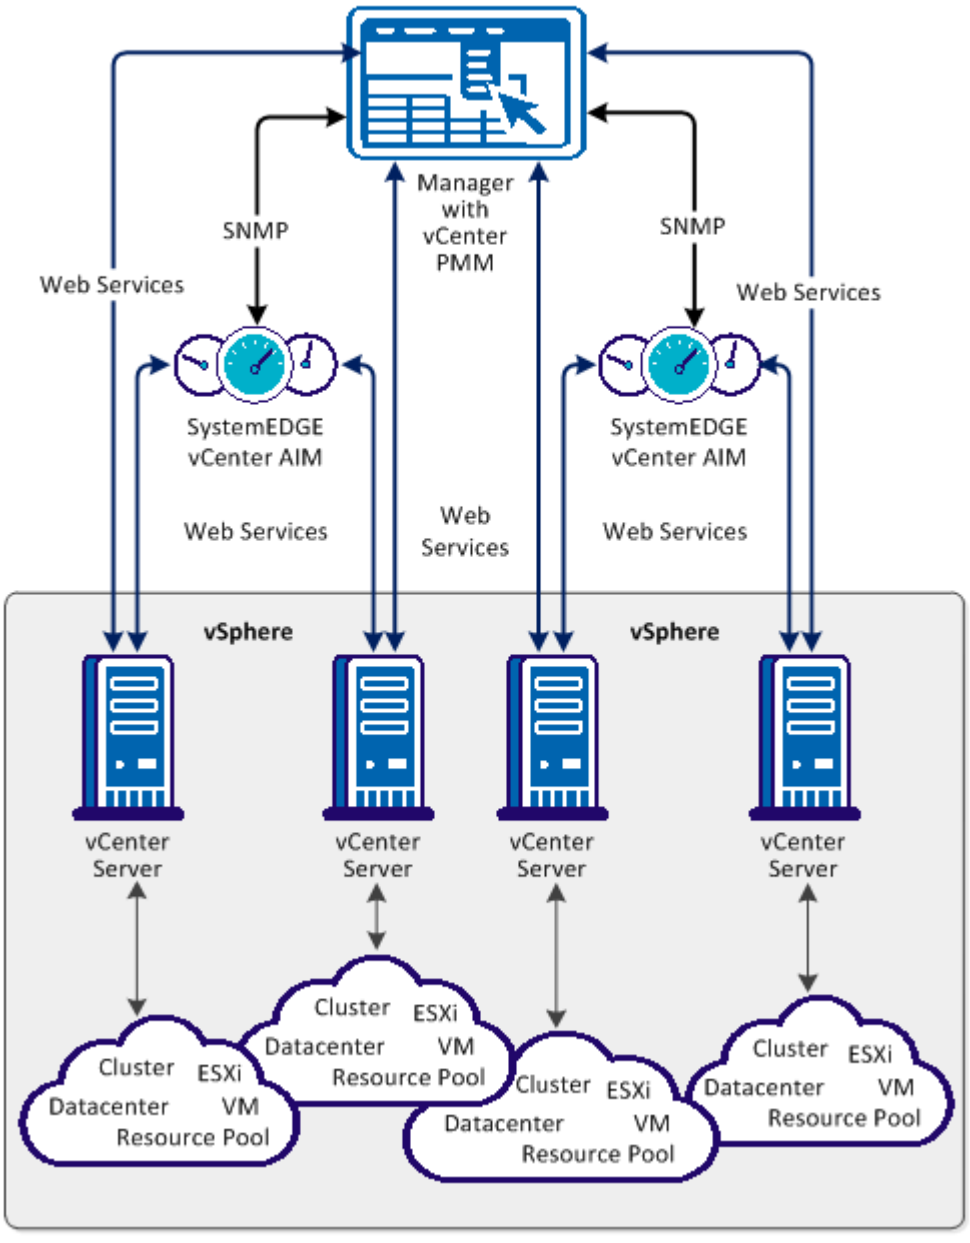

When you have configured the CA Server Automation components successfully, CA Server Automation discovers the new vSphere environment. After a successful discovery, the vCenter Server of the vSphere environment and its virtual infrastructure appear in the resources tree of the CA Server Automation Explore pane. You can then manage the new vSphere environment.

More information:

Add a New vCenter Server Connection to the Manager (see page 620)
Add the AIM Instance for the vCenter Server (see page 624)
Verify the vCenter Server Folder Appearance in the Resources Tree (see page 630)

Add a New vCenter Server Connection to the Manager

You can add a vCenter Server connection using the Administration tab of the CA Server Automation user interface.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.

   The Configuration page appears.

2. Select vCenter Server from the Provisioning section in the left pane.

   The right pane refreshes and displays the managed vCenter Servers, associated vCenter AIM Servers, and the AIM Instance for the vCenter Server.

3. Click  (Add) on the vCenter Servers pane toolbar.

   The New vCenter Server dialog appears.

4. Enter the required connection data (server name, user, password, protocol, port), enable Managed Status (checkbox), and click OK.

   If the network connection has been established successfully, the vCenter Server is added to the top right vCenter Servers pane with a green status icon. CA Server Automation discovers the vCenter Server automatically.

   If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Server Automation 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 621).

More information:

Add the AIM Instance for the vCenter Server (see page 624)
Verify the vCenter Server Folder Appearance in the Resources Tree (see page 630)
Troubleshoot the vCenter Server Connection (see page 621)

Troubleshoot the vCenter Server Connection

The vCenter Server connection has failed. Follow the troubleshooting information indicated in the following diagram:

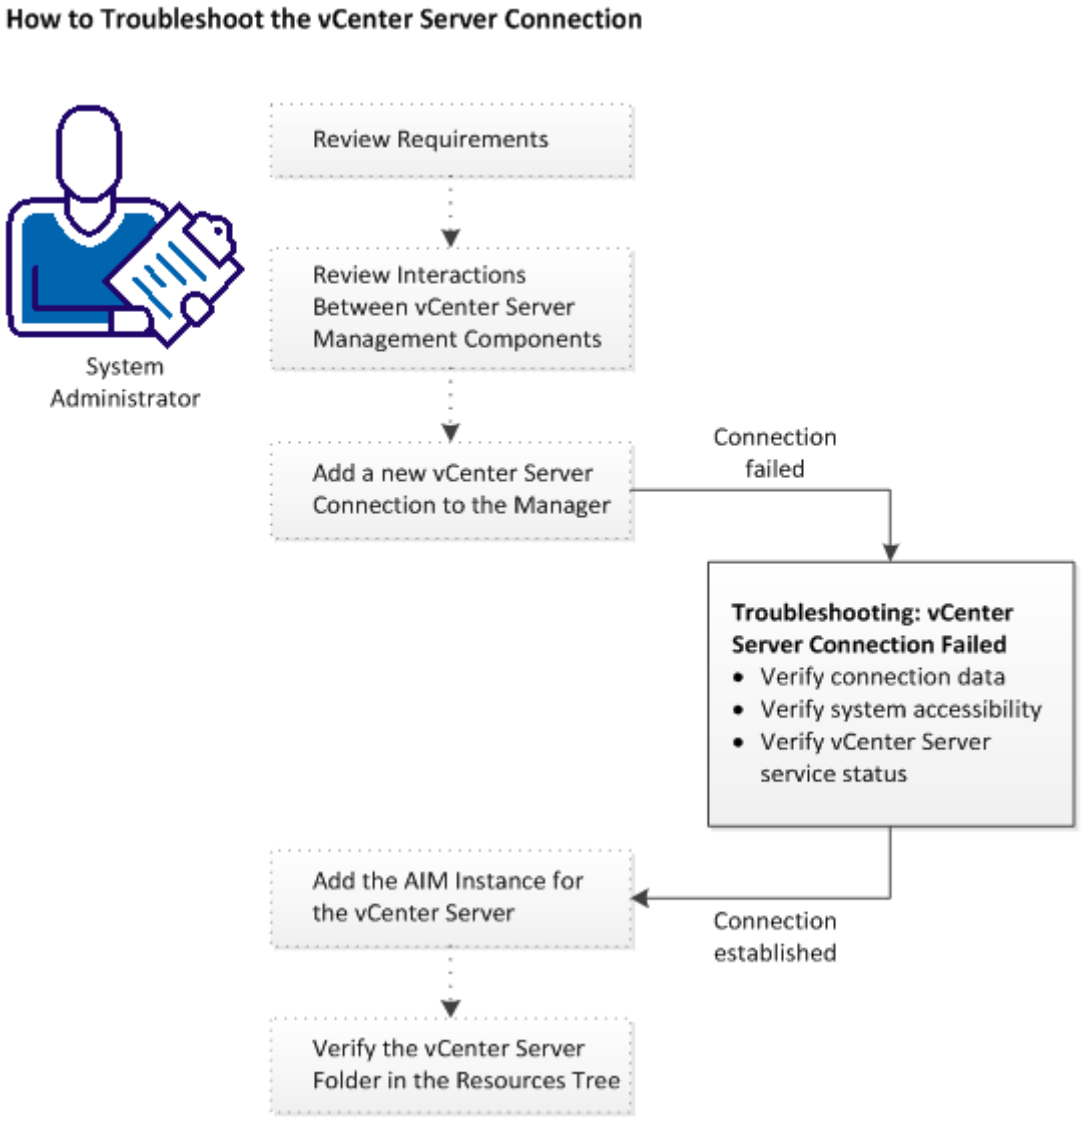

Follow these steps:

Troubleshooting: vCenter Server Connection Failed (see page 622)
Add the AIM Instance for the vCenter Server (see page 624)
Verify the vCenter Server Folder Appearance in the Resources Tree (see page 630)

Troubleshooting: vCenter Server Connection Failed

Symptom:

After I have added a new vCenter Server connection under Administration, Configuration, the validation of the connection to the vCenter Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used vCenter Server connection data (server name, user, password, protocol, port) is still valid. If necessary, update the connection data.
- Verify, if the vCenter Server system is running and accessible.
- Verify, if the VMware Management Service on the vCenter Server system is running properly.

To update the vCenter Server connection data

1. Click (Add) or (Edit) that is associated with the failed connection.

   The New or Edit vCenter Server dialog appears.

2. Add the valid server name, user, password, protocol, and port. Enable Managed Status and click OK.

   The connection data is updated.

3. Click (Validate) in the upper-right corner to validate the new settings.

   If the connection to the vCenter Server cannot be established, continue with the next procedure.

To verify, if the vCenter Server system is running and accessible

1. Open a command prompt on the CA Server Automation 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 Server Automation 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 Server Automation user interface, vCenter Server pane on the manager system and click  (Validate) in the upper-right corner.

   CA Server Automation 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 Server Automation manager, add a vCenter AIM instance to manage the new vCenter Server. CA Server Automation then discovers the entire vSphere environment with all its physical and virtual components, such as vCenter Server, ESX Servers, VMs, and so on.

Follow these steps:

1. Open the CA Server Automation 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 Server Automation 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 Server Automation 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 Server Automation 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 630)
Troubleshoot the vCenter AIM Instance Connection (see page 625)

Troubleshoot the vCenter AIM Instance Connection

The vCenter AIM Connection is in not-ready status. One of the following status icons appears:

Discovery in progress - Wait until the platform manager synchronizes all data.

Multiple instances - This instance is managed by multiple AIMs. Remove the duplicate instance.

Error - Unable to connect to the AIM. Check the network configuration.

Warning - Some instances on the AIM are still in error state.

Disabled - This instance is not managed.

Follow the troubleshooting information indicated in the following diagram:

How to Troubleshoot the vCenter AIM Instance Connection

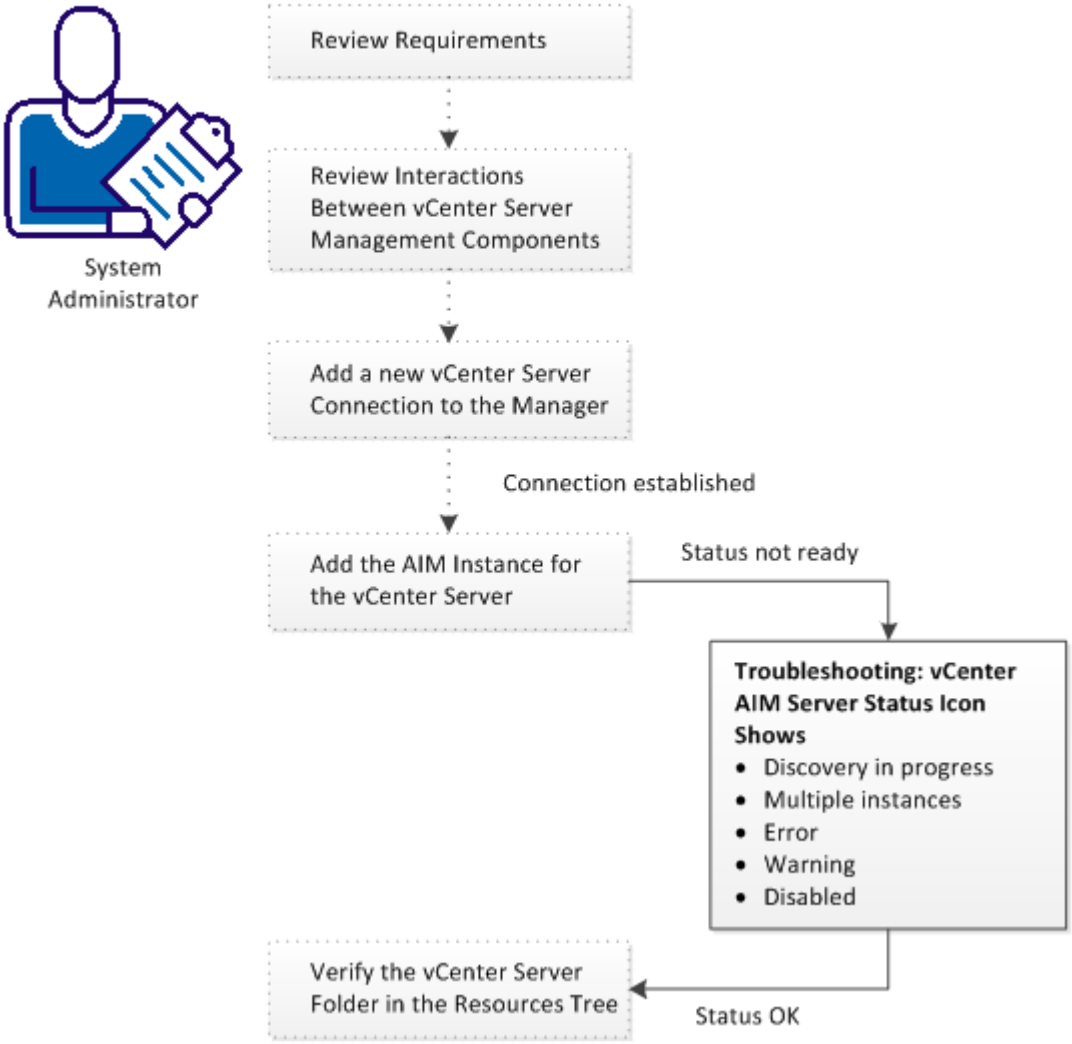

More information:

Troubleshooting: vCenter AIM Instance Status Icon Shows Discovery in Progress (see page 627)
Troubleshooting: vCenter AIM Instance Status Icon Shows Multiple Instances (see page 627)
Troubleshooting: vCenter AIM Instance Status Icon Shows Error (see page 628)
Troubleshooting: vCenter AIM Instance Status Icon Shows Warning (see page 629)
Troubleshooting: vCenter AIM Instance Status Icon Shows Disabled (see page 630)

Troubleshooting: vCenter AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Discovery in Progress).

Solution:

Wait until the discovery process of the vSphere environment has completed. The discovery duration depends on the amount of managed objects related to virtual and physical resources in vSphere. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job has completed, CA Server Automation adds a vCenter Server folder to the resources tree. Then you can start managing vSphere and its entire virtual infrastructure.

Troubleshooting: vCenter AIM Instance Status Icon Shows Multiple Instances

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Multiple AIMs manage this instance).

Solution:

Verify that your CA Server Automation manager manages each vCenter Server with one vCenter AIM instance only. If a CA Server Automation manager manages a vCenter Server through multiple AIM instances, management problems would occur. CA Server Automation stops monitoring the associated vCenter Server.

Decide which AIM instance you want to use to manage the vCenter Server and remove the other instances from the vCenter AIM Servers pane.

Follow these steps:

1. Select the AIM instance you want to delete and click (Delete).

   The Delete Item dialog appears.

2. Click Yes.

   Repeat these steps with other multiple instances until you have unique relationships between manager and AIM instance established.

Troubleshooting: vCenter AIM Instance Status Icon Shows Error

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the vCenter AIM:

- Verify, if the vCenter AIM Server is accessible.
- Verify, if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify, if the vCenter AIM server system is accessible

1. Open a command prompt on the CA Server Automation 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 hosts file on the CA Server Automation 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 Server Automation user interface, vCenter AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

   CA Server Automation validates the vCenter AIM Server connection.

   If the error status remains unchanged, verify whether the data you gathered according to the requirements for this scenario is still valid.

Troubleshooting: vCenter AIM Instance Status Icon Shows Warning

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Warning). Other instances on the AIM are still in error state.

Solution:

No specific actions are required for the associated instance. The warning informs you that other instances on the same AIM are in error state. The warning disappears when the problems of the other instances are resolved. If one or more instances of an AIM are in error state, then all other instances in the AIM show warning state.

Troubleshooting: vCenter AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered vCenter AIM instances in the network, the status icons of several instances show  (Disabled). This vCenter AIM instance is not managed.

This status appears, if CA Server Automation 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 Server Automation 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 Server Automation manager.
- Edit the existing vCenter Server connection and change its managed status to enabled.

Verify the vCenter Server Folder Appearance in the Resources Tree

After a successful configuration and discovery, the new vCenter Server is listed in the Resources Explore pane under the VMware vCenter Server folder.

Follow these steps:

1. Click Resources, Explore.

   The resources tree appears.

2. Expand VMware vCenter Server.

   The managed vCenter Servers appear.

3. Expand the new vCenter Server entry.

   The managed vSphere infrastructure appears: VMware Datacenters, ESX Servers, Resource Pools, VMs, ...

CA Server Automation is now ready to manage the added vSphere environment with its virtual infrastructure.

How to Use Policy Actions to Identify Performance Issues

This scenario provides information about how a system administrator can identify and dynamically address performance issues. This information is meant for System Administrator to optimize the allocation of resource shares of their managed vCenter environments.

The policy actions identify VM resources and dynamically adjust the allocation of CPU shares. *Shares* determine which VM gets resources when there is competition for resources among VM's. Using shares allows dynamic allocation of CPU resources. Each VM is allocated a specified number of shares. The amount of resources the VM receives depends on the proportion of its share against the total number of available shares. The allocation is dynamically changed based on the current usage of CPU resources on the ESX Server host.

If CPU usage of any VM is over 50 percent, allocation of CPU shares increases dynamically. If CPU usage is less than 30 percent, the CPU shares allocation decreases dynamically. The policy component not only identifies the problematic virtual machines but helps ensure dynamic actions which sustain business continuity. Using policy actions helps ensure that resources are allocated where to virtual machines that are in need and deallocate when the need is gone.

How to use policy actions to identify performance issues

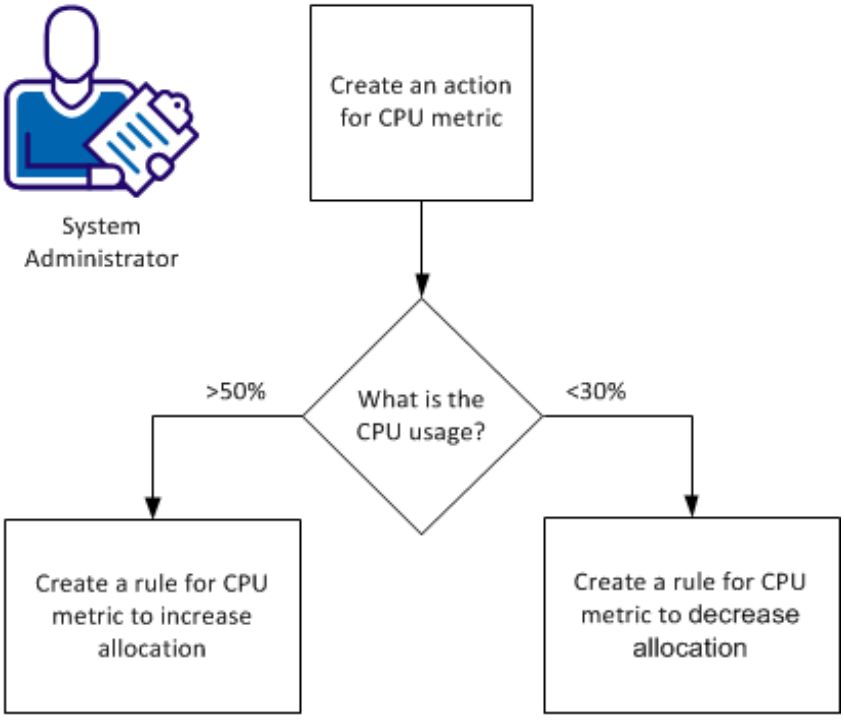

To identify and address performance issues using policy actions, follow these steps:

1. Create an action for CPU metric. (see page 632)
2. If CPU usage is more than 50 percent, create a rule for CPU metric to increase allocation. (see page 633)
3. If CPU usage is less than 30 percent, create a rule for CPU metric to decrease allocation. (see page 633)

Create an Action for CPU Metric

Policy provides the creation of rules and actions that can be used to create policies for the automated management of systems. Custom actions can be created for actions not included in the default library.

Follow these steps:

1. Click Resources.

   The Resources page appears.
2. Open the Explore pane.

   Available groups, services, and systems appear.
3. Click Policy, then click Actions.

   The Actions page appears.
4. Click '+' on upper right side bar to add new action.
5. Enter name of the Action.
6. Select Resource Configuration from Category drop-down list.
7. Select Configure Shares from Type drop-down list.
8. In VC Server field, leave the entry as "%VCServer%" to apply this action on any VM across any VC Server.
9. In VC Data Center field, leave the entry as "%DATACENTER%".
10. In the Target VM Machine field, leave the entry as "%VMNAME%".
11. Select "Set CPU" from Operations drop-down list and enter Values as "10000".

    The number is arbitrary and the share values are set to normal.

    Note: Use higher or lower numbers to increase and decrease the share allocations accordingly.
12. If the changes require approval, enable Help Desk Approval.

    A Message will appear in the Event Console after the Action is created.

    CAAP4521 Policy: Action <action name> was created.

Create a Rule for CPU Metric to Increase Allocation

Creating a rule for CPU metric to increase CPU allocation helps ensure dynamic resource allocation when the usage exceeds the threshold.

Follow these steps:

1. Click on Resources tab, Policy, Rules.
2. Click '+' on upper right side bar to add new rule.
3. Enter Name of rule and click Next.
4. Select the action from the "Action Selection" list for the rule and click Next.
5. Enter the metric-based rule where CPU usage is greater than 50 percent to increase the CPU shares of the VM.

Create a Rule for CPU Metric to Decrease Allocation

Create a Rule for CPU metric to decrease allocation

Follow these steps:

1. Click on Resources tab, Policy, Rules.
2. Click '+' on upper right side bar to add new rule.
3. Enter Name of rule.
4. Select the action from the "Action Selection" list for the rule.
5. Enter the rule-based on metric where CPU usage is less than 30 percent to decrease the CPU shares of the VM.

Scalability Best Practices

This section provides best practices and recommendations for the deployment of CA Server Automation. The purpose of the document is to assist with the planning of a roll out of CA Server Automation within a production environment, with particular focus on:

- Monitoring and CA Server Automation Management of VMware Environments
- Deployment of SystemEDGE and other Monitoring Software
- Initial and on-going configuration of SystemEDGE

The following sections are included:

Remote Deployment and Policy Configuration Overview

CA Server Automation provides a comprehensive solution for remotely deploying the SystemEDGE agent to all managed systems. You can create deployment templates based on provided packages that contain customized installation parameters and simultaneously deploy these templates to numerous managed systems.

In addition CA Server Automation provides a comprehensive solution for the ongoing configuration the SystemEDGE agents running on all managed systems. Policy Configuration allows a library of Policies to be created. These policies are applied to one or more systems running SystemEDGE and SRM AIM. When an agent managed by Policy Configuration is installed, it automatically requests a policy. As a result, the agent runs a controlled and consistent set of Policies. Each agent can then be updated individually, based on the Policy the agent is running, or the service the agent is a member of.

Remote Deployment and Policy Configuration share the Domain Server and Distribution Server technology. This technology provides a solution that is both scalable and able to be distributed across multiple data centers.

The diagram illustrates the basic architecture of the Remote Deployment and Policy Configuration components

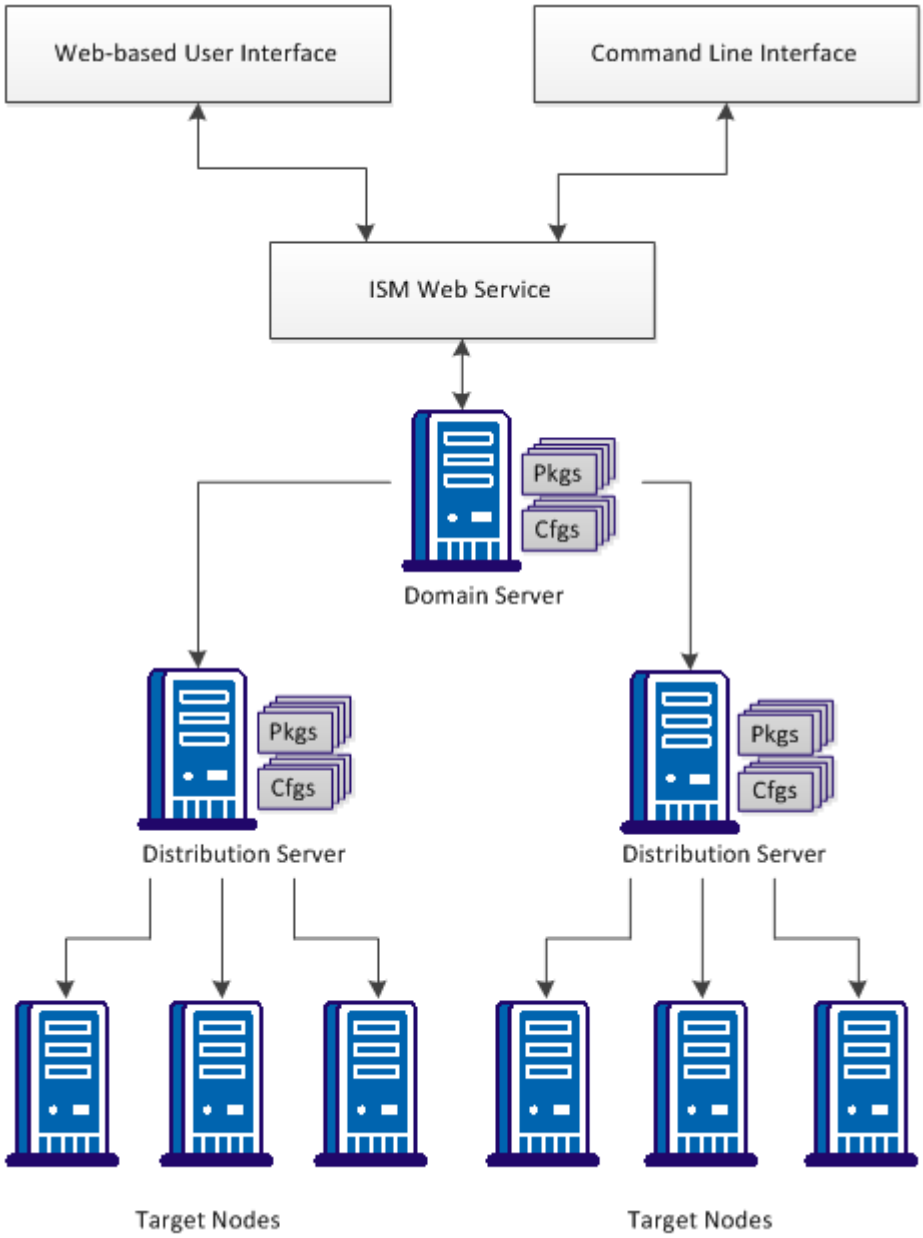

Hardware Specifications

This section lists the minimum hardware specifications for large-scale implementation of CA Server Automation. For larger scale implementations, consider increasing the specification of the management servers.

- Domain Server: 2.6-GHz Dual-core Processor, 4-GB RAM, 100-GB disk.
- Distribution Server: 1-GHz Single Core/Processor/Virtual Processor, 2-GB RAM, 100-GB disk, 100-Mb/sec Ethernet.
- VC AIM Monitoring Server: Dual Core Processor: 2.6-GHz CPU, 4-GB RAM, 100-GB disk.

  Note: 50 percent is the maximum usage allowed for CPU and memory.
- Target System: 1-GHz Single Core/Processor/Virtual Processor, 512-MB RAM, 2-GB disk, Single 100-Mb/sec Ethernet.

Database Considerations

As the managed environment grows larger, more database activity can be expected. The product databases grow in size based on product usage, potentially consuming 30 GB or more, depending on the maintenance that is being done. We recommend the following general rule for data retention: for every 1000 machines in your monitored environment, increase the poll interval by 300 seconds.

Note: Using a dedicated standalone system for the database can improve its performance. Keep the database close to other Virtual Assurance for Infrastructure Managers components on the network, for example, on the same subnet, to improve response times.

Network Considerations

When planning the roll out of CA Server Automation, consider the quality of the network connections to decide where to locate management components. The following items influence the scalability and effectiveness of the solution:

- Network quality: Poor quality results in data loss, causing slow response, or failures.
- Network bandwidth: Lower bandwidth limits the rate at which data can be sent between the components.
- Network latency: Higher latency (delay) limits the rate of data transfer, in a similar way to low bandwidth.
- DNS: Badly configured DNS hinders the deployment and ongoing configuration of the monitoring agents.

We recommend using at least 100-Mb/sec Network Links between management components. Especially when using a remote DB. If the network speed is less than 100Mb/sec, consider introducing additional Distribution Servers collocated with the target systems.

Scalability Recommendations

This section provides information about scalability recommendations and limitations.

Consider the following information:

- Monitoring of VMware Environments (see page 638)
- CA Server Automation Management of VMware Environments (see page 639)
- Remote Deployment and Policy Configuration Recommendations (see page 641)
- Domain Server Recommendations (see page 643)
- Distribution Server Recommendations (see page 643)
- Scalability Use Cases (see page 643)

vCenter AIM Monitoring Recommendations

The SystemEDGE agent provides a plug-in architecture through which it can load optional *application insight modules (AIMs)* when it initializes. AIMs are functional extensions to the SystemEDGE agent. For example, the vCenter AIM enables SystemEDGE to manage vSphere environments through VMware vCenter Servers.

The vCenter AIM (Application Insight Module) is a pluggable component implemented within the SystemEDGE framework. As such, the data that it publishes is available to multiple SNMP Managers. Products such as CA Server Automation Manager, eHealth, and Spectrum IM can leverage this data.

Due to the possibility of this component being utilized outside of the CA Server Automation Manager, scalability recommendations are specified separately. Two main scalability considerations are discussed:

General Recommendation for vCenter AIM Monitoring

The general recommendation for scalability limitations for vCenter AIM Monitoring is as follows:

- Maximum Number of VMs is approximately 240,000 / (x + 6)
- Maximum Number of Objects is approximately 2,000,000 / (x + 6)

  where x is the number of SNMP polls against AIM per hour.

Scalability Limitations in Terms of Monitored Objects

In general, CPU usage is the primary concern when scalability limitations are considered. For vCenter AIM monitoring, consider the three main factors that influence the CPU usage:

- The dynamic nature of the vCenter Servers being monitored.

  The level of activity of vCenter Servers influences the CPU consumption. The following scalability recommendations assume an average level of activity of the vCenter Server being monitored.

- The number of SNMP Managers, and polling intervals of those SNMP Managers.

  Large numbers of SNMP Managers polling the vCenter AIM, or short poll intervals, result in increased CPU consumption.

- The ratio of object count in relation to VM count.

  An *object* is any element of vSphere that is monitored by the vCenter AIM. For example, vCenter AIM monitors Datastores, Virtual Disks, Physical Network Interface Controllers, Virtual Switches, SCSI Controllers, ESX Host Hardware Sensors, and so on. The number of objects directly impacts CPU consumption. Due to the need to maintain the vCenter AIM cache and the additional overhead required for publishing this data. In real-world systems, there are typically from 6 through 11 times as many objects as there are Virtual Machines within a given vCenter.

Based on the preceding factors, a single SNMP Manager monitoring vCenter AIM with a 10-minute poll interval has a limitation of approximately 20000 VMs.

Scalability Limitation in Terms of Monitored Servers

The vCenter AIM functions in a multiple vCenter Server environment. In fact, the framework of the vCenter AIM results in a slight reduction in CPU usage as the number of VMs per vCenter Server decreases. For example, CPU usage is lower against three vCenter Servers, each with 2000 VMs, than against a single vCenter Server with 6000 VMs.

With that noted, there are points at which the responsiveness of the vCenter AIM becomes the scalability limit. In general, vCenter AIM is able to monitor up to ten vCenter Servers.

CA Server Automation vCenter Management Recommendations

The VA-IM manager does far more than simply monitor and publish vCenter data. It has the responsibility for storing and managing historical data, performing active operations against the vCenter, running automation policy based upon vCenter data, reporting, and so on. As such, it often requires more resources than the vCenter AIM, which it uses as its main data collection mechanism.

vCenter Management Limitations in Terms of Virtual Machines

Due to the fact that the vast majority of the CA Server Automation manager resides within a single process space, Operating System limitations tend to be the primary culprit in scalability. Consider the following limiting factors:

- Available Memory: As the number of objects being managed increases, the amount of memory required to cache data and handle messaging increases rapidly. We recommend increasing the memory of manager to at least 8 GB.
- Available CPU: The CA Server Automation manager requires significant CPU resources, especially in times of rapid environmental change, or during initial startup. We recommend supplying additional CPU (3.2 GHz or larger) for moderately large managed environments to improve the responsiveness of automated processes.
- Operating System Limitations: Much of the CA Server Automation Manager resides within a single process space. As a result, the memory addressing space of a 32-bit operating system can become exhausted, even when system memory is not exhausted. To avoid this issue, we recommend using 64-bit processor and operating system for a moderately large managed environment.

Examples:

The following examples provide requirements and scalability limit recommendations for the CA Server Automation manager:

- Minimum Requirements (32-bit, 2.6-GHz CPU, 4-GB RAM, 100-GB disk)

  Scalability limit: 2500 Computer Systems (that is, VMs and ESX Hosts)
- Recommended System (64-bit processor and Operating System, 3.2-GHz CPU, 8-GB RAM, 100-GB disk)

  Scalability limit: 8000 Computer Systems (that is, VMs and ESX Hosts)

Performance Considerations during Initial Discovery

Performing the initial discovery and database load of the vCenter environment you want to manage can be a time consuming process. During this process, the following take place:

1. The vCenter AIM parses the entire vCenter environment and publishes the results for the managers.
2. The CA Server Automation manager retrieves the published data from the vCenter AIM and creates an internal cache for processing.
3. The internal cache is synchronized with the current CA Server Automation manager database contents, with discoveries performed for Computer Systems not currently within the database.

During initial management of the vCenter server, the database has no Computer Systems in the database, so all of these objects are discovered and created. Consider the estimated time to complete initial discovery. Based upon baseline testing, the average throughput is: from eight to nine Computer Systems per minute.

Examples:

The following examples provide the sizes of environments and corresponding estimated times to complete initial discovery:

- 2,500 Computer Systems - approximately five hours
- 8,000 Computer Systems - approximately 15 hours

During this initial population, CPU usage may be high for extended periods.

Note: The initial discovery process takes the vast majority of this time. However, the initial discovery is done once for the lifetime of the product. The vCenter AIM and CA Server Automation internal caching processes, which are done whenever these processes are restarted, are considerably quicker. For example, 2500 Computer Systems are typically published through vCenter AIM and fully cached by the CA Server Automation manager in approximately 5 minutes.

Remote Deployment and Policy Configuration Recommendations

Consider the following aspects and recommendations to improve remote deployment and policy operations:

- Number of target machines

  For optimal performance, limit the deployment job size to 500 target machines in a batch.

- Number of Distribution Servers

  Deployment throughput is better when multiple distribution servers are used.

- Deployment package size

  The smaller the deployment software package the better the throughput. The numbers recommended assume that all of the managed servers are required to have SystemEDGE and Advanced Encryption.

  Note: A typical package size ranges from 10MB through 20MB.

- Quality and speed of the network

  Low bandwidth, high packet loss and high latency between the Distribution server and the target affect the rate and reliability of Deployment and Configuration operations.

- The timescale for the rollout of Monitoring Software to the target systems

  Stagger the deployment of monitoring software using collocated deployment servers. Staggering reduces the load on the network infrastructure. This load reduction can be achieved by creating a number of jobs, or by using the Staggering capability built into the solution.

  If the monitoring software must be deployed over short time, we recommend deploying additional Distribution Servers in the environment.

- Frequency of agents reconfiguration (by Policy Configuration)

  Reconfiguration of SystemEDGE Agents is expected to take around 30 seconds plus from 2 through 10 seconds per agent in a typical network environment. If agents must be reconfigured frequently, we recommend deploying additional Distribution Server in the environment.

- Geographical distribution of the target systems across multiple sites

  Where the target systems are distributed across multiple sites, we recommend deploying a Distribution Server at each location. This recommendation is especially true if the remote data centers use a slow (< 100Mb/sec) or unreliable link. Deploying a local Distribution Server allows all Deployment and Configuration requests to be directed through the on-site Distribution Server. This limits the network traffic between the central and remote site.

- Communication ports

  Remote Deployment and Policy Configuration rely on communications by CA-Messaging for Domain Server to Distribution Server and Distribution Server to Agent communications. CA-Messaging communicates over ports 4104(UDP) and 4105(TCP). For remote sites that are firewall protected, placing a Distribution Server on site allows all CA-Messaging communications to be set up as point-to-point.

  Note: For agent discovery and ongoing monitoring, SNMP communications (typically port 161) are used. Open this port for direct communication from the managed systems to the manager.

- Management of Policy Configuration Policies and Templates

  We recommend organizing the monitoring requirements into Templates based on the different workloads in your environment. Use a maximum of 1000 monitors per Policy or Template. Any number of templates can be applied to a system, but we recommend limiting the number of templates to 100 per system.

- Service Membership

  To ease configuration operations, we recommend organizing servers into Services, with an upper limit of around 500 servers per Service. A server can be a member of multiple Services, and therefore we suggest creating services to represent different workloads. A template can be then applied directly to a Service.

Domain Server Recommendations

A *deployment and configuration domain server* (domain server) manages and controls all deployment and policy configuration operations.

If the number of target systems exceeds 10000, we recommend running multiple instances of CA Server Automation.

Distribution Server Recommendations

A *deployment and configuration distribution (scalability) server* helps ensure that deployment and policy configuration operations are carried out in an efficient and timely manner.

Once the CA Server Automation manager is installed, the next step in the rollout of Remote Deployment and Policy Configuration is to consider the number of Distribution Servers.

For Deployment operations, we recommend using one Distribution Server per 2000 target systems for a typical 100-Mb/sec network environment.

Where the SystemEDGE Agent is not being deployed using Remote Deployment (that is, only Policy Configuration is used), this number can reach 3000 systems per Distribution Server.

Important! We recommend performing a test deployment to at least one system for each distribution server before large deployment operations. We recommend deploying all possible packages using the distribution server to verify that there are no failures with the larger scale deployments.

Scalability Use Cases

This section provides use cases that aim to represent typical production environments. We suggest comparing these cases with your own environment and following the recommendations of the closest one.

Departmental Data Center

In this use case, monitoring is required for 1000 systems, all contained within one data center. All systems are in one location, within the firewall, and fast links exist between the computers.

Recommendations for this environment are:

- Component Installation

  All CA Server Automation manager components can be installed on the same system.

- Initial Deployment

  Two jobs could be created to deploy the monitoring software to all target nodes. Depending on network load, the initial deployment of SystemEDGE (and Advanced Encryption if necessary), would be expected to complete in 8 to 12 hours.

  Once deployment to remote systems has completed, the CA Server Automation Manager discovers SystemEDGE. Policy Configuration delivers an initial policy to each agent. We expect the initial policy delivery to complete approximately eight hours after the completion of the job.

- Service membership

  For ease of maintenance, it is suggested to align the monitored servers into Services, with approximately 200 servers per Service. Services could be aligned to business function, network topology, or other categorization as desired.

- Applying Policies

  If necessary, applying policies to all monitored systems can be performed in a single operation.

Multiple Data Centers

In this use case, 10000 systems, spread across multiple Data Centers, are being managed. The number of systems per data center varies from 500 through 2500. The Data Centers are geographically distributed across multiple locations. The links between the Data Centers including leased lines of less than 100Mb/sec. The systems run various workloads, and are managed by a number of different departments (application owners).

Recommendations for this environment are:

- Component Installation

  Install the CA Server Automation manager components on a dedicated server. This server must meet the minimum supported specification, but a quad-core server with at least 8-GB RAM is recommended.

  We recommend installing the Database on a separate dedicated server with an increased specification of quad-core processor and 8-GB RAM.

  To support the remote data centers, we recommend installing one Distribution Server in each data center. For a data center that contains more than 2000 servers, a second Distribution Server is recommended. Each Distribution Server can scale up to 3000 servers if only used for Policy Configuration.

- Initial Deployment

  The following are relevant factors to consider while planning to deploy using multiple distribution servers:

  - If possible, limit a single deployment job size to a maximum of 500 target systems.
  - Verify that the nearest distribution server is chosen for the deployment operation.
  - Although concurrent deployments are supported within a single distribution server, it is advised to only use concurrent deployments across multiple distribution servers. If you have 4 distribution servers, you could start a job to deploy to 500 machines for every distribution server.
  - When you perform concurrent deployment using the same distribution server, verify that the second job does not deploy to the same set of target machines as the previous one.
  - Where multiple packages are being delivered to many systems, consider splitting the package deliver to multiple jobs. For example, by deploying SystemEDGE and Advanced Encryption first.
  - If you do see failures during a large-scale deployment, verify all prerequisites then use “resubmit job” to retry the deployment.
  - To help with future deployments, we recommend saving the deployments as templates.

- Service membership

  For ease of maintenance, we suggest aligning the monitored servers into Services, with a maximum of 500 servers per Service. For remote data centers, we recommend creating one or more Services (depending on data center size) to represent data center.

  Where multiple departments use particular systems, the systems can be added to multiple services for ease of management by each department.

- Applying Policies

  In this use case, the servers run various workloads. We therefore recommend that the base policy contains only the control settings. Hold your monitoring configuration in templates based on different monitoring requirements. These templates can then be applied to the required systems, either by manually selecting systems, or by applying templates to a service.

  Splitting the monitoring requirements into templates allows the templates to be applied to the required systems, independently of each other. Select the systems manually, or apply templates to a service. We recommend applying templates in batches of 2000 to 2500 systems.

  If the base policy is required to be changed, we recommend applying the policy to systems in batches of 2000 to 2500 systems.

  Note: When templates are used, each delivery of a template or policy involves merging of all assigned templates with the base policy. The next step is delivering the resultant configuration to the agent. Therefore, where multiple templates are applied to a system, the time for delivery may be slightly increased.

Large Environments

In this use case, approximately 21000 Agents are being managed. These agents are spread across three data centers, with each data center containing from 2000 through 10000 targets. The data centers are distributed, but have fast links between them. The managed systems are largely virtualized, run various workloads and are reprovisioned at times.

Recommendations for this environment would be:

- Component Installation

  We strongly recommend running an instance of CA Server Automation in each data center, such that each instance manages up to 10000 systems.

  Within each data center, we recommend installing the CA Server Automation manager components on a dedicated server. This server must meet the minimum supported specification, but we recommend an increased 8-GB RAM.

  We suggest installing the Database on a separate, dedicated server with 8-GB RAM.

  If a single instance of CA Server Automation is required, we recommend using Manager and Database Servers with quad-core processors with 16-GB RAM.

  To support Remote Deployment and Policy Configuration operations, we recommend installing a Distribution Server in each Data Center (one being installed on the Manager system).

- Initial Deployment

  In this use case, we have one additional distribution server in the datacenter. Apart from that one difference in the setup, all the factors that were highlighted in the previous scenario apply equally to this scenario.

- Service membership

  For ease of maintenance, we suggest splitting the monitored servers into multiple Services, with a maximum of 500 servers per Service.

- Applying Policies

  We suggest limiting the base policy to control settings and 'base OS' monitors. Where different images are being used as the basis for virtual machines, a base policy can be created for each OS image. Verify SystemEDGE is configured to request this policy on registration.

  For application-specific monitors, create templates based on individual monitoring requirements. To avoid index conflicts across templates, we suggest defining 'index ranges' up-front for each application. Alternatively, the base policy can be configured to "Automatically Resolve Indexes" under the 'Control settings' section.

  Splitting the monitoring requirements into templates allows the templates to be applied to the required systems, independently of each other. You can either manually select systems, or apply templates to a service. We recommend applying templates in batches of 2000 - 2500 systems.

  If the base policy is required to be changed, we recommend applying the policy to systems in batches of 2000 - 2500 systems.

  Note: When templates are used, each delivery of a template or policy involves merging of all assigned templates with the base policy. The next step is delivering the resultant configuration to the agent. Therefore, where multiple templates are applied to a system, the time for delivery may be slightly increased.

  Important! Contact CA Support if multiple instances of CA Server Automation are deployed and if want to share the created policies between the different instances. CA Support can assist with the export and import of policies and templates between CA Server Automation instances.

How to Apply Policy and Layered Templates to Servers

From the CA Server Automation 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:

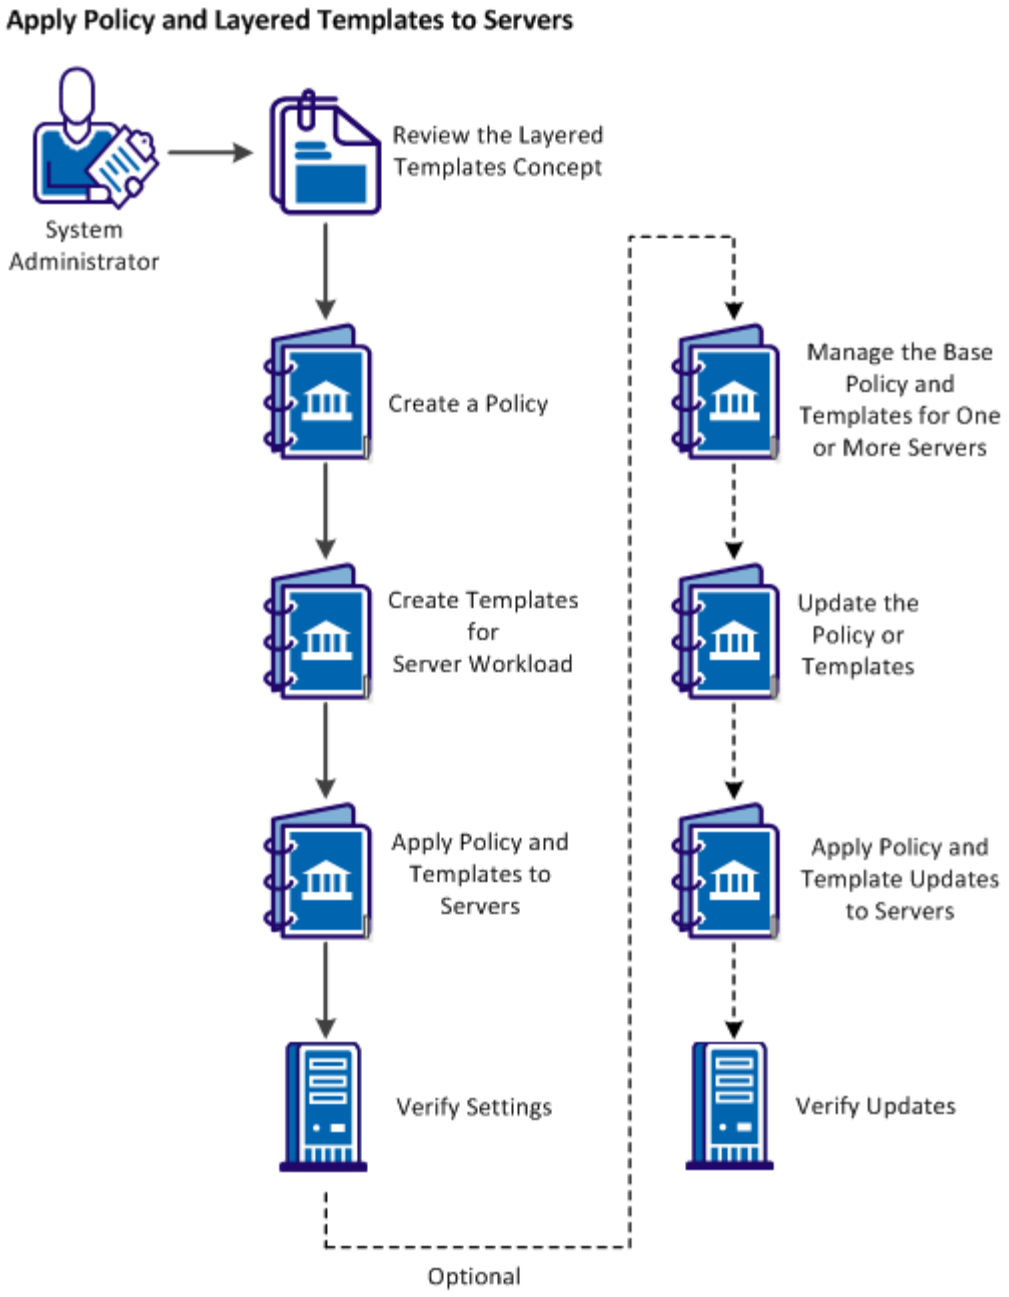

Follow these steps:

Layered Templates Concept

In an enterprise, the workload handled by a server or a server group varies. You can create multiple policies specific to the workload handled by a server or a server group. To assist in the creation of policies, templates are used to create application-specific monitors. The Base Policy and Layered Templates are combined to form a configuration file and applied to servers that you want to monitor. You can add or remove Layered Templates. Template updates can be applied directly to servers, without changing the base policy or reimporting the updated template into the Base Policy.

Example: Apply Base Policy and Templates to Servers

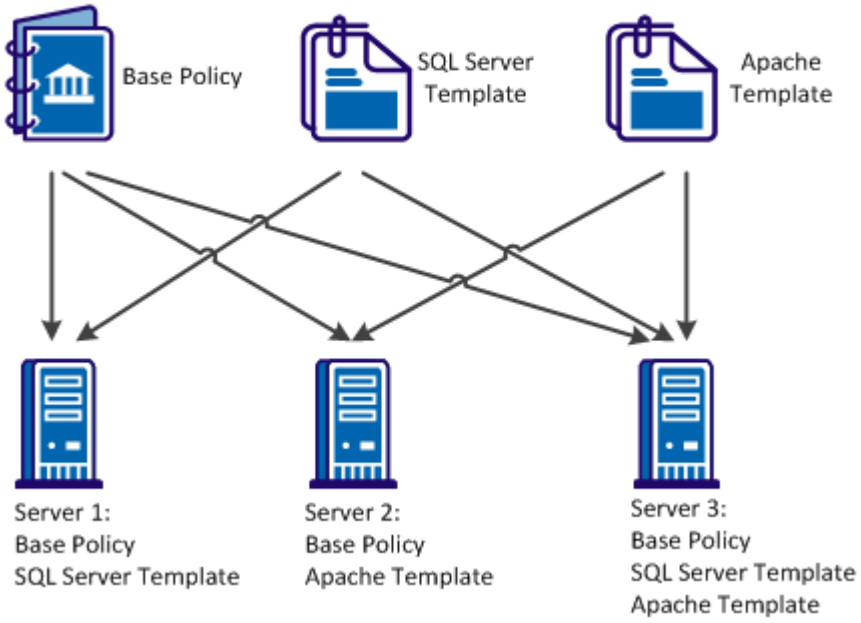

You can use Layered Templates in the following scenarios:

Disparate applications

Create a library of templates for each server running a different set of applications. You can directly apply the template updates to each server.

Dynamic environments

The workload of the servers changes frequently in dynamic environments. You can use Layered Templates to segregate the monitors in logical groups. Based on the workload changes, you can directly apply the logical groups to systems or removed from systems.

Shared servers

In an enterprise setup, servers are shared across multiple departments. Each department manages and monitors applications on the shared server. You can use Layered Templates to independently manage and apply templates to systems of each department.

Application maintenance

You can split monitoring into multiple templates. In a server, you can remove a template for an application not in use, without affecting the monitoring of the remaining system.

Create a Policy

Create a Base Policy to define a set of Monitors, MIB Extensions, Traps & Communities, and Control Settings to control the agent monitoring.

Common settings in Traps & Communities and Control Settings are available for policies only. If you use Layered Templates, common settings are specified in the Base Policy.

Follow these steps:

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

   The Available Policies page appears.

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

   The New SystemEDGE Policy dialog appears.

3. Enter a name and an optional description for the policy, the system type and whether to base it on an existing policy and click Ok.

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

4. Click Save Policy.

   The policy is created and saved.

Note: You can also use the existing default policy as a Base Policy, if necessary.

Define SystemEDGE Policy Control Settings

You can control the following agent behavior using the SystemEDGE policy control settings:

- Security settings
- SNMP settings
- MIB table population
- UNIX settings
- Performance monitoring settings

You can segregate these common control settings from specific server workload configurations by adding them to the Base Policy.

You can apply the control settings defined in the policy to all systems you want to monitor with this configuration.

Follow these steps:

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

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Click Control Settings.

   The Controls page appears.

4. (Optional) Click Use Defaults.

   The default selections pane appears. You can change the default settings.

5. Configure the following control settings:

SNMP

Lets you define the following basic SNMP properties:

Bind Address

Specifies an interface to which the agent binds and listens for incoming SNMP requests. Valid addresses are IPv4 or IPv6 address.

Note: The corresponding default _port is specified during installation.

Bind Port

Specifies the trap port the agent binds to for sending SNMP traps. If no bind_address is specified, the agent binds to all available UDP addresses.

Default: Port selected by the system

IP Family

Specifies the agent communication method: IPv4 only, IPv6 only, or both. By default, the agent tries using IPv4 and then IPv6.

FIPS Mode

Specifies the agent to use FIPS-compliant encryption. Select Non-FIPS Mode to enable the CA eTrust Public Key Infrastructure libraries, and if this method fails, fall back to the internal minimum security solution. Select FIPS Co-existence Mode to enable FIPS-compliant encryption, and if this method fails, fall back to the CA eTrust Public Key Infrastructure Libraries. I 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 657).

   The control settings are defined.

9. Click Save Policy.

   The policy is saved.

More Information:

Configure Object Aggregation (see page 313)

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 366)

Define Traps and Communities

SNMP settings define the communities that the agent uses and the destinations to which it sends traps.

Follow these steps:

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

   The Available Policies page appears.

2. Select the policy in the Available Policies table.

   The Summary page for the policy appears.

3. Click the Traps and Communities tab.

   The Communities page appears.

4. Select one of the following and click Action, select Apply:

   - Include only Server specific SNMP settings
   - Include Server specific SNMP settings and all Default settings
   - Select 'Include Server specific SNMP settings and selected Default settings

The SNMP settings are updated and the community page in the Communities table displays the following:

Name

Specifies the name of the community string.

Port

Specifies the port of SNMP.

SNMP Version

Specifies the SNMP version that the community uses.

Access Rights

Specifies whether the community should have read write or read-only permissions.

Note: Add at least one read-only and one read write community.

Community/User

Specifies the community name.

Authentication Protocol

Specifies the protocol to authenticate SNMPv3 data.

Privacy Protocol

Specifies the protocol to authenticate SNMPv3 data.

Access Control List

Specifies a space separated list of IP addresses to restrict community usage to those addresses only. If you leave the list blank, the agent grants access to any system that uses the associated community name. Access lists are only for communities that use SNMPv1.

Note: For information about defining SNMPv2c and SNMPv3 access lists, see the *SystemEDGE User Guide*.

5. (Optional) Add, update, or delete other communities as necessary.
6. Click Save Policy.

   The policy is saved.
7. Click Trap Destinations.

   The Trap Destinations page appears.

8. Define a trap destination using the following controls and click Add:

Trap Type

Specifies the type of trap to send, depending on the SNMP version.

Destination

Specifies the IPv4 or IPv6 address to which to send traps.

Port

Specifies the UDP port to which to send traps.

Community

Specifies the community name sent with the traps.

Encoding

(Optional) Specifies how to include the source address you defined in the Trap Source field of the Control Settings pane in traps. This parameter is important if the trap source translates to an IPv6 address. Enter the encoding parameter in a three digit format XYZ, assuming leading zeros.

Default: 000

X

Controls extending the four byte IPv4 source address field (SNMPv1 traps only). Enter 0 to not extend the source address field to include the 16 byte IPv6 address, and enter 1 to extend the source address field.

Y,Z

Controls the inclusion of source information into the trap's varbind (Y) or UDP packet (Z; SNMPv1 traps only). Enter one of the following for these digits:

0: Do not modify the trap's varbind or the outer UDP packet.

1: Include the trap_source parameter as is in the varbind or packet (IPv4/IPv6 address or host name).

2: Include the trap_source parameter preferably as an IPv4 address (then IPv6 address, then host name).

3: Include the trap_source parameter preferably as an IPv6 address (then IPv4 address, then host name).

4: Include the trap_source parameter preferably as a host name (then IPv4, then IPv6).

5: Follow the preference for 2 and include the host name.

6: Follow the preference for 3 and include the host name.

7: Follow the preference for 1 and include the host name (if trap_source is an IPv6 address).

Trap Source

(Optional) Specifies the IPv4 or IPv6 address or the host name to use as trap source.

Default: Global Trap

The trap destination appears in the Defined Trap Destinations table.

9. (Optional) Add, update, or delete other trap destinations as necessary.
10. Click Save Policy.

    The policy is saved.

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

Create Templates for Server Workload

Create templates that are specific to the workload of a server. You can specify Monitors and MIB Extensions.

Follow these steps:

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

   The Template List page appears.
2. Click + (New) on the Template List toolbar.

   The New SystemEDGE Monitoring Template dialog appears.
3. Enter a name and an optional description for the template, the system type, and whether to base it on an existing template, and click Ok.

   The template is created, and the Summary page appears.
4. A template is a collection of monitors and MIB extensions. To add monitors to the template, see the section Add Monitors to the Template or the Policy (see page 662). To add MIB extensions to the template, see section Define MIB Extensions (see page 673).
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 662)
   - Create a Process Monitor (see page 665)
   - Create a Log File Monitor (see page 667)
   - Create a Windows Event Monitor (see page 668)
   - Create a History Monitor (see page 670)
   - Create a Process Group Monitor (see page 671)

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

   The monitor is loaded to the policy or the template.

Create a Threshold Monitor

Create a threshold monitor that lets the agent monitor the servers or the server groups against specified thresholds. The agent sends a trap when thresholds are breached.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

   The Template List page appears.

2. Select the template in the Template List.

   The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Threshold.

   The Threshold Monitors page appears.

5. Click + (New) on the Threshold Monitors toolbar.

   The Threshold Monitor Details: New dialog appears.

6. Configure the following threshold settings:

   Index

   Defines the table index that you want to use.

   Platform

   Specifies the platform.

   Description

   Defines an optional description.

   Object Class

   Specifies the object class to monitor. The values refer to the available MIB tables.

   Object Class Name

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

   Object Attribute

   Specifies the object attribute to monitor. The values refer to the available attributes of the table selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this threshold monitor.

   Object Attribute Name

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

   Object Instance

   Specifies the object instance to monitor. This value, for example, .3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this threshold monitor. For some object classes, the name of the instance itself can be given (for example, C: instead of .3, or /var for a Unix machine).

   Object Instance Name

   Defines the object instance name to use for the object state model. Value is an arbitrary string, for example, SysVol_C.

   Interval

   Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

Severity

Specifies the severity to use for the object state model.

Operator

Specifies the operator to use.

Value

Defines the value to use.

Sample Type

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

   The Threshold Monitor settings are saved.

8. Click Save Template.

   The Threshold Monitor is loaded to the template.

Create a Process Monitor

Create a process monitor that lets the agent monitor a process, service, or process table objects against specified thresholds. The agent sends a trap when thresholds are breached or the state of a process (running or stopped) changes.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

   The Template List page appears.

2. Select the template in the Template List.

   The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Process.

   The Process Monitors page appears.

5. Click + (New) on the Process Monitors toolbar.

   The Process Monitor Details: New dialog appears.

6. Configure the following process settings:

   Index

   Defines the table index that you want to use.

   Platform

   Specifies the platform.

   Description

   Defines an optional description.

   Object Class Name

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

   Object Attribute

   Specifies the object attribute to monitor. The values define the available attributes for process monitoring.

   Object Attribute Name

   Defines the object attribute name to use for the object state model. Value is an arbitrary string, for example, MemUsedPercent.

Object Instance

Specifies the object instance to monitor. This is the regular expression (dependent from optional settings) to use for matching processes by name, or Windows services by name. Pattern should uniquely match a single process (service). Arguments can be included (see optional settings).

Object Instance Name

Specifies the object instance name to use for the object state model. Value is an arbitrary string, for example, ApacheServer.

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

Severity

Specifies the severity to use for the object state model.

Operator

Specifies the operator to use.

Value

Defines the value to use.

Sample Type

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

   The Process Monitor settings are saved.

8. Click Save Template.

   The Process Monitor is loaded to the Policy.

Create a Log File Monitor

Create a log file monitor that lets the agent monitor any UTF-8 encoded system or application log file by searching for strings specified as regular expressions. The agent sends a trap when a match occurs.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

   The Template List page appears.

2. Select the template in the Template List.

   The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Log File.

   The Log File Monitors page appears.

5. Click + (New) on the Log File Monitors toolbar.

   The Log File Monitor Details: New dialog appears.

6. Configure the following process settings:

   Index

   Defines the table index that you want to use.

   Monitor Type

   Specifies the monitor type that you want to use.

   Platform

   Specifies the platform.

   Description

   Defines an optional description.

   Log File/Directory Name

   Defines the path to the file or the directory to monitor.

   Search Filter

   Specifies the search filter.

   Interval

   Defines the evaluation interval for the monitor in minutes

   Severity

   Specifies the significance of the monitor on a match.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings page lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

   The Log File Monitor settings are saved.

8. Click Save Template.

   The Log File Monitor is loaded to the Policy.

Create a Windows Event Monitor

Create a windows event monitor that lets the agent monitor the Windows event log entries using different filters (event source). The agent sends a trap when a match occurs.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

   The Template List page appears.

2. Select the template in the Template List.

   The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Windows Event.

   The Windows Event Monitors page appears.

5. Click + (New) on the 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 673). MIB Extensions within templates are supported for the purposes of applying the MIB Extensions directly to monitored systems. MIB Extensions for use within policies should be created directly in the Policy itself.

Add MIB Extensions to a Template or a Policy

Define MIB extensions for a template or policy using the policy configuration feature.

Follow these steps:

1. Click the Resources tab, open the Configure pane, and expand Monitoring Templates or Policies.
2. From the Templates List or Available Policies page, click the template or policy name.

   The Summary page appears.

3. Click the MIB Extensions tab.

   The MIB Extensions page appears.

4. Define the MIB extension attribute using the following controls and click Add:

   Index

   Defines the attribute leaf number.

   Type

   Specifies the attribute type.

Extension Command

Defines the full path or the name (including parameters) of the script or binary to execute.

Access Rights

Specifies the attributes access rights.

5. Click the Windows Performance tab.

   The Windows Performance pane appears.

6. Define the Windows Performance attributes using the following controls and click Add:

   Index

   Defines the attribute leaf number.

   Type

   Specifies the attribute type.

   Object

   Specifies the performance registry object.

   Counter

   Specifies the performance registry counter.

   Instance

   Defines the performance registry instance.

7. Click the Windows Registry tab.

   The Windows Registry pane appears.

8. Define Windows Registry attribute using the following controls and click Add:

   Index

   Defines the attribute leaf number.

   Type

   Specifies the attribute type.

   Key

   Defines the registry key in HKEY_LOCAL_MACHINE.

   Value

   Defines the attribute value.

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

9. Click Save Template or Policy.

   The configuration is saved.

Delete Monitors from Templates or a Policy

You can delete a monitor from a policy or template.

Follow these steps:

1. Click the Resources tab, open the Configure pane, and expand Monitoring Templates or Policies.
2. From the Templates List or Available Policies page, click the template or policy name.

   The Summary page appears.
3. Click the Monitors tab.

   The Summary page appears with a list of monitors managed by the policy.
4. Click the appropriate monitor tab and select the monitor you want to delete.
5. Click Action and select Delete.

   A warning message appears.
6. Click Ok to confirm the deletion.
7. (Optional) Repeat the process for any additional monitors.
8. Click Save Policy.

   The monitor is deleted from the Policy.

   Note: You cannot delete a template, or the policy with the template which is used by a server or a server group.

Apply Policy and Templates to Servers and Verify Settings

After you create the template, you can apply the policy, with the template directly to the servers or the server groups across the enterprise.

Follow these steps:

1. Select the policy in the Available Policies table or select a template from the Template List.

   The Summary page for the policy or template appears.
2. Select the Managed Machines tab.

   The list of managed machines appears.

3. Click Action and select Apply.

   Tabs appear for selecting systems on which to apply the policy. The 'Update machines running this policy/template' tab lets you apply the policy to systems that are already running the policy or template. The 'Apply to Machines not running this policy/templates' tab lets you apply policy or a template to systems.

4. (Optional) Do one of the following from the 'Update machines running this policy/template' tab:

   - Select 'Update all machines using this policy/template' to deploy the policy or template on all systems currently running it. This option is useful if you have made configuration changes that you want to apply globally.
   - Select 'Update selected groups of machines' to update only systems that meet any of the following criteria:
     - Running an out of date version of the policy or template
     - Exceptions have occurred
     - Running a current version of the policy

     Select any of these options. Exceptions occur when a user applies a point configuration change to an agent that is not represented in the applied policy or template.
   - Select Advanced and manually select systems in the Select Machines pane to which to reapply the policy or template.

5. (Optional) Select systems from the 'Apply to Machines not running this policy/template' tab to which to apply the policy or template.

6. Click Apply Policy or Apply Template.

   The application is initiated.

7. Verify if the servers behave as expected. If necessary, you can update and apply the updated policies and templates.

(Optional) Manage the Base Policy and Templates for One or More Servers

Manage the templates and the base policy for a single or multiple servers. You can replace the current base policy, add templates, or remove templates.

Follow these steps:

1. Click the Resources tab, open the Explore pane, and select the server for which you want to change the policy configuration.

   The resources page for the server appears.

2. Select Monitoring Software, Policies.

   The table displays the list of policies and templates applied to the server.

3. Click  (Modify Policy) to replace the current base policy for this server by another available base policy.

   The Modify Policy dialog appears listing all available base policies.

4. Select the appropriate policy and click Apply.

   The new base policy for the selected server has been applied. The status of the policy changes from Delivery Requested, Delivered, to Configured.

5. Click  (Modify Template) to add or remove templates from the configuration of the selected server.

   The Modify Templates dialog appears listing the available templates in the left pane and the applied templates in the right pane.

6. Select the templates that you want to add or remove, use the arrows to make your assignments, and click Apply.

   The new set of templates has been applied to the configuration. The status of the templates change from Delivery Requested, Delivered, to Configured.

   The new configuration has been applied.

You can also manage multiple servers as a group.

Follow these steps:

1. Create a service at the datacenter level that specifies the group of servers.

   The new service appears in the Explore pane.

2. Select the service.

   The service page appears.

3. Select Monitoring Software, Policies.

   The table displays the list of policies and templates applied to the servers.

   The following steps are identical to the procedure for single servers.

4. Complete the configuration.

(Optional) Update the Policy or Templates

If necessary, you can update the existing policy or templates by adding or deleting monitors from the policy or the template. The update procedures are similar to the creation process.

Follow these steps:

1. Add or delete monitors that are specific to the server workload. To add monitors to the template or the Policy, see Add Monitors to the Template or the Policy (see page 662). To delete monitors from the Policy, see Delete Monitors from Templates or a Policy (see page 675).
2. Define MIB Extensions (see page 673).
3. Define SystemEDGE Policy Control Settings (see page 652).

   The policy or templates are updated.

(Optional) Apply Policy and Template Updates to Servers and Verify Updates

After you update the template, apply the template updates directly to servers or server groups across the enterprise.

Follow these steps:

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

   The Summary page displays a list of the SystemEDGE Monitoring Templates.
2. Select the Template Name.

   The Summary page appears with the Template information.
3. Click Action and select Apply.

   Tabs appear for selecting machines on which to apply the monitoring template. The 'Update machines running this template' tab lets you apply the monitoring template to machines that are already using the template. The 'Apply to Machines not running this template' tab lets you apply the monitoring template to machines without using any template.
4. (Optional) Under Existing Machines, select machines from the 'Update machines running this template' tab options.
5. (Optional) Under Selected Machines, select the machines to which the template is re-applied.
6. (Optional) Select machines from the 'Apply to the Machines not running this Template' tab to apply the template.

7. Click Apply.

   The template application is initiated and the view Status link appears.

8. Click View Status link to verify whether the SystemEDGE monitoring template updates are applied to servers.

   The page appears with the list of servers to which the SystemEDGE monitoring template updates are applied.

   The Layered Template Updates have successfully been applied to the servers or the server groups.

9. Verify if the servers behave as expected. If necessary, you can update and apply the updated policies and templates again.

Appendix A: FIPS 140-2 Encryption

This section contains the following topics:

FIPS Overview

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 Server Automation uses the Advanced Encryption Standard (AES) adapted by the US government. CA Server Automation incorporates the RSA Crypto-J v3.5 and Crypto-C ME v2.0 cryptographic libraries, which have been validated as meeting the FIPS 140-2 Security Requirements for Cryptographic Modules.

Appendix B: Troubleshooting

This section contains the following topics:

CA Server Automation Troubleshooting

This section contains troubleshooting topics for CA Server Automation.

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

More information:

Unable to Connect to Microsoft SQL Server

Symptom:

Attempts to authenticate credentials to the Microsoft SQL Server 2005 and 2008 (Evaluation Edition) fail during product installation. The error message, Failed to establish connection to MSSQL displays.

Solution:

This issue occurs because TCP/IP is disabled by default on the Evaluation Edition. Enable TCP/IP.

OpenSSL Software Compatibility Issues

Symptom:

You install software that uses OpenSSL on the same system as CA Server Automation and then experience compatibility issues (particularly if the software installs library files in the System32 directory).

Solution:

Remove incompatible OpenSSL versions.

Note: Before you remove OpenSSL versions, verify that other applications do not use them.

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.

Log In with Different User Credentials on the Same Computer Using IE8

Symptom:

If you log in with two different user accounts in two different windows on the same computer, the login information of one user account overwrites the other. The product behaves inconsistently with the security attributes of one of the users. This inconsistency happens because the UI stores login information in the Internet Explorer 8 browser context, and the browser shares context with other browser instances by default.

Note: Earlier versions of Internet Explorer do not have this issue.

Solution:

To work around these issues that can occur as a result of this sharing, perform one of the following actions:

- Use different computers for different user logins.
- Start Internet Explorer 8 with the option *-nomerge*. Enter this command on a Windows command line or create or modify a shortcut:
  `C:\Program Files\Internet Explorer\iexplore.exe" -nomerge`

  This option tells Internet Explorer 8 to use a separate context for each window. Separate tabs within the same Internet Explorer 8 window always share context.

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.

DB Transaction Log Sizes Increase Unexpectedly

Symptom:

In data centers with numerous managed objects, configuration changes, and metrics data collection activities, the Management DB and Performance DB transaction logs can increase unexpectedly. This issue can cause disk space to become low in environments with limited resources.

Solution:

To resolve this issue, see the KB article on the Microsoft Support website http://support.microsoft.com/kb/873235

The transaction log files, aom2.ldf and dpm.ldf, are located in the directory C:\Program Files\Microsoft SQL Server\...\MSSQL\Data in default Microsoft SQL Server installations.

Note: If the database log file is reduced in size, restart the Apache service to improve performance.

Password Changes May Cause Authentication Errors

In some situations, changing a password for Active Directory, CA EEM, Microsoft SQL, and the system user causes issues with CA Server Automation.

Active Directory Password Expiration Causes Log in Issues

Symptom:

I cannot display the CA Server Automation user interface.

Solution:

If your CA Server Automation installation is configured to connect to Active Directory, the user who installs CA Server Automation is automatically registered with CA EEM. The registration of this user lets CA Server Automation authenticate users from the Active Directory domain. If the password for this user changes, users are no longer able to log in to the CA Server Automation user interface because CA EEM cannot authenticate them.

Resolve this issue by changing the password for the user who installed the product.

To change the user password

1. Click Start, Programs, CA, Embedded Entitlements Manager, EEM UI to log in to the CA EEM user interface.
2. Select Admin from the application drop-down list, leave the user name EiamAdmin, enter your password, and click Log In.
3. Click Configure, and then EEM Server.

4. Click Global Users/Global Groups in the left pane and leave the "Reference from an external directory" option selected.
5. Leave Microsoft Active Directory for Type, enter a new password in the Password and Confirm Password fields, and click Save.
6. Select Start, Program Files, CA, CA Server Automation, CA Server Automation Command Prompt. Run the following command from the CA Server Automation command prompt:

   ```
   dpmutil -set -sysuser
   ```

   You are prompted for the CA EEM user name and password.

Note: The Apache HTTP Server 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 with the dpmutil command can be used to log in to the CA EEM UI. Retry the dpmutil command using valid credentials.

CA EEM Password Change Causes Authentication Failure

Symptom:

When I start CA Server Automation after changing the CA EEM password, the services are not running.

Solution:

If you change the CA EEM administrator password (EiamAdmin), CA Server Automation does not start properly. All services appear to be down. Because the product stores the CA EEM credentials, change them in CA Server Automation using one of the following procedures:

If CA Server Automation is using native security

1. Run dpmutil –set –eiam and specify the new credentials.
2. Verify the system (sys_service) credentials in CA EEM. If they have been changed, run dpmutil –set –sysuser.
3. Recycle the CAAIPApache and CAIPTomcat services.

If CA Server Automation is using Active Directory

1. Run dpmutil –set –sysuser with either the same credentials as in step 1 or another AD user that has CA EEM admin rights.
2. Recycle the CAAIPApache and CAIPTomcat services.

Note: The Apache HTTP Server 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 with the dpmutil command can be used to log in to the CA EEM UI. Retry the dpmutil command using valid credentials.

SQL User Password Change Causes Blank UI

Symptom:

I cannot see the CA Server Automation user interface after changing the Microsoft SQL user password.

Solution:

If you are using Microsoft SQL Authentication and you change the password for the Microsoft SQL user (usually the *sa* password), the CA Server Automation UI is blank or displays Microsoft SQL error messages. Because the product stores the Microsoft SQL user credentials, change them using the following procedure.

To change the SQL user credentials

1. Select Start, Program Files, CA, CA Server Automation, CA Server Automation Command Prompt. Run the following command from the CA Server Automation command prompt:

   ```
   dpmutil -set –mgmtdb
   ```

   You are prompted for the database server, version, port, and the credentials used for accessing the main product tables.
2. If the performance database uses the same database server and SQL user as the management database, run the following command:

   ```
   dpmutil –set –perfdb
   ```

   You are prompted for the server name, administrator user name and password, database type, database instance, and database port.
3. Recycle the CAAIPApache and CAIPTomcat services.

System User Password Change Causes Blank UI

Symptom:

I cannot see the CA Server Automation user interface after changing the system user password.

Solution:

The user sys_service is created with the Native Security installation. If you change the password for this user, CA Server Automation displays a blank UI and all services do not run. Because the product stores the sys_service credentials, you must change them using the following procedure.

To change the system user credentials

1. Select Start, Program Files, CA, CA Server Automation, CA Server Automation Command Prompt. Run the following command from the CA Server Automation command prompt:

   ```
   dpmutil -set -sysuser
   ```

   You are prompted for the user name and password.

2. Recycle the CAAIPApache and CAIPTomcat services.

Discovering Large Networks

Symptom:

Discovery can fail to discover networks that contain more than 1024 nodes (for example, a Class B network with a subnet mask of 255.255.0.0).

Solution:

Run network discovery on smaller subnets (for example, a Class C network with a subnet mask of 255.255.255.0). If discovery fails on a large network, perform the following procedure to clean up the discovery database:

1. Stop the following services from Windows Services Control on the server where Network Discovery Gateway is installed:
   - Network Discovery Gateway
   - CA Server Automation Windows service
   - Network Discovery Gateway Agent
   - Network Discovery Gateway Server
   - Apache 2.2
2. Open the folder where Network Discovery Gateway is installed. The file is located in the following path:

   ```
   [CA Server Automation_installation_drive]:\Program Files\CA\SC\Network Discovery Gateway
   ```
3. Delete the *.sq3 files.
4. Start Network Discovery Gateway and the CA Server Automation Windows service in the following order:
   1. Network Discovery Gateway Agent
   2. Network Discovery Gateway Server
   3. Apache 2.2

Discovery Does Not Identify Operating System

Symptom:

A discovered system and operating system is classified as *other* instead of Windows.

Solution:

If a firewall is enabled, the operating system cannot be identified because Internet traffic is blocked. Turn the firewall off to classify the operating system as Windows.

VMs Not Being Discovered

Symptom:

VMware vCenter VMs are not being discovered.

Solution:

Verify that VMware Tools is installed on the VM in your VMware environment.

New System Name is not Displayed

Symptom:

When I rename a system, run discovery for the system and assign the system to a CCA server, the old name of the system is still displayed.

Solution:

For the new name to display, the DNS server has to refresh the table that maps IP addresses to system names. Wait until the DNS server refreshes the table and run discovery again.

Scheduled Jobs do not Run

Symptom:

Scheduled jobs do not run. When a scheduled job is run, the service controller and the initiation component must be active and accessible from the server that is running the job to help ensure the job runs. Symptoms of this issue include the following:

- Dashboard messages do not occur.
- Scheduled Jobs list shows the job status as Not Available.
- Your job does not run.

Solution:

Review the log file for the command line program that is running the job and look for the following entry: "Failed to get job id." Examples of log files are dpmpolicycli.log, dpmccmcli.log, and so on.

Note: For more information about log files, see the *Reference Guide*.

Local and Remote Monitors Do Not Show the Same Values

Symptom:

Local and remote monitors do not show the same values for the same attributes at the same time.

Solution:

For seamless local and remote monitoring, the monitored object names have been chosen identical. However, the APIs used are different and can return different values.

SystemEDGE on a remote machine runs independently from the RM AIM on the server, and the start point of their poll schedulers cannot be synchronized.

The monitored metrics are highly volatile. For example, two samples can differ with a high probability.

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 Server Automation user interface and log back in.

Accessing the CA Process Automation Server Requires Credentials Even After Configuration

Symptom:

You set the CA Process Automation EEM user name and password using the dpmutil -set -itpam-cfg-eem command. However, when accessing the CA Process Automation server from CA Server Automation, you are still prompted for credentials. This issue results from export regulations affecting the Java Cryptography Extension Policies of the JDK 6 environment.

Solution:

Resolve this issue by downloading jce_policy-6.zip from The Oracle Sun Developer Network (SDN) and applying the following JAR files to your installation:

- local_policy.jar
- US_export_policy.jar

Replace the files installed by CA Server Automation with the downloaded files.

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.

VM Reservation Fails: Could Not Find Computer UID for Software Delivery

Symptom:

After you provision a VM and then deploy the Software Delivery agent, the agent does not register back to the Software Delivery server with a unique computer UID. The computer UID is required to identify the computer system. An event message similar to the following displays in the UI dashboard:

```
A reservation task has failed. Reservation ID: n; System Name: host_name Task: 2
Software installation; Reason: "The status of a system preparation job has been
updated: Target computer = host_name, Description = DCRM request, Previous job status
= Scheduled, Current job status = Failed, Could not find computer UID for SD Agent
deployed on host_name".
```

Solution:

This failure has two possible causes:

1. When deployment of the Software Delivery agent takes longer than normal. To resolve this issue, increase the amount of time the Software Delivery service waits before failing the task. Increase the amount of time by increasing the number of times the computer UID lookup is attempted before failing the operation.
   a. Open the casdaconf.cfg file located on the CA ITCM or CA DSM domain manager where the Software Delivery service was installed.
   b. Locate the following configuration file setting, increase the value, and save the file:

      ```
      SD_DSM_Find_Computer_Retry_Count =3
      ```

   c. Restart Apache on the system where the Software Delivery service is installed for the change to take effect, and then retry the reservation operation.
2. An issue with the Common Application Framework (CAF). To resolve this issue, restart CAF using the following steps:
   a. Open a Command Prompt and type “caf stop”
   b. After the CAF services are stopped, type “caf start”
   c. After the CAF services are restarted, retry the reservation operation.

No Cisco UCS Manager in Explore Pane

Symptom:

The Cisco UCS PMM and AIM were configured, and the user interface had time to populate. However, the Cisco UCS Server does not appear in the CA Server Automation Explore pane with chassis, blade, interconnect, and organization information.

Solution:

To verify the UCS Manager name

1. Open the Cisco user interface.
2. On the Cisco Admin page, find the Cisco Java UI system name and verify whether that name is resolvable in the DNS. If the name is not resolvable, update the <drive>:\WINDOWS\system32\drivers\etc\hosts file with the correct name and IP address of the UCS Manager.
3. Reconfigure the UCS AIM with the correct UCS Manager system name.
4. Register the UCS Manager.
5. Register the UCS AIM.
6. Verify that the Cisco UCS Manager appears in the Explore pane.

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.

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 Folder Does Not Display in UI

Symptom:

After the product installation, the user interface does not display the vCenter Server folder.

Solution:

- Verify on the manager system if the Apache service is running. Start the Apache service if it is stopped.
- Verify on the SystemEDGE vCenter AIM system that the SystemEDGE service is running. Start the service if it is stopped.
- Verify that the vCenter AIM is configured correctly. You can verify the configuration from the Administration tab in the user interface or through NodeCfgUtil on the vCenter AIM system.
- Discover the server that is running the vCenter AIM from the CA Server Automation manager.

To discover the server from the user interface

1. Select the Resources tab, and then select the Management tab. The Discovery subtab is selected and Discovery type is set to System by default.
2. Enter the fully qualified domain name for the AIM system name or the IP address in the System Name field.
3. Click OK.

   After a short time, events relating to the specified system display in the Events windows on the Dashboard tab.

Resetting the vCenter Server Password Causes Data Collection to Fail

Symptom:

After resetting the VMware vCenter Server password for the user that CA Server Automation 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.

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

Duplicated Zone Entries in the Managed Folder

Symptom:

CA Server Automation has discovered multiple Solaris Zones hosts. In the Managed folder of the Explorer, zones with the same name appear.

Solution:

If zones with the same name belong to different Solaris Zones, "duplicated" entries can appear. The zones listed in the Managed folder are different objects with the same name. In the Solaris Zones folder, these zones appear under their hosts and can uniquely be identified.

Solaris Zones Folder

```
ZoneHost1
|-- ZoneA
ZoneHost2
|-- ZoneA
```

Managed Folder

```
Managed
|-- ...
|-- ZoneA
|-- ZoneA
|-- ...
```

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.

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.

Troubleshooting: vCenter Server Connection Failed

Symptom:

After I have added a new vCenter Server connection under Administration, Configuration, the validation of the connection to the vCenter Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used vCenter Server connection data (server name, user, password, protocol, port) is still valid. If necessary, update the connection data.
- Verify, if the vCenter Server system is running and accessible.
- Verify, if the VMware Management Service on the vCenter Server system is running properly.

To update the vCenter Server connection data

1. Click (Add) or (Edit) that is associated with the failed connection.

   The New or Edit vCenter Server dialog appears.

2. Add the valid server name, user, password, protocol, and port. Enable Managed Status and click OK.

   The connection data is updated.

3. Click (Validate) in the upper-right corner to validate the new settings.

   If the connection to the vCenter Server cannot be established, continue with the next procedure.

To verify, if the vCenter Server system is running and accessible

1. Open a command prompt on the CA Server Automation 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 Server Automation 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 Server Automation user interface, vCenter Server pane on the manager system and click  (Validate) in the upper-right corner.

   CA Server Automation 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.

Troubleshooting: vCenter AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows (Discovery in Progress).

Solution:

Wait until the discovery process of the vSphere environment has completed. The discovery duration depends on the amount of managed objects related to virtual and physical resources in vSphere. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job has completed, CA Server Automation adds a vCenter Server folder to the resources tree. Then you can start managing vSphere and its entire virtual infrastructure.

Troubleshooting: vCenter AIM Instance Status Icon Shows Multiple Instances

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows (Multiple AIMs manage this instance).

Solution:

Verify that your CA Server Automation manager manages each vCenter Server with one vCenter AIM instance only. If a CA Server Automation manager manages a vCenter Server through multiple AIM instances, management problems would occur. CA Server Automation stops monitoring the associated vCenter Server.

Decide which AIM instance you want to use to manage the vCenter Server and remove the other instances from the vCenter AIM Servers pane.

Follow these steps:

1. Select the AIM instance you want to delete and click (Delete).

   The Delete Item dialog appears.

2. Click Yes.

   Repeat these steps with other multiple instances until you have unique relationships between manager and AIM instance established.

Troubleshooting: vCenter AIM Instance Status Icon Shows Error

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the vCenter AIM:

- Verify, if the vCenter AIM Server is accessible.
- Verify, if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify, if the vCenter AIM server system is accessible

1. Open a command prompt on the CA Server Automation 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 hosts file on the CA Server Automation 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 Server Automation user interface, vCenter AIM Server pane on the manager system and click (Validate) in the upper-right corner.

   CA Server Automation validates the vCenter AIM Server connection.

   If the error status remains unchanged, verify whether the data you gathered according to the requirements for this scenario is still valid.

Troubleshooting: vCenter AIM Instance Status Icon Shows Warning

Symptom:

After I have added a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows (Warning). Other instances on the AIM are still in error state.

Solution:

No specific actions are required for the associated instance. The warning informs you that other instances on the same AIM are in error state. The warning disappears when the problems of the other instances are resolved. If one or more instances of an AIM are in error state, then all other instances in the AIM show warning state.

Troubleshooting: vCenter AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered vCenter AIM instances in the network, the status icons of several instances show  (Disabled). This vCenter AIM instance is not managed.

This status appears, if CA Server Automation 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 Server Automation 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 Server Automation manager.
- Edit the existing vCenter Server connection and change its managed status to enabled.

CA Configuration Automation Agent Stops During Installation

Symptom:

The CA Configuration Automation agent installation stops when I install it.

Solution:

On some systems, the Windows DEP option prevents javaw.exe from running. The CA Configuration Automation agent installation stops because it uses javaw.exe. To resolve this issue, follow these steps:

1. Open the Control Panel in Windows and double-click System.

   The System Properties dialog opens.

2. Click Advanced, and then click Settings in the Performance section.

   The Performance Options dialog opens.

3. Click Data Execution Prevention.
4. Select the Turn on DEP option for essential Windows programs and services only, and click OK.

Deleted OS Images from ITCM are not Deleted from CA Server Automation

Symptom:

After deleting all of the OS Images from ITCM, they no longer appear in ITCM, but still display in CA Server Automation.

Solution:

To resolve this issue, delete the OS images from the Software Package Library in addition to ITCM.

CA NSM Fails to Receive Traps on Windows

Symptom:

After a few days of receiving traps from CA Server Automation, the CA NSM Event Console stops displaying traps in the Event Console Log. The issue is in the Windows SNMP Trap Service. WinSNMP applications stop responding when you run third-party security scanning software program on a Windows Server 2003-based computer.

Solution:

To resolve this issue, see the KB article on the Microsoft Support website http://support.microsoft.com/kb/931565.

CA Service Desk Manager Exception Error

Symptom:

When CA Service Desk Manager is down, the connection status in the Administration, Configuration page displays the following message indicating that CA Service Desk Manager cannot be connected:

"ServiceDeskClientAdapter validateSDUser() exception: ; nested exception is: java.net.ConnectException: Connection refused: connect".

Solution:

To resolve this issue, restart CA Service Desk Manager.

ESX Job Status is Current But OS Installation Not Complete

Symptom:

I am using <asdm>, and the user interface displays my ESX/ESXi4.1 provisioning job status as "current." The boot image is deployed to the target system, and the OS image installation has not started.

Solution:

Wait until the OS image installation is finished, and the OS is ready for use.

Rapid Server Imaging (RSI) Troubleshooting

This section contains troubleshooting topics for Rapid Server Imaging.

RSI Server Errors

Symptom:

The RSI server reports the following errors when you provision to virtual hardware:

```
Failure - Could not communicate with server at xxx.xxx.xxx.193/8011: [Errno 104]
Connection reset by peer

Event watcher timer (45 seconds) expired waiting for ['FTShutdownHeartbeat']
```

Solution:

These are hypervisor performance issues. Review your hypervisor configuration settings to maximize performance. Then, retry provisioning the VM.

RSI and ITCM Servers

Symptom:

If both ITCM and RSI servers are installed, provisioning tasks (offline capture, capture driverset, deploy) can fail with the following error:

```
Check server for boot configuration errors. (Is network boot enabled?)
```

Also, if you are at the server console, you can see the CA-Unicenter ManagedPC Boot Server answering PXE boot requests rather than the RSI server.

Solution:

To configure the ACL setting on the ITCM server to ignore boot request for servers that are not registered with itself, follow these steps:

1. From Explorer, navigate to Control Panel, Configuration, Configuration Policy.
2. Right-click on Default Computer Policy, and select Un-Seal.
3. Navigate to DSM, Scalability Server, OSIM, ManagedPC, Server.
4. In the right panel, select "User Answer Control List" and change the Value to 2.
5. Navigate back to Default Computer Policy, and select Seal.

Capture or Deploy Fails

Symptom:

I get an error during an offline capture or deploy when using a newly registered network.

```
Network is Unreachable
```

Solution:

On the RSI server, register the depot to associate the newly registered networks (boot or external). See the *RSI Server Administration Guide* on DVD3.

RSI: Image is not Deployed Because of Timeout

Symptom:

A timeout occurs during an RSI image deployment and the deployment fails.

Solution:

Redeploy the image.

Invalid X Configuration When Provisioning Linux Images

Symptom:

When moving a Linux image from one server to another, the X configuration is invalidated and requires reconfiguration. After image deployment and during restart of the Linux server, a message displays that the X configuration is incorrect and prompts you to retry. For SUSE Linux Enterprise Server Linux distribution, the boot process is delayed until the X reconfiguration is completed. This delay results in a timeout failure when capturing or deploying images.

Solution:

To resolve this issue, respond Yes to the prompts to reconfigure the X configuration.

Linux and UNIX Provisioning Fails When Using Resizing Option

Symptom:

The *-scale* option in the *dpmrsi deploy image* command increases or decreases the file system sizes of the source image to fit on the disk of the target server. The RSI server calculates the target size of the file system by using the used space in each file system of the source disks to set the sizes of the new file system proportionally. This method can result in some file systems with insufficient space created on the new disk. On UNIX and Linux systems, this method results in a small /tmp file system. If /tmp is too small, provisioning fails.

Solution:

To resolve this issue, restart the provisioning operation without the *-scale* option and resize the file systems manually after provisioning is complete. You can clear the scale image option when deploying an image from the UI Actions drop-down menu.

Registered Server Fails

Symptom:

If a server is already registered in Racemi® DynaCenter®, registering the hypervisor fails with the following error:

`CAAM2228 RSI task RegisterHypervisor-73bef2acde failed; error Slot` *`server_id`* `is in use` *`RSI_server_name`*

Solution:

To resolve this issue, unregister the server_id in use, and then register the hypervisor.

RSI: Remote Server Discovery Fails

Symptom:

Occasionally, after a successful Rapid Server Deployment request, the automatic discovery of the target server fails. The following event is displayed in the imaging summary tab:

```
CAAM0515 The status of a discovery operation has been updated: Target machine = ,
Status=Progressing, Attempt nn
```

Solution:

The RSI agent running on the target server fails to communicate the IP address for discovery. To resolve this problem, recycle the RSI agent on the target server and manually issue the discover command for the target server.

RSI Agent Error During Image Capture or Deploy

Symptom:

When I submit a request for image capture or deployment, the following error occurs in the CA Server Automation UI:

```
Error: Image deployment failed: Error encountered when contacting RSI agent. Please
check if the agent is running.
```

Solution:

To resolve this issue, restart the RSI agent on the target server.

Solaris SPARC Provisioning Fails on Discovery of DVD

Symptom:

The Solaris SPARC agent image erroneously discovers a DVD as storage. A DVD inserted in the DVD drive causes the provisioning operation to fail.

Solution:

Remove the DVD from the drive and restart the provisioning operation.

RSI Imaging Fails with SSL Errors

Symptom:

My RSI imaging tasks are failing (capture, deploy, driver collection, and driver set capture) with SSL errors.

Solution:

Although there are several possible causes for this error, try the following actions to correct the issue:

1. Navigate to the CA Server Automation program files:

   C:\Program Files\CA\ServerAutomation\conf

2. Open the file, caimgconf.cfg.
3. Increase the value of the retry count parameter from 3 to 6:

   CONFIG_KEY_IMG_SSL_SOAP_ERROR_MAX_RETRY_COUNT.

4. Restart the Apache service.

Target Server Stops Responding During Reboot After Image Capture or Deployment

Symptom:

The target server running an agent image stops responding during a reboot after image capture or deployment.

Solution:

The RSI server releases the agent image assigned to it before the server has completed shut down because of a timing issue. To avoid this issue, assign individual agent images for each target server OS type before starting image capture and deployment.

1. Run the following commands to create agent images for all the supported OS types:

   ```
   # . /opt/race/share/conf/buildout.conf;
   /opt/race/share/conf/provisionmgr.sh
   ```

   Sample Output:

   ```
   Sourcing confData

   Server class ProvisionMgr-Linux2.6 exists

   Server class ProvisionMgr-CentOS4-x86_64 exists

   Server class ProvisionMgr-CentOS5-i686 exists

   Server class ProvisionMgr-CentOS5-x86_64 exists

   Server class ProvisionMgr-Solaris10-sun4u exists

   Server class ProvisionMgr-Solaris9-sun4u exists

   Server class ProvisionMgr-Solaris8-sun4u exists

   Server class ProvisionMgr-Solaris10-i86pc exists
   ```

2. Run the following command on the RSI server to assign an agent image to each target server:

   ```
   # dccmd assign agent <server_id> <ostype>
   ```

Windows Driver Collection Issue with IBM Servers

Symptom:

When provisioning a Windows image to an IBM server using drivers collected from the IBM Server Guide, exclamation points or question marks appear on some noncritical devices in the Windows Device Manager. To locate the Windows Device Manager, click Start, All Programs, Administrative Tools, Computer Management, Device Manager. Additionally, the Windows driver collection does not produce the desired results.

Solution:

To resolve this issue, collect the IBM drivers directly from the server.

Windows Provisioning Fails

Symptom:

Occasionally, Windows provisioning fails even though a sufficient amount of disk space exists on the target server. The following error is displayed:

```
Could not populate from ntfsclone image _.ntfsclone.gz
```

Solution:

For deployment operations, the disk space on the target server must be equal to or greater than the disk space on the source server. To resolve this issue, enable the scale image option when you deploy an image in the UI from the Actions drop-down menu. You can also use the *dpmrsi deploy image* command with the *-scale* option to run the provisioning operation again. **Note:** For more information about dpmrsi commands, see the Rapid Server Imaging Commands section in the *Reference Guide*.

Reservation Manager Troubleshooting

This section contains troubleshooting topics for Reservation Manager.

More information:

Amazon Machine Images Are Not Available for Selection

Symptom:

When an administrator starts the Add AMI Image wizard, no AMI instances are available for selection. The connection status to the EC2 server is fine.

Solution:

By default, only AMIs that are owned by the owner specified in the EC2 server connection information are available for adding to the inventory. You can change the default so that public AMIs published by Amazon or other users can be made available for use. To allow public AMIs to be added to the inventory, run dpmutil to reconfigure Amazon EC2.

Chargeback Calculations Are Lower or Higher Than Reservation Amounts

Symptom:

If a user adds resources for an existing reservation, they are charged for the full amount of resources for the entire 24 hour period. If a user adds resources to an existing reservation, and then returns those resources before the end of the day, there is no extra charge.

Solution:

By default, chargeback costs are calculated once daily at midnight. To charge for usage of less than a day, increase the value of the Chargeback Calculation Frequency configuration setting. For more information, see Configure Chargeback (see page 471).

Installation Target Cannot be Resolved

Symptom:

A software installation task fails and an event in the Reservation Events table indicates that the “SD Agent installation job” failed. The reason listed is “Target cannot be resolved”.

This message indicates that the CA Software Delivery application was unable to access the target system using the name it was given. This problem can occur if the DNS has not been updated with the name of the newly provisioned virtual machine.

Solution:

If this issue occurred due to a delay in updating the DNS, restart the software installation task from the Reservation Task page.

If the task fails again, log on to the CA DSM server and ping the target system by name. If it is not reachable, investigate why the Software Delivery server DNS is not being updated with the name of the target VM.

Log In with Different User Credentials on the Same Computer Using IE8

Symptom:

If you log in with two different user accounts in two different windows on the same computer, the login information of one user account overwrites the other. The product behaves inconsistently with the security attributes of one of the users. This inconsistency happens because the UI stores login information in the Internet Explorer 8 browser context, and the browser shares context with other browser instances by default.

Note: Earlier versions of Internet Explorer do not have this issue.

Solution:

To work around these issues that can occur as a result of this sharing, perform one of the following actions:

- Use different computers for different user logins.
- Start Internet Explorer 8 with the option *-nomerge*. Enter this command on a Windows command line or create or modify a shortcut:

  ```
  C:\Program Files\Internet Explorer\iexplore.exe" -nomerge
  ```

  This option tells Internet Explorer 8 to use a separate context for each window. Separate tabs within the same Internet Explorer 8 window always share context.

No Resources Available Message When Requesting a VM

Symptom:

After selecting a virtual machine template, a user sees the warning message "No resources are currently available for your selection." The user cannot continue to the next step in the reservation wizard. This message is typically displayed when an end user has access to the virtual machine template but is not authorized to create virtual machines on an ESX server or cluster in the same data center as the template.

Solution:

To resolve this issue, perform the following steps:

1. Identify the data center where the selected virtual template resides. This data center is displayed in the Location column of the System Images table.
2. Identify the organizational unit to which this user belongs.
3. Identify the virtual resource pools that are defined in the same data center as the selected virtual machine template. Provide access to the organizational unit the user belongs to.
4. Alternatively, create a new virtual resource pool that contains one or more ESX servers or clusters that are located in the same data center. Provide access to the organizational unit.

Password Change Causes Error Message

Symptom:

After I change my password, I see the following message in CA Server Automation:

```
The security token could not be authenticated or authorized.
```

Symptom:

If you logged in to CA Server Automation before the password change, you may see this message. Log out, and then log back in.

Tier Label Changes on VMware Datastore

Symptom:

When I assign a tier label to a datastore, the tier label changes for other datastores with the same name.

Solution:

When setting up VMware datastores in Reservation Manager, you can assign a tier label to each datastore. If you have different datastores that have the same name, the tier label is applied to all and cannot be changed. The only workaround is to rename the datastores with unique names.

Unable to Find Package Entries for Personality AutoDeploy

Symptom:

A software installation task fails with the error “Could not find any package entries for personality AutoDeploy”.

This message indicates an issue with the definition of software that is to be automatically deployed to the system being provisioned. The list of software that is to be automatically deployed is defined in the casdaconf.cfg (see page 388) file which is located on the server where the Packaging component was installed.

Solution:

Verify that at least one software package is configured to be installed for each operating environment and that the AUTODEPLOY definitions for that operating environment are sequentially numbered.

More information:

Software Delivery Configuration File (see page 388)

Unable to Retrieve Information from vCenter

Symptom:

A warning message is displayed to the end user that reads “No resources are currently available for your selection”. The message also includes the following information “Unable to retrieve information from Virtual Center for *template_name*”.

This message is displayed when a user has been granted access to the virtual machine template selected but the template is not available in VMware vCenter. This situation can occur if the template has been deleted or orphaned since adding it to the Reservation Manager inventory.

Solution:

To resolve this issue, perform *one* of the following steps:

1. If the template has been deleted from your VMware vCenter server, remove all access to the Reservation Manager inventory item associated with this template.
2. If the template has been orphaned, use the VMware Infrastructure Client to correct the problem.
3. If the template was renamed, either rename it back to the original name or remove all access to the Reservation Manager inventory item.

If none of the preceding steps resolve the issue, check if the CA Server Automation connection to the vCenter server is down. Log in (see page 25) to the CA Server Automation user interface and check the vCenter Server connection status that is displayed under the Administration, Configuration tabs.

VM Resources are not Available for Dates Requested

Symptom:

A warning message is displayed to the end user that reads “Unable to fulfill request: Virtual machine resources are not available for the dates requested”. The end user cannot continue to the next step in the reservation wizard.

This message can be displayed for the following reasons:

- The user has already requested the maximum amount of virtual machines that they are allowed to reserve for the time period specified.
- Reservation Manager has determined that none of the ESX servers or clusters that are available to this user have free capacity to accommodate the request for the reservation period.

Solution:

Use one of the following solutions:

1. The user can change the reservation period to a time when resources are available.
2. If the user was prevented from reserving a new virtual machine due to the limit on the number of VMs they are allowed to reserve, increase the limit by modifying the Maximum Systems setting defined for one or more virtual resource pools to which the user has access.
3. If the user was prevented from reserving a new VM because the ESX server capacity limit is too low, modify the associated resource pool. Select Allow memory overcommitment on the Properties tab, and enter a percentage. For more information, see Set Overcommitment of Memory on ESX Server or Cluster (see page 506).

VM Reservation Fails Because of CPU Limitation

Symptom:

A reservation fails when more virtual CPUs are requested than VMware allows for the ESX server chosen. Messages like the following are displayed:

```
Virtual machine requires X CPUs to operate, but the host hardware only provides X.

This virtual machine is configured with an unsupported number of virtual CPUs.
```

Solution:

You can reduce the number of virtual CPUs for the ESX servers in your environment by changing a configuration setting. See the field *Virtual CPU Limit* in the topic Set Limits on Virtual Machine Resources (see page 492).

VM Reservation Fails in a Clustered Environment

Symptom:

A reservation fails in a clustered environment. Messages such as the following are displayed:

```
Reason: Datastores are not configured correctly in virtual resource pool Name.
Resolution:  Please make sure to specify datastores that can be used to deploy to ESX server Server_name.

Details: Exception: An attempt to deploy a virtual machine to Server_name cannot be performed as no datastores are available for use.
Please edit the virtual resource pool Name to specify datastores that can be used when deploying virtual machines to Server_name.
```

Solution:

Use one of the following solutions:

- Each ESX server needs a datastore defined for use when creating new virtual machines. Update the resource pool to associate the listed ESX server with a datastore.
- Verify that the VMware vCenter cluster name does not contain a slash (/) character.

Glossary

AIM

See *application insight module*.

AIP

See *automation integration platform*.

Amazon Elastic Compute Cloud (EC2)

The *Amazon Elastic Compute Cloud (EC2)* provides data center services from Amazon.com for developers. For more information, visit http://docs.amazonwebservices.com.

AOM

See *automation object model (AOM)*.

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.

AppLogic appliance

An *AppLogic appliance* (virtual appliance) is a self-contained virtual object that provides a particular function inside an application.

AppLogic application

An *AppLogic application* is a single system object that includes everything necessary to run a specific distributed application. The application includes the application code, HTML pages, templates and scripts, databases and content, and also operating systems, middleware, file storage, load balancers, firewalls, and all configuration information needed to reconstruct and run the application on an AppLogic grid. In addition, each application carries a defined resource budget - a minimum set of hardware resources including CPU, memory, and bandwidth required to run the application, and the maximum resource quota allowed for it. Applications are assembled using self-contained software components called virtual appliances.

AppLogic assembly

An AppLogic *assembly* is a packaged structure of interconnected appliances that can be manufactured on demand and used in exactly the same way as one would use an appliance. The assembly support is hierarchical, allowing users to create assemblies of assemblies of assemblies and so on.

AppLogic component

An AppLogic component is an instance of an AppLogic appliance or composed of AppLogic assemblies. An AppLogic component can have multiple instances of appliances or can contain only one instance of an appliance.

AppLogic grid

An *AppLogic grid* is a system which runs on a server array assembled from commodity servers connected using a gigabit ethernet switch. Some (or all) of the servers are expected to have directly attached storage - inexpensive IDE/ATA/SATA hard drives which AppLogic uses to provide a distributed storage pool for applications. Grids are a collection of servers which host AppLogic applications.

AppLogic template

An *AppLogic template* is a model application used to provision multiple instances of the application using different resource settings and quotas for each instance.

automation integration platform (AIP)

The a*utomation integration platform* is a management platform based on Web Services and ActiveMQ.

automation object model (AOM)

An *automation object model* is a database that stores managed entities. It is based on a CIM schema, which is a model for describing management data. See also *common information model (CIM)*.

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 Server Automation 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.

CA AppLogic

CA AppLogic is a cloud computing platform for composing, running, and scaling distributed applications.

chassis (UCS)

Hardware frame that holds Cisco UCS switches and blades.

CIM

See *common information model (CIM)*.

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 Server Automation 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.

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.

common information model (CIM)

A *common information model (CIM)* provides schemas for databases that store information about such things as systems, networks, and devices. A CIM implementation lets different management applications collect data from a variety of sources.

container (Solaris)

A Solaris *Container* provides complete runtime environments for applications. Resource management and Solaris Zones are parts of a container.

datacenter (VMware)

A *datacenter* serves as a container for your hosts, virtual machines, resource pools, or clusters. If their virtual configurations meet the requirements of specific departments, datacenters can represent organizational structures, such as geographical regions or separate business functions. You can also use datacenters to create isolated virtual environments for testing or to organize your environment.

datastore (VMware)

A *datastore* specifies a virtual representation of combinations of underlying physical storage resources in a datacenter. These physical storage resources can be provided by local disks on a server, by SAN disk arrays, and so on.

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.

EC2

See *Amazon Elastic Compute Cloud (EC2)*.

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.

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.

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.

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).

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.

organizational unit

An *organizational unit* (org unit) is a group of users. Org units provide security by giving users access to objects like resource pools, system images, and templates.

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.

provisioning

After discovery, *provisioning* can find a physical or virtual machine, add an operating system and image, and make it available for use. When you require specific machine characteristics, the product can provision a machine to meet your needs.

remote deployment

Remote deployment provides the ability to remotely deploy and configure monitoring agents to multiple systems in one operation throughout your enterprise.

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.

security group (Amazon EC2)

A *security group* is a term that Amazon uses to describe the IP filtering to a running instance. For more information, visit http://docs.amazonwebservices.com.

service-oriented architecture (SOA)

Service-oriented architecture (SOA) is a way of programming that creates small services instead of large applications for various business functions. These services provide flexibility and reusability because different departments can share common functions.

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.

snapshot

A *snapshot* is a record of a virtual machine at a certain point in time. Snapshots let Reservation Manager users restore VMs to their previous state without contacting an administrator. Snapshots are useful in development and testing environments. Because administrators can control whether taking snapshots is allowed, snapshots may not be available at all sites.

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.

SOA

See *service-oriented architecture*.

storage tiers

See *tiers.*

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.

tiers

In Reservation Manager, *storage tiers* are classifications for the data stores associated with each disk. Tiers generally indicate different levels of performance of the data store on which a VM and its hard drives are created.

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).

vCenter Server (VMware)

VMware *vCenter Server* provides the central point of control for configuring, provisioning, and managing a virtual vSphere environment. vCenter Server runs as a service on Microsoft Windows Servers and Linux Servers.

vCenter Server Agent (VMware)

The VMware *vCenter Server Agent* connects ESX Servers with a vCenter Server.

vCenter Server Database (VMware)

The VMware *vCenter Server Database* stores persistent information about the physical servers, resource pools, datacenters, and virtual machines managed by the VirtualCenter.

virtual disk (VMware)

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

virtual LAN or VLAN

See *virtual local area network*.

virtual local area network

A group of hosts that communicate like hosts attached to the same broadcast domain, even if they are not in the same physical location. You can group end stations on a virtual local area network (VLAN) regardless of whether they are on the same network switch. You can configure VLAN connections using software instead of physically relocating devices.

virtual machine, VM (VMware)

A *virtual machine (VM)* is a software-based computer that runs an operating system and applications like a physical computer. A virtual machine consumes resources dynamically on its physical host, depending on its workload. Because virtual machines are flexible computing units, their deployment comprises a wide range of environments like datacenters, clusters, cloud computing, test environments, desktops, or laptops. Their primary strength lies in datacenters, where they are used for server consolidation, workload optimization, and energy efficiency.

virtual NIC (VMware)

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

virtual switch (VMware)

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

vNetwork Distributed Switch, vDS (VMware)

A *VMware vNetwork Distributed Switch* abstracts the configuration of virtual switches from the host to the datacenter level. A vNetwork Distributed Switch operates as a single virtual switch that spans across all hosts in a datacenter which are associated with that switch. vNetwork Distributed Switches consist of distributed port groups which are similarly configured to port groups on standard switches, but extend across multiple hosts. These properties allow virtual machines to maintain a consistent network configuration as they migrate among multiple hosts.

Like a vNetwork Standard Switch, each vNetwork Distributed Switch is a network hub that VMs can use. A vNetwork Distributed Switch can forward traffic internally between VMs or link to an external network by connecting to physical NICs (uplink adapters). For further details, see the vNetwork Distributed Switches documentation at http://pubs.vmware.com.

CA Server Automation 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 Server Automation 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.

XML-RPC

Allows software that runs on different operating systems or in different environments to make procedure calls over the internet. *XML-RPC* uses HTTP as the transport protocol and XML for encoding.

Index

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X