

CA Server Automation

Administration Guide

Release 12.7



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

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

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

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

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

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

The manufacturer of this Documentation is CA.

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

Copyright © 2012 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

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 ITAM)
- 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 Software Delivery, a component of CA IT Client Manager
- CA Spectrum® Infrastructure Manager (CA Spectrum)
- CA SystemEDGE
- Racemi DynaCenter®

Contact CA Technologies

Contact CA Support

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

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

Providing Feedback About Product Documentation

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

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

Contents

Chapter 1: Introduction 17

Related Publications	17
Conventions	18

Chapter 2: Overview 21

Architecture	21
Databases	25
Management DB	25
Performance DB	26
User Interface.....	26
Access the User Interface.....	27

Chapter 3: Managing Users and User Groups 29

User Access Control.....	29
Active Directory.....	29
Native Security	30
Password Management.....	33
Change the CA EEM Administrator Password (EiamAdmin)	33
Change the Database Administrator (sa) Password	34
Change the System User Password for Native Security	35
Change the System User Password for Active Directory Security.....	36
User Group Management	37
Search for Users or User Groups.....	37
Create a User Group.....	38
Assign Users to Groups	39
Assign External Directory User Groups to User Groups	39
Set User Group Privileges.....	41
Set User Group Permissions.....	41
Set User Group Permissions for Services	42
Set Run Command Script Privileges	42
Import External Directories.....	43
Delete User Groups.....	43
Assign User Groups Access Rights to Services	44
Remove Users or User Groups from a User Group	44

Chapter 4: Managing Systems Performance 47

Systems Management	47
Discovery	49
Discover a System	49
Delete a System	50
Discover a Network	51
Enhanced Discovery and SNMP Information	52
Cancel Network Discovery	53
Rediscover a Network	54
Delete a Network	54
Services	54
Create a Service	55
Edit a Service	56
Remove Server from Services	57
Delete Services	58
Managed and Unmanaged Resources	58
Unmanage Managed Resources	59
Manage Unmanaged Resources	59
Delete Managed Resources	60
SystemEDGE Features	60
Systems Management MIB	62
State Management Model	64
Stateless Monitoring	65
Managed Mode and Unmanaged Mode	66
Application Insight Modules (AIMs)	66
Agent Configuration	68
Monitoring Software Settings	70
Security and Maintenance	70
Enable Maintenance Mode	71
Service Response Monitoring	72
SRM Tests	73
Agent Visualization	75
View SystemEDGE Monitors	75
View Managed Object States	76
View SRM Tests	77

Chapter 5: Managing SystemEDGE and Application Insight Modules (AIMs) 79

User Permissions and Access Requirements Reference	79
Cisco UCS	80
Citrix XenServer	81
Hyper-V	82

IBM High Availability Cluster Multiprocessing	82
IBM PowerVM	83
Microsoft Cluster Server	84
Oracle Solaris Zones	84
Agent Remote Deployment.....	85
Remote Monitoring.....	86
SystemEDGE and Advanced Encryption	87
VMware vCenter	87
VMware vCloud	88
How to Configure SNMP and Access Control Lists	88
SNMP Consistency.....	88
Global and Server-level SNMP Settings.....	89
How to Configure SNMPv1/v2 Settings and Access Control Lists	91
How to Manage Server-level SNMP Settings	102
How to Configure SNMPv3.....	106
How to Deploy SystemEDGE and AIMS	112
Overview	112
Configuration	114
Scalability	117
Deployment Packages	119
Using Remote Deployment	132
Specific Remote Deployment Use Cases.....	142
Deployment Jobs.....	149
Infrastructure Deployment Process	150
How to Configure SystemEDGE and Service Response Monitor Through Policies and Templates	157
Configuration Overview	158
How to Apply Policy and Layered Templates to Servers	161
How to Create and Apply an Autowatcher to a System.....	195
How to Monitor User-specific Metrics (MIB Extensions)	202
How to Monitor a Specific Windows Performance Registry Metric	204
How to Create SRM Policy.....	207
Discovering the Agents	208
Common Usage of Policy Configuration Functions.....	208
How to Change the Configuration Mode for SystemEDGE	263
Review Requirements	264
Review Managed Mode and Unmanaged Mode Details	264
Verify the Current Configuration Mode of SystemEDGE	265
How to Change SystemEDGE from Managed Mode to Unmanaged Mode	267
How to Change SystemEDGE from Unmanaged Mode to Managed Mode	270
Verify the SystemEDGE Configuration Mode.....	272

Chapter 6: Managing Virtual Environments

275

VMware vSphere and vCenter Server	275
Monitored vSphere and vCenter Server Resources	276
How to Configure the vCenter Server Management Components	278
Device Management for VMs	292
Fault Tolerance for Virtual Machines	294
Hot-plug Support for VMs	299
Logical Volumes in Virtual Machines	301
Resource Allocation	301
How to Use Policy Actions to Identify Performance Issues	305
vApp Support	307
vCenter Server in a Cluster	313
Virtual Standard Switches and Virtual Distributed Switches in the vNetwork Panel	313
VMware vCenter Provisioning and Common Use Cases	320
VMware vCloud	335
How to Configure the vCloud Director Management Components	337
Remote and Multi-instance vCloud Director Support	350
vCloud Folder Structure	350
vApp Support in vCloud	350
vCloud Organizations	352
IBM PowerVM (LPAR)	354
How to Configure the PowerVM Management Components	355
calpara.xml File	369
IBM PowerVM Server Administration Overview	374
LPAR Monitoring	376
Add a Logical Partition for an IBM AIX Computer	377
IBM PowerVM Management	380
Citrix XenServer	388
How to Configure XenServer Management Components	390
How to Prepare Windows Templates for XenServer Provisioning	399
Manage VM Status (XenServer)	404
Provision a Citrix XenServer Virtual Machine	405
Solaris Zones	406
How to Configure the Solaris Zones Management Components	407
Solaris Zones Management	418
Microsoft Hyper-V Server	424
How to Configure Hyper-V Management	426
Hyper-V Management	438
Cisco UCS	446
Interaction Between Cisco UCS Management Components	447
Add a Cisco UCS Manager Server	448

Register a UCS AIM Server	449
How to Configure the Cisco UCS Management Components	450
Cisco UCS Management	461

Chapter 7: Configuring Resources 475

Add Proxy Servers	475
Cisco UCS Server	476
Configure the Cisco UCS AIM from the Command Line	477
AppLogic Management	478
AppLogic Prerequisites	479
Add AppLogic Servers	479
Discover AppLogic Grids	480
Provision AppLogic Applications	481
Manage AppLogic Resources	482
How to Configure AIX NIM Imaging	484
Install NIM Adapter on AIX NIM Server	485
Edit the ca_post_install.sh script File	485
Update the Hashed Password Variable	486
Increase the size of the /tmp and /opt filesystems	486
Start or Stop the NIM Adapter Daemon	487
Configure NIM Master Server	488
Synchronize NIM Master Servers	488
Dynamic NIM Machine Resource Support	489
Solaris JumpStart Provisioning	489
Overview	490
JumpStart Prerequisites	490
JumpStart Adapter Installation	491
JumpStart for Solaris	493
How To Create a Solaris 8 Image	495
Configure CA Network Automation	507
Configure the CA Network Automation Server	507
Storage Provisioning Features	508
Storage Prerequisites	509
How to Configure Storage	510
How to Configure Software Delivery	511
Automating Processes with CA Process Automation	512
CA Process Automation Prerequisites	513
Configure CA Process Automation for Single Sign-On	513
Access the CA Process Automation User Interface	515
Configure a CA Process Automation Process	516
Event Forwarding	518

Configure Windows for SNMP	518
Configure SNMPv1 Traps by Editing the sysedge.cf File	519
Configure CA Server Automation to Forward Events.....	521
SNMP V3 Engine ID	521
Configure SNMP Management Servers.....	522

Chapter 8: Monitoring Clusters 523

Microsoft Cluster Service	523
How to Configure Microsoft Cluster Service Management Components	524
Register a Cluster	535
Remove a Cluster	535
Modify Cluster Properties	536
Microsoft Cluster Service Management.....	536
IBM High Availability Cluster Multiprocessing	537
Interaction Between IBM HACMP Management Components	538
Configure SSH.....	539
Configure HACMP AIM with NodeCfgUtil in Dialog Mode	539
Configure HACMP AIM with NodeCfgUtil in Command Mode	540
CA IBM HACMP SystemEDGE AIM Traps.....	541

Chapter 9: Agent-less Monitoring 543

Remote Monitoring.....	543
Interaction Between Remote Monitoring Components	544
Advantages of Remote Monitoring.....	545
Features and Benefits	545
Architecture	547
Use Case Scenario	549
Configuration Prerequisites	550
Configuring Remote Monitor Systems	551
Managing Systems Using Remote Monitoring.....	554

Chapter 10: Using Rules and Actions 561

Rules and Actions	561
Configure CA SDM	562
Configure the CA SDM Ticket Status Setting	563
Rule Planning.....	564
Create a Rule	564
Use a Predefined Action Type	567
Create a Custom Action	649
Define an Action Sequence	650

Define a Schedule.....	651
Create Automation Policy	653
Use Cases for Policies	653
Use Case: Adding a Server to a Service	653
Use Case: Adding a New Rule to a Service	654
Use Case: Defining an Action	654
Configuring Data Collection	655
Key Points About Metrics Collection	656
Configure Data Collection for a Data Center.....	658
Configure Data Collection for a Server.....	659
Configure Data Collection for a Virtual Resource	661
Configure Performance Thresholds	663
Configure the Metric Filter.....	664

Chapter 11: Provisioning Resources 667

Imaging Services.....	667
Service Provisioning	668
How to Provision Services	668
How to Deploy a Wiki Web Page	690
How to Deploy Oracle WebLogic Server	704
CA Software Delivery.....	711
Understanding Packaging.....	712
Software Delivery Configuration File	714
Changing Agent Versions	718
Amazon EC2 Provisioning.....	719
Supported Features.....	720
Prerequisites	721
How to Configure and Provision Amazon EC2 Resources	722
Bare Metal Provisioning to a Cisco UCS Blade	728
LPAR Provisioning for IBM AIX	728
IBM AIX Provisioning with NIM	729
Prerequisites	729
Add an IBM AIX Client System Using a Resource Group	730
Add an IBM AIX System Using an Individual Resource.....	732
Using the MKSYB Utility	733
Rapid Server Imaging	740
RSI Prerequisites	741
Register the RSI Server.....	743
Capture RSI Images	744
Deploy RSI Images.....	746
Migrate Servers with RSI	748

Manage RSI Images and Drivers.....	750
Automate RSI Actions.....	751
How to Deploy Rapid Server Imaging Using CA ITCM	751
How to Set Up Rapid Server Imaging for AppLogic	759
How to Backup and Restore with RSI.....	761
How to Perform Disaster Recovery with RSI	762
Rapid Server Imaging Support for UCS.....	764

Chapter 12: Setting Up Reservation Manager 765

Reservation Manager Prerequisites	765
Prepare Your Environment for Reservation Manager	766
Prepare CA Server Automation for Reservation Manager	768
Setup and Configuration	768
Predefined Content and Configuration for Service Provisioning	769
Make Virtual Machines Available to Users	769
Make Hyper-V Virtual Machines Available to Users	777
Make Amazon Machine Images Available from EC2.....	777
Enable AppLogic Application Reservation.....	783
Logical Partitions	783
Configure IBM PowerVM Logical Partitions	787
Make Physical Systems Available to Users.....	787
Make Services Available to Users.....	792
Public Templates for End Users	793
Use Static IP Addresses	794
Configure Email Notifications.....	797
Configure Announcements	798
Let Users Perform Some Administrative Tasks	799
User Access to Reserved Systems	800
User Management.....	803
Organizational Units.....	803
Multi-Tenancy Environment	805
VLAN Scoping	809
Administration.....	809
Access the Reservation Manager User Interface	810
Approve or Reject Reservation Requests	811
Extending Reservations.....	811
Resource Allocation and Forecast Charts.....	812
Run Frequently Used Reports	813
Suspend and Restart the Scheduling of Tasks	814
Suspension and Restart of Individual Tasks	815
Chargeback.....	816

Configure Chargeback	818
Configure Chargeback by Resource	818
Configure Chargeback for Storage Tiers.....	820
Configure Chargeback by Tier for Amazon EC2.....	821
Configure Chargeback by Tier for IBM PowerVM Logical Partitions.....	821
Select a Chargeback Tier for IBM PowerVM Logical Partitions	822
Configure Chargeback by Resource for AppLogic Applications.....	823
Configure Chargeback Display.....	823
Customization	824
Configure the Contact Hyperlink.....	824
Use the Reservation Manager Mobile App	824
Configure Online Help	825
Customize the Home Page	826
Email Customization.....	827
Enable or Disable Inheritance of Resources from the Public Org Unit	833
Enter Home Page Welcome Text	834
Specify a Timeout Value.....	834
Set Limits on Virtual Machine Resources.....	835
Set Over Commitment of Memory on ESX Server or Cluster	835
Specify a Folder for VMware Virtual Machines.....	836
Specify the Maximum Number of NICs per Virtual Machine	836
Add New Virtual Machines to a Service	837
Configure Reservations	837
Configure Services.....	840
Configure Snapshots	841
Disable Software Deployment	841
Modify the Physical System Allocation Policy	842
Specify When to Send Pending Approval Request Notification.....	843
Set Automatic Cancellation of Unapproved Reservations	843
Allow Users to Select Storage Tiers.....	844

Chapter 13: Scalability Best Practices 845

Scalability Overview	845
Hardware Specifications.....	846
Database Considerations.....	846
Network Considerations.....	847
Remote Deployment and Policy Configuration Overview.....	847
Scalability Recommendations	849
vCenter AIM Monitoring Recommendations.....	849
CA Server Automation vCenter Management Recommendations	850
LPAR AIM Monitoring Recommendations	852

Solaris Zones AIM Monitoring Recommendations.....	853
Remote Deployment and Policy Configuration Recommendations.....	854
Appendix A: FIPS 140-2 Encryption	863
FIPS Overview.....	863
Appendix B: Configure AIMs with NodeCfgUtil	865
NodeCfgUtil Overview	865
Configure AIMs with NodeCfgUtil in Dialog Mode	867
Configure AIMs with NodeCfgUtil in Command Mode	870
Appendix C: Troubleshooting	873
CA Server Automation Troubleshooting	873
Adjusting Poll Interval Settings for Solaris Zones Environments.....	875
Attributes Show a Value of Zero	875
Browsers Do Not Display Consecutive Spaces in Events.....	875
CA NSM Fails to Receive Traps on Windows	876
Cisco UCS Folder Does Not Display in UI.....	876
DB Transaction Log Sizes Increase Unexpectedly	876
Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero	877
Domain Server is not available.....	877
Empty Task ID for the dpmvc virtualswitch Command	878
Local and Remote Monitors Do Not Show the Same Values	878
Remote Deployment to Solaris Lists SPARC and x86 Systems	879
Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear	879
Resetting the vCenter Server Password Causes Data Collection to Fail.....	879
Solaris Zones AIM Reset if a Monitored System is Down.....	880
Status Icon of Component Shows Not Configured.....	880
Unable to Connect to Microsoft SQL Server	880
Upgrading SystemEDGE	881
User Interface is Unresponsive on Provisioning and Policy Screens.....	881
vCenter Server AIM Attributes Show Zero	881
VM Usage Values Do Not Update Immediately After Power Down.....	882
Blank Query Results Tab after Upgrade	883
Accessing the CA Process Automation Server Requires Credentials After Configuration	884
CA DSM Functionality is Not Fully Supported	884
CA Configuration Automation Agent Stops During Installation	885
Deleted OS Images from CA ITCM are Not Deleted from CA Server Automation.....	885
Discovering Large Networks.....	886
Discovery Does Not Identify Operating System	886

Duplicated Zone Entries in the Managed Folder.....	887
Error When Installing CA DSM Agent and Asset Management Plug-in.....	887
ESX Job Status is Current But OS Installation Not Complete.....	888
ESX Machines Fail to Discover.....	888
Log In with Different User Credentials on the Same Computer Using IE8	888
No Cisco UCS Manager in Explore Pane	889
New System Name is not Displayed.....	890
OpenSSL Software Compatibility Issues.....	890
Password Changes May Cause Authentication Errors	890
Rapid Server Imaging (RSI) Troubleshooting.....	893
Reservation Manager Troubleshooting.....	899
Scheduled Jobs do not Run	905
CA SDM Exception Error.....	906
Software Delivery Adapter Errors	907
SSP - The Home content does not display in Internet Explorer 9	908
vCenter Server Folder Does Not Display in UI.....	908
VM Reservation Fails: Could Not Find Computer UID for Software Delivery.....	909
VMs Not Being Discovered.....	909

Glossary	911

Index	923

Chapter 1: Introduction

This section contains the following topics:

[Related Publications](#) (see page 17)

[Conventions](#) (see page 18)

Related Publications

The CA Bookshelf provides the following CA Server Automation publications:

Administration Guide

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

Installation Guide

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

Reference Guide

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

CA Process Automation Connector Reference Guide

Provides detailed information about CA Process Automation connectors and use cases.

Online Help

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

Reservation Manager Help

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

Release Notes

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

Provides installation and configuration details of SRM.

SystemEDGE User Guide

Provides end-user information about the SystemEDGE agent.

SystemEDGE Release Notes

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

In addition, the CA Bookshelf supplies the following Rapid Server Imaging (RSI) server guides:

- *RSI Server Administration Guide*
- *RSI Server Installation Guide*
- *RSI Server for AppLogic User and Installation Guide*
- *RSI Server Release Notes*

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

Conventions

This guide uses the following conventions:

Case-Sensitivity

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

Cross-References

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

Guide Name

Indicates the name of another guide.

"Chapter Name"

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

Synonyms

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

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

Syntax

Syntax and user input use the following form:

Italic

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

{a|b}

Indicates a choice of mandatory operands, a or b.

[] or [[]]

Indicates optional operands.

Syntax Example

The following example uses these conventions:

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

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

Default Directory

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

Installation Path

Install_Path used in path statements indicates the directory in which CA Server Automation or components of CA Server Automation are installed.

Defaults:

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

Chapter 2: Overview

This section contains the following topics:

[Architecture](#) (see page 21)

[Databases](#) (see page 25)

[User Interface](#) (see page 26)

Architecture

CA Server Automation is a policy-based product that 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) that analyzes your data center continuously to ensure that your servers are provisioned optimally to perform required tasks. Use the web-based CA Server Automation user interface to manage your data center and to obtain detailed information about each managed system in your data center.


CA Server Automation leverages the following CA technologies:

- CA Service Desk Manager (CA SDM) for request escalation and resolution
- CA IT Client Manager (CA ITCM) for operating system and application deployment
- CA Embedded Entitlements Manager (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 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 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 Software Delivery server. These components must also be registered back to one main service controller for proper functioning.

Provisioning

Provides a uniform integration point for CA ITCM, OS Installation Management technology (OSIM), VMware vCenter, Microsoft 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 ITCM 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 SDM 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 using 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 SDM for issue management and CA Software Delivery for image and package deployment 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; CA Software Delivery to troubleshoot issues with package deployment; or CA Configuration Automation to perform advanced change and configuration management functions. 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, create policy, schedule jobs, and so on. CA Server Automation Start menu shortcuts are only available on the CA Server Automation server. Use the Start menu to access product features such as the user interface and CLI command window. To access the interface from a different server, 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/UI`

servername

Identifies the CA Server Automation server.

port

Identifies the Apache Tomcat Server port.

Default: 8443

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

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

The Dashboard appears.

Chapter 3: Managing Users and User Groups

This section contains the following topics:

[User Access Control](#) (see page 29)

[Password Management](#) (see page 33)

[User Group Management](#) (see page 37)

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.

Security Considerations for Active Directory

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

The requirements are described in a Microsoft Knowledge Base article.

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

Native Security

Native Security lets the CA EEM administrator create users, user groups, and policies specifically for CA 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: If you receive a security certificate request, bypass it and continue. To eliminate these messages, acquire a certificate from the vendor of your choice and apply it to the server. For information about installing security certificates, see the Apache Tomcat website.

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

The User Name field is populated with EiamAdmin.

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

The CA 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 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.

User Group Management

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

More information:

[Search for Users or User Groups](#) (see page 37)

[Create a User Group](#) (see page 38)

[Assign Users to Groups](#) (see page 39)

[Assign External Directory User Groups to User Groups](#) (see page 39)

[Set User Group Privileges](#) (see page 41)

[Set User Group Permissions](#) (see page 41)

[Set User Group Permissions for Services](#) (see page 42)

[Set Run Command Script Privileges](#) (see page 42)

[Import External Directories](#) (see page 43)

[Delete User Groups](#) (see page 43)

[Assign User Groups Access Rights to Services](#) (see page 44)

[Remove Users or User Groups from a User Group](#) (see page 44)

Search for Users or User Groups

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

To search for users or user groups

1. Click Administration.

The Administration page appears.

2. Click User Group.

The User Groups page appears.

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

The user group page appears in the right pane.

4. Click Membership.

The User/User Group page appears.

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

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

Create a User Group

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

To create user groups

1. Click Administration.

The Administration page appears.

2. Click User Groups.

The User Groups page appears.

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

4. (Optional) Type a Description.

5. Click Save.

The new user group appears in the left pane.

More information:

[Assign Users to Groups](#) (see page 39)

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 37)
6. Select the user to add from the Available User/User Group pane, and click the right arrow .
The user name moves to the Selected User/User Group pane.
7. Click Save to finish adding users.
Users are granted the access privileges of their user group.

Assign External Directory User Groups to User Groups

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

To assign external directory user groups to user groups

1. Click Administration.
The Administration page appears.
2. Click User Groups.
The User Groups node appears on the left pane.

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

The user group page appears on the right pane.

4. Select the Membership submenu.

A series of membership panes appears below the tab.

5. Select User Group from the Identity list.

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

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

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

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

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

8. Click Save to finish adding user groups.

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

Set User Group Privileges

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

Note: For information about giving users access to CA Server Automation, see the *Administration Guide*.

To set user group permissions

1. Click Administration.
The Administration page appears.
2. Click User Group.
The User Groups page appears.
3. Select the user group for which to set permissions, and click the Privileges tab.
The Privileges page appears.
4. Click the check boxes for the tabs and actions to which you want to grant or restrict access, and click Save.
The user group privileges are updated.

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

Set User Group Permissions

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

To set user group permissions

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

Set User Group Permissions for Services

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

To set user group permissions

1. Click Administration.
The Administration page appears.
2. Click User Group.
The User Groups page appears.
3. Click a user group in the left pane.
4. Click Service Access.
The right pane displays the resources for which services are enabled or disabled.
5. Enable or disable resources for service access as needed.
6. Click Save.
The Service Access list is updated.

Set Run Command Script Privileges

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

To set run command script action privileges

1. Click Administration.
The Administration page appears.
2. Click User Group.
The User Groups page appears.
3. Expand User Group, and select a user group from the list.
Tabs appear in the right pane.
4. Select the Privileges tab.
A list of privileges appears, with check boxes to select or unselect privileges.
5. Expand the Policy folder, select Run Command Script, and click Save.
The command script privileges are updated.

Import External Directories

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

To import an external directory

1. Click Administration.
The Administration page appears.
2. Click User Group.
The User Groups page appears.
3. Expand User Groups and select a user group from the list.
4. Select Membership.
The Users/User Group page appears.
5. Select User Groups from the Identity drop-down list, type a name or partial name of an external directory in the Value text box, and click Search.

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

Delete User Groups

You can delete user groups that you no longer need.

To delete user groups

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

Assign User Groups Access Rights to Services

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

To assign user groups access rights to services

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

Remove Users or User Groups from a User Group

You can remove users and user groups from an existing CA 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.

Chapter 4: Managing Systems Performance

This section contains the following topics:

[Systems Management](#) (see page 47)

[Discovery](#) (see page 49)

[Services](#) (see page 54)

[Managed and Unmanaged Resources](#) (see page 58)

[SystemEDGE Features](#) (see page 60)

[Service Response Monitoring](#) (see page 72)

[Agent Visualization](#) (see page 75)

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 Systems Performance LiteAgent

Supported Virtual Environments and Operating Systems



Note: RHEV and ADES are available for CA Virtual Assurance only.

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 60)

[Service Response Monitoring](#) (see page 72)

[Interaction Between Remote Monitoring Components](#) (see page 544)

[Agent Configuration](#) (see page 68)

[Agent Visualization](#) (see page 75)

[Security and Maintenance](#) (see page 70)

Discovery

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

You can specify a single system to discover, manage, or assign this system to a service.

To discover a system

1. Select Resources, Manage, Discover System.

2. Complete the System Name field to specify the name or IP address of the server.
3. (Optional) Click Next.

The Enhanced Discovery and SNMP Information dialog appears.

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

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

Delete a System

Deleting a discovered system removes the system from CA Server Automation.

To delete a system

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

A list of discovered systems appears.

3. Select the system or multiple systems to delete and select Delete from the Actions drop-down menu.

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

A message prompts for confirmation.

4. Click Yes.

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

Discover a Network

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

172.24.143.0/24

You can also use wildcards and ranges:

172.24.143.*

172.24.143.{1-255}

To discover a network

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

Network Name

Specifies the name of the network.

Network Address

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

Exclude Address

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

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

Ping Sweep

Discovers all IP addresses in the network.

DNS

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

4. (Optional) Click Next.

The Enhanced Discovery and SNMP Information dialog appears.

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

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

Enhanced Discovery and SNMP Information

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

To discover using enhanced information

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

The Specify Discovery Type and Target section appears.

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

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

Enhanced Discovery

Select this option to specify enhanced credentials for discovery.

Discovery Credentials

Select *one* of the options to specify credentials.

Specify Credentials

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

Select Saved Credentials

Select the existing saved credentials from the Available list.

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

Override SNMP Defaults

Select this option to override the SNMP defaults for discovery.

SNMP Settings

Select *one* of the options to specify credentials.

Specify Credentials

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

Select Saved Credentials

Select the existing saved credentials from the Available list.

4. Click Finish

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

Cancel Network Discovery

You can cancel a network discovery that is in progress.

To cancel network discovery

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

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

Rediscover a Network

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

To rediscover a network

1. Select Resources, Manage, Manage Discovered Networks.

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

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

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

Delete a Network

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

To delete a network

1. Select Resources, Manage, Manage Discovered Networks.

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

2. Select a network to delete and click - (Delete) on the Network List toolbar.

A message prompts for confirmation.

3. Click OK.

The network is deleted from the network list.

Services

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

More information:

[Create a Service](#) (see page 55)

[Edit a Service](#) (see page 56)

[Remove Server from Services](#) (see page 57)

[Delete Services](#) (see page 58)

Create a Service

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

To create a service

1. Click Resources, and open the Explore pane.
2. Select a parent service node, such as Data Center or CA Server Automation Services.
3. Right-click Management, New Service.

The Service: New dialog appears.

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

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

Service Priority

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

Example:

ServiceA: Priority 3

ServiceB: Priority 1

ServiceC: Priority 2

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

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

Lag

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

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

Lower and Upper Threshold %

Specifies the lower and upper thresholds of the entire service.

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

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

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

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

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

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

The servers are added to the Selected Servers section.

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

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

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

Edit a Service

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

To modify a service

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

The Service: Edit dialog appears.

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

Service Priority

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

Example:

ServiceA: Priority 3

ServiceB: Priority 1

ServiceC: Priority 2

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

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

Lag

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

Lower and Upper Threshold %

Specifies the lower and upper thresholds of the entire service.

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

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

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

The servers are added to the Selected Servers section.

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

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

7. Click Save.

The Servers lists are updated.

Remove Server from Services

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

To remove a server from a service

1. Click Resources.

The Resources page appears.

2. Expand the Data Center folder and the CA Server Automation Services folder in the Explore pane.

The discovered and managed resources in the Data Center appear.

3. Right-click the server and select Management, Remove from Service.

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

4. Click OK.

The server is removed from the service.

Delete Services

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

To delete a service

1. Select Resources, Manage, Manage Services.
A list of services appears in the right pane.
2. Select a service and click - (Delete) on the Services toolbar.
A message prompts for confirmation.
3. Click Yes.
The service is deleted.

Managed and Unmanaged Resources

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

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

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

More information:

[Unmanage Managed Resources](#) (see page 59)

[Manage Unmanaged Resources](#) (see page 59)

[Delete Managed Resources](#) (see page 60)

Unmanage Managed Resources

You can stop managing currently managed servers.

Follow these steps:

1. Click Resources.

The Resources page appears.

2. Expand the Data Center folder and the CA Server Automation Services folder in the Explore pane.

The discovered and managed resources in the data center appear.

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

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

4. Click OK.

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

Manage Unmanaged Resources

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

To manage a resource

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

The list of unmanaged resources appears.

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

A message confirms that the resource is Managed.

5. Expand the Managed folder.

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

Delete Managed Resources

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

To delete a managed resource

1. Click Resources.
2. The Resources page appears.
3. Expand the Data Center folder and the CA Server Automation Services folder in the Explore pane.

The discovered and managed resources in the data center appear.

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

A message prompts you to confirm deletion.

5. Click OK.

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

SystemEDGE Features

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

SystemEDGE supports monitoring metrics from the following MIBs:

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

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

Self monitoring

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

Process and service monitoring

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

Process group monitoring

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

Log file and directory monitoring

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

Windows event monitoring

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

History collection

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

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

More Information

[Systems Management MIB](#) (see page 62)

[State Management Model](#) (see page 64)

[Configure Object Aggregation](#) (see page 216)

[Stateless Monitoring](#) (see page 65)

[Managed Mode and Unmanaged Mode](#) (see page 66)

Systems Management MIB

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

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

System Group (sysedgeSystem)

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

Mounted Devices Table (devTable)

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

Kernel Configuration Group (kernelConfig)

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

Boot Configuration Group (bootconf)

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

Streams Group (streams)

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

User Table (userTable)

Contains information about the user accounts on the system.

Group Table (groupTable)

Contains information about the user groups on the system.

Process Table (processTable)

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

Who Table (whoTable)

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

Remote Shell Group (remoteshell)

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

Kernel Performance Group (kernelperf)

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

Interprocess Communication Tables (msgqueueTable, shmемTable, semTable)

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

Message Buffers Allocation Table (mbufAllocTable)

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

Streams Buffers Allocation Table (strbufAllocTable)

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

I/O Buffer Cache Group (ioBufferCache)

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

Directory Name Lookup Cache Group (dnlc)

Contains information about directory and file name cache performance.

AIX Logical Partition Group (logicalPartition)

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

Trap Community Table (trapCommunityTable)

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

NT System Group (ntSystem)

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

RPC Statistics Group (rpc)

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

NFS Statistics Group (nfs)

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

Disk Statistics Table (diskStatsTable)

Contains information about disk I/O.

CPU Statistics Table (cpuStatsTable)

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

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

State Management Model

The SystemEDGE agent supports a state management model for self monitors and process monitors fully integrated with the overall CA 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.

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 Unmanaged Mode

When you deploy SystemEDGE (or install it on a standalone basis), you can specify to run the agent in managed mode. In managed mode, the agent is managed by the CA 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.

Application Insight Modules (AIMs)

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

AIM for VMware vCenter Server

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

AIM for VMware vCloud Director

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

AIM for IBM PowerVM (LPARs)

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

AIM for IBM High Availability Cluster Multiprocessing

Provides capabilities to monitor an IBM High Availability Cluster Multiprocessing (HACMP) system.

AIM for Microsoft Hyper-V

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

AIM for Microsoft Cluster Services

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

AIM for Solaris Zones

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

AIM for Cisco Unified Computing System (UCS)

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

AIM for Citrix XenServer

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

AIM for Service Response Monitor

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

AIM for Remote Monitoring

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

More Information

[NodeCfgUtil Overview](#) (see page 865)

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 69)
[Configuration Overview](#) (see page 158)

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 an existing self or process monitor.

Monitoring Software Settings

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

Follow these steps:

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Select a system or a service.
3. Click the Resources tab, and then Monitoring Software.
The Machine Details pane appears.
4. Modify the settings according to your needs, and click Apply:

System Description

Defines a description of the system.

System Contact

Defines a contact for the system.

System Location

Defines the system location.

SysEdge Log Level

Specifies the SystemEDGE log level.

The settings are updated.

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 71)

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.

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 73)

SRM Tests

The SRM AIM provides the following response time tests:

Active Directory

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

Custom

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

DHCP

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

DNS

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

File I/O

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

FTP and TFTP

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

HTTP and HTTPS

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

LDAP

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

NIS

Verifies that NIS map requests are being processed.

NNTP

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

Ping

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

Email

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

SNMP

Verifies that SNMP agents are responding to SNMPv1 GET requests.

SQL Query

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

TCP

Verifies that systems are listening for and processing connection requests.

Virtual User

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

Agent Visualization

The CA 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 76)

[View SystemEDGE Monitors](#) (see page 75)

[View SRM Tests](#) (see page 77)

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 69) 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 Managed Object States

The CA Server Automation user interface displays all SystemEDGE managed objects for managed systems.

To view managed object states, click Resources, expand an appropriate folder in the Explore tree, and select a 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 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

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

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

This section contains the following topics:

[User Permissions and Access Requirements Reference](#) (see page 79)

[How to Configure SNMP and Access Control Lists](#) (see page 88)

[How to Deploy SystemEDGE and AIMs](#) (see page 112)

[How to Configure SystemEDGE and Service Response Monitor Through Policies and Templates](#) (see page 157)

[How to Change the Configuration Mode for SystemEDGE](#) (see page 263)

User Permissions and Access Requirements Reference

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

This documentation is intended for:

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

More information:

[Cisco UCS](#) (see page 80)

[Citrix XenServer](#) (see page 81)

[Hyper-V](#) (see page 82)

[IBM High Availability Cluster Multiprocessing](#) (see page 82)

[IBM PowerVM](#) (see page 83)

[Microsoft Cluster Server](#) (see page 84)

[Oracle Solaris Zones](#) (see page 84)

[Agent Remote Deployment](#) (see page 85)

[Remote Monitoring](#) (see page 86)

[SystemEDGE and Advanced Encryption](#) (see page 87)

[VMware vCenter](#) (see page 87)

[VMware vCloud](#) (see page 88)

Cisco UCS

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

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

Note: We recommend giving the UCS user admin privilege.

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

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

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

Communication Ports

HTTP Port: 80

HTTPS Port: 443

Citrix XenServer

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

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

Communication Ports

HTTPS Port: 443

SNMP Port: 161

Hyper-V

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Requires local Administrator Account.

SCVMM Monitoring

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

Communication Port

Windows RPC Endpoint Mapper Port: 135

DCOM/WMI Port: dynamically assigned during RPC Endpoint negotiation

IBM High Availability Cluster Multiprocessing

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring HACMP

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

- clstat
- clRGinfo -s
- cldump
- clIsnw
- cltopinfo
- clIsif
- clshowsrv -v
- vmstat

Communication Port

Secure Shell TCP Port: 22

IBM PowerVM

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

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

Monitoring Hardware Management Console (HMC)

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

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

Monitoring Virtual IO Server (VIOS)

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

Monitoring Integrated Virtualization Manager (IVM)

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

Note: IVM Monitoring requires IVM configuration.

Communication Port

Secure Shell TCP Port: 22

Microsoft Cluster Server

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

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

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

Communication Port

Windows RPC Endpoint Mapper Port: 135

DCOM/WMI Port: dynamically assigned during RPC Endpoint negotiation

Oracle Solaris Zones

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

Requires root user access.

Communication Port

Secure Shell TCP Port: 22

Agent Remote Deployment

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

Installation on Windows

Requires Windows System Administrator privileges.

Installation on Linux

Requires root access or use of sudo or pfexec.

Cross-platform Remote Deployment

Uses Infrastructure Deployment (ID).

ID manager component

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

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

Communication Ports

Remote Deployment (Windows)

CIFS UDP Port: 137 (Inbound/Outbound)

CIFS UDP Port: 138 (Inbound/Outbound)

TCP Port: 135 (Inbound)

CIFS TCP Port: 139 (Inbound/Outbound)

CIFS TCP Port: 445 (Inbound/Outbound)

CAM UDP Port: 4104 (Inbound/Outbound)

CAM TCP Port: 4105 (Configurable)

Remote Deployment (UNIX, Linux)

CAM UDP Port: 4104 (Inbound/Outbound)

Secure Shell TCP Port: 22 (Inbound)

TCP Port: 135 (Inbound)

CAM TCP Port: 4105 (Configurable)

Remote Monitoring

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Remote Monitoring

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

Communication Port

Windows RPC Endpoint Mapper Port: 135

DCOM/WMI Port: dynamically assigned during RPC Endpoint negotiation

As a best practice the Remote Monitoring systems must be a member of an AD Domain. This membership lets you use a domain account and avoids the need to define local user accounts on each RM System. Create a CARMuser domain account that is a member of the Domain Admins group of the AD Domain.

When user credential settings are prompted for during RM installation, provide the domain account with the password. For any system member of this domain, no additional configuration is required.

Note: If necessary, you can restrict the CARMuser access rights so the user is not a member of the Domain Admins group. In this case, configure WMI Namespace access and DCOM access. For more information about defining WMI Namespace access and DCOM access, see the Microsoft website.

SystemEDGE and Advanced Encryption

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Installation on Linux or Unix

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

Monitoring

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

Communication Ports

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

UDP Trap Port: 162 (Outbound)

SystemEDGE in Managed Mode uses CAM:

CAM UDP Port: 4104

CAM TCP Port: 4105

VMware vCenter

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

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

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

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

Communication Port

HTTPS Port: 443

VMware vCloud

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

AIM Installation

Requires Windows System Administrator privileges on the AIM host.

Monitoring

(Valid for AIM) Requires the System Administrator role.

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

System Administrator@System

Grants full access.

Organizational Access@org_name

Limits operation at the organization level and the role assignment.

Communication Port

REST API Port: 8443

How to Configure SNMP and Access Control Lists

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

More information:

[SNMP Consistency](#) (see page 88)

[Global and Server-level SNMP Settings](#) (see page 89)

[How to Configure SNMPv1/v2 Settings and Access Control Lists](#) (see page 91)

[How to Manage Server-level SNMP Settings](#) (see page 102)

[How to Configure SNMPv3](#) (see page 106)

SNMP Consistency

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

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

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

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

More information:

[Global and Server-level SNMP Settings](#) (see page 89)

Global and Server-level SNMP Settings

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

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

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

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

- Port
- Read Community
- Read-Write Community

Alternatively, you can edit the fields inline.

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

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

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

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

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

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

More information:


[How to Configure SNMPv1/v2 Settings and Access Control Lists](#) (see page 91)

[How to Manage Server-level SNMP Settings](#) (see page 102)

[How to Configure SNMPv3](#) (see page 106)

How to Configure SNMPv1/v2 Settings and Access Control Lists

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



Follow these steps:

[Review Requirements \(SNMPv1/2\)](#) (see page 92)

[Review SNMP Configuration and Policy Relationships](#) (see page 92)

[Specify Global SNMP Settings and Access Control Lists](#) (see page 94)

[Apply Global SNMP Settings and Access Control Lists to Policies](#) (see page 95)

[\(Optional\) Specify Access Control Lists at the Policy Level](#) (see page 96)

[\(Optional\) Specify SNMP Settings and Access Control Lists at the Server Level](#) (see page 97)

[Distribute Policies to Server Groups](#) (see page 98)

[Example for Three Server Groups](#) (see page 99)

Review Requirements (SNMPv1/2)

Review the following requirements before you start configuring the SNMP settings on CA Server Automation:

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

More information:

[Review SNMP Configuration and Policy Relationships](#) (see page 92)

Review SNMP Configuration and Policy Relationships

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

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

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


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

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

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

The following diagram illustrates the policy architecture:

Policy Architecture



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

More information:

[Specify Global SNMP Settings and Access Control Lists](#) (see page 94)

[Apply Global SNMP Settings and Access Control Lists to Policies](#) (see page 95)

[\(Optional\) Specify Access Control Lists at the Policy Level](#) (see page 96)

[\(Optional\) Specify SNMP Settings and Access Control Lists at the Server Level](#) (see page 97)

[Distribute Policies to Server Groups](#) (see page 98)

[Example for Three Server Groups](#) (see page 99)

Specify Global SNMP Settings and Access Control Lists

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

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

Follow these steps:

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

SNMP settings page appears.

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

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

3. Specify the required data to create an SNMP Settings object and click OK.
4. Select the SNMP Settings object to which you want to add an ACL and click the Edit icon.

Edit SNMP Settings dialog with the ACL panel appears.

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

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

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

More information:

[Apply Global SNMP Settings and Access Control Lists to Policies](#) (see page 95)

[\(Optional\) Specify Access Control Lists at the Policy Level](#) (see page 96)

[\(Optional\) Specify SNMP Settings and Access Control Lists at the Server Level](#) (see page 97)

[Distribute Policies to Server Groups](#) (see page 98)

Apply Global SNMP Settings and Access Control Lists to Policies

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

Follow these steps:

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

The Policy page appears.

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

The SystemEDGE page appears listing the available policies.

3. (Optional) Click  to create a policy.

The New SystemEDGE Policy dialog appears.

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

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

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

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

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

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

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

- Server Read
- Server Write

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

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

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

7. Click Save Policy.

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

More information:

[\(Optional\) Specify Access Control Lists at the Policy Level](#) (see page 96)

[\(Optional\) Specify SNMP Settings and Access Control Lists at the Server Level](#) (see page 97)

[Distribute Policies to Server Groups](#) (see page 98)

(Optional) Specify Access Control Lists at the Policy Level

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

Follow these steps:

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

The ACL dialog opens.

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

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

4. Click Save Policy.

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

More information:

[\(Optional\) Specify SNMP Settings and Access Control Lists at the Server Level](#) (see page 97)

[Distribute Policies to Server Groups](#) (see page 98)

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

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

Follow these steps:

1. Navigate to Resources, Explore in the user interface.
The Explore pane appears.
2. Expand the Explore tree and right-click the system for which you want to specify SNMP credentials and an Access Control List.
3. Select Policy, Configure SNMP Settings from the pop-up menu.
The SNMP Settings dialog lists the valid SNMP settings for that system.
4. Click Add.
The New SNMP Settings dialog appears.
5. Specify Name, Port, Community String, the type of operation (read-only or read-write), SNMP version, port, timeout, and retry limit. Use the port number of the installed SystemEDGE on the server. Click OK.
6. Close the dialogs, change to the selected system page, and click the SNMP Access Control tab.
The specified system-specific SNMP Community Settings are listed.
7. Click the Edit link from the Access Control List column.
The system-specific Access Control List dialog appears.

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

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

9. Click Save.

Distribute the policy to the appropriate server group.

More information:

[Distribute Policies to Server Groups](#) (see page 98)

Distribute Policies to Server Groups

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

Follow these steps:

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

The Policy page appears.

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

The SystemEDGE page appears listing the available policies.

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

The policy page opens.

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

4. Click Action, Apply.

The Select Machines page appears.

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

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

The new settings are applied to the target systems.

More information:

[Example for Three Server Groups](#) (see page 99)

Example for Three Server Groups


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

The datacenter consists of the following server groups:

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

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

Access Control List Relationships



Follow these steps:

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

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

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

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

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

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

development-write: port 161, read-write access, community_admin_, ACL: va-im.ca.com, im1.ca.com
2. Create three policies (one for each server group) that are based on the default policy: infrastructure, sales, and development
3. Change to the infrastructure policy page, select the third option to apply global SNMP settings from the table:
 - Custom Selection
4. Add infrastructure-read and infrastructure-write global SNMP objects to the infrastructure policy.
5. Save the policy.

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

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

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


More information:

[Review SNMP Configuration and Policy Relationships](#) (see page 92)

How to Manage Server-level SNMP Settings

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

How to Manage server-level SNMP Settings



Follow these steps:

[Review Requirements \(Server-level\)](#) (see page 102)

[Add Server-level SNMP Settings](#) (see page 103)

[Apply the Package Wrapper SNMP Settings as Server-level Settings](#) (see page 104)

Review Requirements (Server-level)

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

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

Add Server-level SNMP Settings

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

Follow these steps:

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Expand the Data Center folder, then any subfolder, and select the server that you want to configure.
3. Right-click and select Policy.
The Policy submenu appears.
4. Click Configure SNMP Settings.
The Configure SNMP Settings dialog opens showing the server-level settings.
5. Do one of the following options:
 - Select the check box for an existing metric from the list, and click the tool icon (Edit) to modify an existing entry.
 - Click Add to create an SNMP entry at the server level.The New SNMP Settings dialog appears.
6. Complete the following fields, and click OK:

Name

Describes the SNMP credentials that are being defined.

Port

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

SNMP version

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

Community String (for SNMP v1/v2)

Specifies the SNMP community string.

Security User (for SNMP v3)

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

Access Type

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

Timeout

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

Default: 10 seconds

Retry limit

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

Authentication (for SNMP v3)

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

Privacy (for SNMP v3)

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

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

Apply the Package Wrapper SNMP Settings as Server-level Settings

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

Follow these steps:

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

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

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

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

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

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

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

The policy pane appears.

4. Select the appropriate items under Server Communities.

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

5. Select the appropriate items under Default Communities.


The Default Communities represent the global SNMP settings.

6. Change to Resources, Deployment, Packages and open the SystemEDGE package wrapper.
7. Set the SNMP credentials, select the policy that you want to apply to the SystemEDGE agents after installation on the managed node, and save the wrapper.
8. Create a Deployment Job and deploy the SystemEDGE package with its policy to the managed nodes.

How to Configure SNMPv3

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

How to Configure SNMPv3



Follow these steps:

[Review Common Requirements \(SNMPv3\)](#) (see page 107)

[Review SNMPv3 Configuration Details](#) (see page 107)

[Create a Global SNMPv3 Object](#) (see page 108)

[Create a Policy](#) (see page 109)

[Apply the Policy](#) (see page 111)

[Verify the SNMPv3 Settings in the System Summary](#) (see page 111)

Review Common Requirements (SNMPv3)

Review the following requirements before you start configuring the SNMP settings on CA Server Automation:

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

Review SNMPv3 Configuration Details

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

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

Create a Global SNMPv3 Object

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

Follow these steps:

1. Navigate to Administration, SNMP in the user interface.
SNMP settings page for global objects appears.
2. Click  to create an SNMP Settings object.
The New SNMP Settings dialog appears.
3. Set SNMP Version to SNMPv3.
The SNMPv3-related fields appear in the dialog.
4. Complete the following fields, and click OK:

Name

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

Port

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

SNMP version

Specifies SNMPv3 (already set in the previous step).

Security User

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

Access Type

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

Timeout

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

Default: 10 seconds

Retry limit

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

Authentication

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

Privacy

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

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

Create a Policy

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

Follow these steps:

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

The Policy page appears.

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

The SystemEDGE page appears listing the available policies.

3. Click  to create a policy.

The New SystemEDGE Policy dialog appears.

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

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

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

- Include only server-specific SNMP settings
- Include server-specific SNMP settings and all default (global) settings
- Custom Selection

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

6. Select Custom Selection.

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

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

The selected SNMPv3 objects are associated with the policy.

8. Click the Trap Destinations tab.

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

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

Depending on the selection, the following fields appear:

Destination

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

Port

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

Username

Specifies the SNMPv3 user with which to send the trap.

Encoding

Specifies the type of encoding to use when sending traps.

Default: 000

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

Context

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

Timeout

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

Retries

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

10. Fill out the fields and click Add.

A new entry appears in the Trap Destinations table.

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

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

The policy is saved with the appropriate trap destination.

You can distribute the policy to the appropriate server group.

Apply the Policy

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

Follow these steps:

1. Navigate to Resources, Configuration in the user interface.
The Policy page appears.
2. Expand Policy, Policies, SystemEDGE in the navigation pane.
The SystemEDGE page appears listing the available policies.
3. Select the policy that you have previously saved with appropriate SNMPv3 settings.
The policy page opens.
4. Click Action, Apply.
The Select Machines page appears.
5. Select all systems which you want to configure with that policy, and click Apply.
You can view the delivery status or return to the Policy page.
The new settings are applied to the target systems.

Alternative

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

Verify the SNMPv3 Settings in the System Summary

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

Follow these steps:

1. Expand the components tree in the Explore pane.
2. Select a managed system to which you have applied the SNMPv3 settings and open Summary.
The Machine Status Information appears.
3. Verify that the Active SNMP Credentials field shows the SNMPv3 global object.

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

How to Deploy SystemEDGE and AIMs

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

More information:

[Overview](#) (see page 112)

[Configuration](#) (see page 114)

[Scalability](#) (see page 117)

[Deployment Packages](#) (see page 119)

[Using Remote Deployment](#) (see page 132)

[Specific Remote Deployment Use Cases](#) (see page 142)

[Deployment Jobs](#) (see page 149)

[Infrastructure Deployment Process](#) (see page 150)

Overview

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, cancellation, and filter 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:



About Packages

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

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

Deployment Components

This section lists and briefly describes the deployment key components:

Domain Server

The domain server is the repository for all configuration and control data. The server is responsible for managing configuration and software package data required for deployment operations and it manages all configuration operations. Detailed event data is passed between the domain and distribution servers during the deployment process. A single Domain Server is responsible for maintaining the status of all Distribution Server jobs.

Distribution Server

The distribution server controls the Infrastructure Deployment Manager (IDManager) server located on the same machine. The architecture allows for multiple distribution servers offering deployment services.

Infrastructure Deployment

Infrastructure Deployment initiates and manages deployment jobs. During the deployment process, the Infrastructure Deployment Manager (IDManager) provides access to remote systems and the Infrastructure Deployment Primer (ID Primer) provides the mechanism to remotely install agent software packages. The IDPrimer is used to transfer the deployment package data to the target computer and run the installation. All subsequent deployments to the same target computer can use the existing IDPrimer installation. The IDManager controls all the deployment operations and handles job status.

Configuration

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

More information:

[Deployment Dashboard Views](#) (see page 115)

[Enhanced Search Functionality for Remote Deployment](#) (see page 116)

[Job Status Filter](#) (see page 116)

[Change the Domain Server a Distribution Server Connects To](#) (see page 117)

Deployment Dashboard Views

The following views are available on the Dashboard for tracking deployment metrics:

Deployment Task Summary

Displays a pie chart and a list showing the number of completed, active, pending, and failed deployment tasks.

Unresolved Deployment Tasks

Displays a list of deployments that did not complete successfully. You can click the job ID to view details about why the task is unresolved.

Active Deployment Tasks

Displays a list of deployment tasks that are currently active. Details include the associated deployment job, target, package, and current state. You can click the task ID for details about the current status.

Deployment Package Summary

Displays a bar chart showing the number of deployments for each deployment package type.

Completed Deployment Jobs

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

Enhanced Search Functionality for Remote Deployment

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

The benefits are as follows:

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

Follow these steps:

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

Example:

Deploy or Remote or remote deployment.

A list of Remote Deployment links appears.

2. Select the appropriate Remote Deployment operation.
Perform the Remote Deployment operation.

Job Status Filter

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

Follow these steps:

1. Select Resources, Deploy.
The Deployment pane displays the Packages, Templates, and Jobs folders.
2. Click Jobs.
The job details appear in the right pane.
3. (Optional) Select/unselect the check box for Job status columns.
The customized columns appear.
4. Select/unselect the Filters of the columns.
Jobs appear per the filter selection.

Change the Domain Server a Distribution Server Connects To

In the situation where the network address of the domain server machine changes after the original installation, it is necessary to reconfigure the distribution server to connect to a new network address.

Prior to making the configuration change shown below it is important to verify that the new network address is connectable from the distribution server. If the distribution cannot make a connection to the domain server using the new address then deployment functionality will not work correctly.

To change the domain server a distribution server connects to

1. From the Start menu, open Administrative Tools, Services.

The Services user interface appears, listing the installed services.

2. Right-click CA SM Distribution Server and select Properties.

The Properties dialog appears.

3. Click Stop to stop the service.

4. Enter the following parameter into the Start parameters field:

`-m domainserver`

The *domainserver* parameter specifies the IP address or DNS name of the domain server.

5. Click Start

The distribution server will now attempt to connect to the domainserver address entered.

Scalability

The deployment system provides a degree of scalability using multiple distribution servers as a scalability layer. Each distribution server communicates with one IDManager instance. The IDManager can manage multiple component deployments to multiple target computers. CA Server Automation supports many simultaneous deployments because of this federated model.

More Information

[Remote Deployment and Policy Configuration Recommendations](#) (see page 854)

[Scalability Use Cases](#) (see page 856)

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 require less time.

Recommendations:

- Verify that target servers typically meet the requirements for deploying software remotely.
- Install additional distribution servers local to the target location.
- Deploy using distribution servers that are local to targets if possible.
- Schedule deployments to start during periods of low network traffic if possible.

See Scalability Best Practices for details.

More information:

[Scalability Best Practices](#) (see page 845)

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.

Deployment Packages

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

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

The following deployment packages are available:

CCA Agent

Provides the CCA Agent.

Performance Agent (CA Systems Performance LiteAgent)

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

SystemEDGE

Provides the core SystemEDGE agent.

SystemEDGE Advanced Encryption

Provides a FIPS 140 compliant encryption package for SystemEDGE.

SystemEDGE AIX LPAR

Provides the AIM for IBM PowerVM (LPAR).

SystemEDGE CXEN

Provides the AIM for Citrix XenServer.

SystemEDGE HACMP

Provides the AIM for IBM High Availability Cluster Multi-Processing.

SystemEDGE Hyper-V

Provides the AIM for Microsoft Hyper-V.

SystemEDGE MSCS

Provides the AIM for Microsoft Cluster Support (MSCS).

SystemEDGE RM

Provides the Remote Monitoring AIM.

SystemEDGE Solaris Zone

Provides the AIM for Oracle Solaris Zones.

SystemEDGE SRM

Provides the Service Response Monitor AIM.

SystemEDGE UCS

Provides the AIM for Cisco UCS.

SystemEDGE VC

Provides the AIM for VMware vCenter.

SystemEDGE VCLLOUD

Provides the AIM for VMware vCloud Director.

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 a set of default values for the chosen software package. If a package requires mandatory parameters, specify these parameters and save the settings before you deploy the package.

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

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

SystemEDGE

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

- Port
- Read Community
- Read-Write Community

Alternatively, you can edit the fields inline.

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

Install Path

Defines the root installation directory for the package.

Data Path

Defines the data directory for the package.

Shared Path

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

Port

Defines the SystemEDGE port number.

Default: 161

Description

Defines the SNMP system description.

Location

Defines the SNMP system location.

Contact

Defines the SNMP System contact.

Read Community

Defines the SNMP read-only community string.

Default: public

Read-Write Community

Defines the SNMP read-write community string

Trap Community

Defines the SNMP trap community string.

Trap Destination

Defines the SNMP trap destination host name.

Trap Port

Defines the SNMP trap port.

Default: 162

Privilege Separation User (UNIX/Linux)

Specifies the user name under which credentials the agent run during SNMP communication.

This entry instructs the agent to run SNMP communication under another user account. The agent also uses the default group of this user as an effective group.

Default: The agent operates using root account.

Start Agent check box

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

Suppress Reboot check box

Specifies whether to suppress any automatic reboot that can otherwise occur at the 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 parameters, see the *Installation and Deployment* chapter in the *SystemEDGE User Guide*.

SystemEDGE Advanced Encryption

Suppress Reboot check box

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

CA SystemEDGE AIX LPAR

LPAR Host

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

Username

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

Password

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

CA SystemEDGE CXEN

CXEN Hostname

Specifies the hostname to use for Citrix XenServer integration.

CXEN Username

Specifies the username to use for Citrix XenServer integration.

CXEN Password

Specifies the password to use for Citrix XenServer integration.

Suppress Reboot check box

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

CA SystemEDGE HACMP

HACMP hostname

Specifies the hostname to use for connecting to HACMP. Specify an HACMP host name to deploy this package.

HACMP username

Specifies the username to use for connecting to HACMP. Specify an HACMP user name to deploy this package.

HACMP password

Specifies the password to use for connecting to HACMP. Specify an HACMP password to deploy this package.

HACMP Port

Defines the HACMP port number.

Default: 22

Suppress Reboot check box

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

CA SystemEDGE Hyper-V

Suppress Reboot check box

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

CA SystemEDGE MSCS

MSCS Hostname

Specifies the hostname to connect to the cluster.

MSCS Username

Specifies the username to connect to the cluster.

MSCS Password

Specifies the password to connect to the cluster.

Suppress Reboot check box

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

CA SystemEDGE RM

Default WMI username

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

Default WMI password

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

Suppress Reboot check box

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

CA SystemEDGE SRM

Allow scripts check box

Specifies whether to allow running scripts as tests.

Allow File I/O tests check box

Specifies whether to allow running file I/O as tests.

Allow untrusted SSL check box

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

Disable user TOS check box (Windows)

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

Suppress Reboot check box (Windows)

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

CA SystemEDGE Solaris Zone

Zones Host

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

Username

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

Password

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

CA SystemEDGE UCS

UCS hostname

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

UCS username

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

UCS password

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

UCS protocol

Specifies what protocol to use, HTTP or HTTPS.

Port

Defines the UCS port number.

Default: 80 for HTTP or 443 for HTTPS.

Suppress Reboot check box

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

CA SystemEDGE VC

vCenter hostname

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

vCenter username

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

vCenter password

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

vCenter port

Defines the vCenter port number.

Default: 443

vCenter Protocol

Specifies what protocol to use, HTTP or HTTPS.

Suppress Reboot check box

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

CA SystemEDGE VCLLOUD

VCLLOUD hostname

Specifies the hostname to use for connecting to vCloud.

VCLLOUD username

Specifies the username to use for connecting to vCloud.

VCLLOUD password

Specifies the password to use for connecting to vCloud.

VCLLOUD port

Defines the vCloud port number.

Default: 443

Suppress Reboot check box

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

CA Systems Performance LiteAgent

Shared Path

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

Install Path

Defines the root installation directory for the package.

Suppress Reboot check box

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

More Information

[Create a New Package Wrapper](#) (see page 133)

[Modify a Package Wrapper](#) (see page 134)

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 131) section.

The package library is located in the following directory:

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

The directory tree layout is defined by the requirements of the IDManager component. The package library itself consists of a top-level packages directory which contains two sub-directories, Public and Private. The Public directory contains all the deployable software packages.

```

..
└── SM
    ├── CA_LiteAgent
    ├── CA_ProcProbe_Utility
    ├── CA_SystemEDGE_AdvancedEncryption
    ├── CA_SystemEDGE_Core
    │   ├── 5.7.0
    │   │   └── ENU
    │   │       ├── AIX_aix
    │   │       ├── HPUX_hp
    │   │       ├── HPUX_ia64
    │   │       ├── Linux_ia64
    │   │       ├── Linux_ppc
    │   │       ├── Linux_x86
    │   │       ├── Solaris_sparc
    │   │       ├── Solaris_x86
    │   │       ├── Windows_ia64
    │   │       ├── Windows_x64
    │   │       └── Windows_x86
    ├── CA_SystemEDGE_CXEN
    ├── CA_SystemEDGE_HACMP
    ├── CA_SystemEDGE_HyperV
    ├── CA_SystemEDGE_LPAR
    ├── CA_SystemEDGE_MSCS
    ├── CA_SystemEDGE_RM
    ├── CA_SystemEDGE_SRM
    ├── CA_SystemEDGE_UCS
    ├── CA_SystemEDGE_VCLOUD
    ├── CA_SystemEDGE_Zone
    └── CA_VMVCAIM

```

The top-level Public directory has five sub-directories:

Component Name

Must be the IDManager instance name, which for CA 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 5.7.0

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.

Package Filter

If upgrades have been applied to the server, Remote Deployment can show an increasing number of package versions. The default behavior of this release is to show only the latest packages available. As a consequence, the data in the Packages - Details tab is also filtered to show only the latest versions of each package. Selecting a wrapper from this panel, expands the tree at the selected wrapper position.

If you want to see all packages, you can override the default filtering behavior by the check box "Display Latest Package Versions Only" in the Package Information Panel. When enabled (default), it filters out any older package versions from the left tree and the Package Details Tab.

To change the filtering behavior

1. Select/Unselect the check box for "Display Latest Package Versions Only" in the Package Information Panel.
2. Refresh the Deployment view in the user interface.

The latest package versions appear.

Note: All other locations within the UI where package versions are displayed are not affected.

Deployment Package Configuration File

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

The configuration files provide the following:

- A localizable description of the installation package.
- A mechanism by which package dependencies are encoded in a machine-readable format.
- Documenting the publicly accessible installation parameter types.
- Additional context to the parameter types so a level of validation may be performed within the UI.
- A mapping between parameter names and the tokens used to represent those parameters in the packages installation program, in a platform independent manner.
- Specifies how the installation materials should be executed on a target machine.
- Mapping between the installer exit codes and those understood by the deployment system.

Localized elements of the `pkginfo.xml` file is optionally provided using either side by side locale-specific files or embedded within a single file. The file name that matches with `pkginfo_PLATFORM.xml` is loaded to obtain localized message data.

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

```
%AllUsersProfile%\CA\SM\domainserver\Deployment\Packages\SM\CA_SystemEDGE_Core\5.7.0\ENU
```

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

Using Remote Deployment

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

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.

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*.

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.

Copy a Package Wrapper

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

To copy 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 copy.

The wrapper details appear in the right pane.

Right-click a wrapper name in the Deployment pane and select Copy. You can also select Copy from the Actions drop-down menu.

The Copy dialog appears.

3. Enter a new name for the package wrapper, an optional description, and click Ok.
The new package wrapper appears in the deployment pane.
4. Edit the properties according to your needs and click Save.

Delete a Package Wrapper

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

To delete 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 delete.
The wrapper details appear in the right pane.
3. Right-click a package name in the Manage Resource pane and select Delete. You can also select Delete from the Actions drop-down menu.
A warning message appears.
4. Click Yes to confirm the deletion.
The package wrapper is deleted.

Rename a Package Wrapper

You can rename a package wrapper.

To rename 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 rename.
The wrapper details appear in the right pane.
3. Right-click a package wrapper name in the Manage Resource pane and select Rename. You can also select Rename from the Actions drop-down menu.
The Rename dialog appears.
4. Enter a new name and click Ok.
The package wrapper is renamed.

Create a Deployment Job

To deploy agents to systems, create a deployment job. Deployment jobs contain the details that are required for CA Server Automation to deliver the deployment packages to the appropriate systems at the appropriate time.

Follow these steps:

1. Select Resources, Deployment.

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

2. Right-click the Jobs folder in the Manage Resource pane and select Create New Job. You can also select the Jobs folder and Click + (New) on the Job Status toolbar.

The Jobs wizard appears.

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

The Package Selection page appears.

4. Select a platform and the packages you want to deploy.

5. (Optional) Click the Details tab. The Package Wrapper Details dialog appears and lets you edit the package properties in-line. Package wrappers can be in an incomplete or invalid state and all fields (including port and communities) can be corrected through in-line editing. Click Edit, modify the package wrapper properties, click Save, and then click OK.

The package wrapper properties are updated.

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

The Machine Selection page appears.

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

The Machines Selected page appears.

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

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

The Advanced page appears.

9. Set the distribution server to manage the deployment. If not set, it is automatically chosen.
10. Select the scheduling options for the job; Immediate, Staggered over a specific time period, or Scheduled for a specific time in the future.

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

12. Click Next.

The Summary page appears.

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

The job is created.

Note: You can save the job as a template after you create it. A template saves the package and machine selections so that you can easily reuse them for subsequent jobs.

Specify Read-Write Community Prior To Deployment

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

Specify read-write community prior to deployment

1. Select Resources, Deployment.

2. Open the Deployment pane.

Available deployment groups appear.

3. Expand Packages, the specific package type, and the wrapper platform, and select the wrapper to modify.

The wrapper details appear in the right pane.

4. Specify the read-write parameters in the Read-Write Community field and click Save.

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

More Information:

[Add Server-level SNMP Settings](#) (see page 103)

[Apply Policy to Machines](#) (see page 257)

Specify Read-Write Community Post-Install

To get full SystemEDGE monitoring and management functionality you must specify a valid read-write SNMP community for the SystemEDGE agent. If you have already deployed the SystemEDGE agent, you can add the read/write community string post-installation. You can do this either by creating a global SNMP entry that can be used to monitor and manage multiple systems, or by creating a server-specific SNMP entry.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Traps and Communities tab.

The Communities page appears.

4. Select 'Include only Server-specific SNMP settings'.

5. Click Save Policy.

The policy is saved.

To create a global SNMP entry

1. Click Administration.

The Administration page appears.

2. Click Configuration.

The Configuration page appears.

3. Click SNMP.

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

The Edit SNMP Settings dialog appears.

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

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

The Available Policies page appears.

7. Select the policy in the Available Policies table.

The Summary page for the policy appears.

8. Click the Traps and Communities tab.

The Communities page appears.

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

10. Click Save Policy.

The policy is saved.

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

More Information:

[Add Server-level SNMP Settings](#) (see page 103)

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 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.

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

Follow these steps:

1. Select Resources, Deploy pane.

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

2. Click the Jobs.

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).

Resubmit a Deployment Job

You can resubmit a failed or partially failed deployment job.

To resubmit a deployment job

1. Select Resources, Deployment.

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

2. Click the Jobs folder.

The Job Status pane appears.

3. Click the job you want to resubmit.

The Job Information page appear.

4. Click Actions and select Resubmit.

The Deployment wizard appears at the Package Selection screen.

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

The Machines Selected page appears.

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

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

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

The Advanced page appears.

8. (Optional) Modify the Scheduling options for the job; Immediate, Staggered over a specific time period, or Scheduled for a specific time in the future.

9. (Optional) Select Redeploy previously deployed packages to force all packages (including packages previously successfully deployed) to be deployed, and click Next.

The Summary page appears.

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

The job is resubmitted.

View Deployed Packages

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

To view deployed packages

1. Select Resources.

The Explore pane appears.

2. Select a system or a service.

The Summary page appears.

3. Select Monitoring Software, Deployment.

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

- Task ID
- Job ID
- Package
- Platform
- Wrapper
- Wrapper version
- Stated by
- Start time
- End Time
- Status
- Status extended

View Deployment History

Deployment history information is available from the following places:

Deployment Pane

Displays a count of completed, active, pending, and failed deployment tasks and a summary of successful deployment. Click the top-level Deployment folder to access this view.

Jobs pane

Displays a table of all created deployment jobs that list the job name, 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.

Specific Remote Deployment Use Cases

Deploying/Installing SystemEDGE Agents Using Custom Ports

Deployment of SystemEDGE agents to a non-standard port requires a number of settings to be configured. To ensure the manager can discover and manage the system once it has been deployed, perform the following actions:

1. Update the package wrapper:
 - If you are using the remote deployment solution, you must first configure the package wrapper to specify the port to be used. Navigate to Provisioning, Deployment in the user interface and change the Port field. The Write Community string can also be updated here.

2. Update SNMP Community strings in CA 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.

Reconfigure the SystemEDGE Agent Port

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

Reconfigure the SystemEDGE agent using the control panel

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

The SystemEDGE Setup Wizard opens.

2. Click Next.

The Reinstallation Type page opens.

3. Select Reinstall and click Next.

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

4. Click Next.

The SystemEDGE SNMP Port Number page opens.

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

6. Review the settings, and click Reinstall.

The SystemEDGE agent is reconfigured to use port number 1691.

Reconfigure the SystemEDGE agent port using remote deployment/policy configuration

1. Follow the steps in the chapter [Create a Deployment Job](#) (see page 136). In Step 5 of the wizard, select "Redeploy previously deployed packages".

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

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

2. Do *one* of the following:

Create server-specific SNMP entries for the server:

1. In the CA Server Automation UI, click the Resources tab, open the Explore pane, select the Machine_Name.

The Machine_Name is selected.

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

The SNMP Settings page appears.

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

The New SNMP Settings page appear.

4. Enter the required details, and click Ok.

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

Ensure that global SNMP settings exist and update the policy:

In the CA Server Automation UI, navigate to Administration, SNMP and add a new entry for the required port. When the settings are correct, edit the policy by navigating to the Resources tab, open the Configure pane, and select the policy. Then click Traps & Communities, Communities and select the middle option, "Include server-specific SNMP settings and all Default settings". Save the policy by clicking Save Policy. You should now [apply the policy to the system](#) (see page 257). You can check the community strings in use by the agent using the SystemEDGE control panel applet on the agent machine.

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

More Information:

[Add Server-level SNMP Settings](#) (see page 103)

Remote Deployment to UNIX/Linux Using Non Privileged User Account

If you want to use a nonprivileged user account, consider the following requirements about the sudo configuration:

- Sudo must not enforce that the executed program has a valid pseudo terminal that is attached to it. To disable such validation for a particular user (if it is globally enabled), add the line “Defaults:\$username !requiretty” to the /etc/sudoers file. Replace \$username by the actual username that is used for Remote Deployment.

The standard way to edit the file is using the visudo command. The visudo command invokes \$EDITOR. When editing is finished, it verifies the syntax of the file. If the result is not valid, visudo blocks saving the file.

- Sudo must not ask the user for a password before running the elevated program. To achieve this behavior, the NOPASSWD: keyword must be present on the line giving privileges to the user.
- Sudo must be allowed to run the necessary commands or all. Configuration entries (lines in /etc/sudoers) satisfying the previous requirements are, for example:

```
$username ALL=(ALL) NOPASSWD: ALL
```

or

```
$username ALL = NOPASSWD: /usr/bin/id,/bin/sh /tmp/idprimer/PifInst *
```

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

On Solaris, consider the following requirements for pfexec:

- Any local user can be given profile “Primary Administrator” with the following command
- Any nonlocal user can be given profile “Primary Administrator” by manually adding an entry in the file /etc/user_attr:

```
usermod -P “Primary Administrator” {user}
```

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

Agent Configuration Without Write Community

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

- 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.

Deployment to Windows Vista, Windows 2008 and Windows XP Computers Running Firewall Software

To enable deployment of agents to computers running firewall software, consider the following:

- If the firewall of a target computer running Windows Vista or Windows 2008 operating system is *off* (disabled) and deployment to the computer fails, create or set the following registry variable so that it is a DWORD type with a value 0x1:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\LocalAccountTokenFilterPolicy
```

This is required because User Account Control (UAC) in Windows Vista or Windows 2008 does not automatically grant administrative rights to local users. This occurs even though the local users are members of the Administrators group.

Note: Setting this value will result in remote UAC access token filtering being disabled.

Setting this value is only worthwhile if the user has a local administrator account on the computer running Windows Vista or Windows 2008. Domain administrators will not benefit from this change.

- If the firewall of a target computer running Windows Vista or Windows 2008 is *on* (enabled), the following ports should be opened in addition to file sharing ports, to enable deployment to that computer:

UDP ports

CAM: 4104

File and printer sharing, and so on: 137, 138

TCP ports

IDManager: 135

File and printer sharing, and so on: 139, 445

- If deployment still fails, the following Outbound Rules in the firewall for Windows Vista or Windows 2008 should be fully enabled:
 - Remote Assistance
 - Network Discovery
 - File and Printer Sharing
 - Core Networking
- To enable agent deployment to Windows XP computers that run firewall software you must perform the following actions manually:
 1. Change Security Policy Network Access: Sharing and security model for local accounts from 'Guest only - local users authenticate as Guest' to 'Classic - local users authenticate as themselves'.

The Classic model allows fine control over access to resources and prevents network logons that use local accounts from being mapped to the Guest account, which typically has Read Only access to a given resource.
 2. Configure the following firewall settings:
 - Allow File and Printer Sharing
 - Open UDP Port 4104
 - Open TCP Port 135

Deployment Jobs

To deploy agents to target systems, 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. You can create new jobs using the remote deployment job wizard accessible from several locations. Choose one of the following:

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

When you create a deployment job, you specify the following information:

Job information

Includes the job name and whether to base the job on an existing template.

Deployment package

Includes the platform, the packages to deploy, and the specific wrappers for each package.

Machine information

Includes the systems to which to deploy the packages and system credentials required to establish a connection.

IP Address

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

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

Deployment time

Specifies when to perform the deployment: immediately, staggered over a specific time period, or scheduled for a specific time in the future.

You can also save the job as a template after you create it. A template saves the package and machine selections so that you can easily reuse them for subsequent jobs.

Infrastructure Deployment Process

When executing a deployment, the primary steps of the process are as follows:

1. From the administrator computer, the infrastructure deployment client component issues a request to the IDManager to install an agent on a list of one or more target computers. The deployment manager may be running on a computer that is remote from the client. The list of targets can consist of explicit machine names or IPv4 addresses.

Note: Only discovered resources can be deployed to.

For deployment to succeed to each target computer, it is important to verify that its name, whether entered explicitly or obtained from a container, is suitable for resolving to the address of the target as seen on the deployment manager computer. If, for example, the list of targets retrieved from a directory is not fully qualified with network domain names, deployment may not be able to proceed in certain network configurations.

2. A check is made to see if the IDPrimer is already installed on the target computer. If not, IDPrimer will be installed first on the target computer. The IDManager tries to deliver the IDPrimer installation package. The delivery method used depends on the target operating environment and the security that has been enabled on it. After the IDPrimer image is copied across to the target computer, its installation is initiated.

As some operating systems do not have a method for remote invocation of the IDPrimer installation, in which circumstances the IDPrimer installation may have to be performed manually.

3. The IDPrimer installer installs itself and the CA Messaging (CAM) component on the target computer. Once the IDPrimer is installed and IDManager has received the 'installation complete' signal from the target computer, package deployment can be initiated. An IDManager that has previously installed an IDPrimer and has authenticated with it can deploy packages without needing to resupply user names or passwords. On subsequent deployments, IDPrimer uses asymmetric cryptographic keys to authenticate and limit access to those managers from which we have already gained access.

More Information

[Prerequisites for Automatically Deploying CA Server Automation Infrastructure](#) (see page 151)

[Notes on Infrastructure Deployment Using IPv6 Addresses](#) (see page 153)

[Protocols for Transferring Packages Employed by IDManager](#) (see page 154)

[Manual Installation of the Infrastructure Deployment Primer Software](#) (see page 154)

[Deployment Primer Installation on Windows](#) (see page 154)

[Deployment Primer Installation on Linux or UNIX](#) (see page 155)

[Provide the Deployment Management Certificate to a Primer Installation](#) (see page 155)

[Deployment Management Certificate on Windows](#) (see page 155)

[Deployment Management Certificate on Linux or UNIX](#) (see page 155)

[Compatibility Libraries for Linux](#) (see page 156)

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 146)
- 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 146)

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 151).

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.

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

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

More information:

[Configuration Overview](#) (see page 158)

[How to Apply Policy and Layered Templates to Servers](#) (see page 161)

[How to Create and Apply an Autowatcher to a System](#) (see page 195)

[How to Monitor User-specific Metrics \(MIB Extensions\)](#) (see page 202)

[How to Monitor a Specific Windows Performance Registry Metric](#) (see page 204)

[How to Create SRM Policy](#) (see page 207)

[Discovering the Agents](#) (see page 208)

[Common Usage of Policy Configuration Functions](#) (see page 208)

Configuration Overview

You can configure managed agents and apply the configuration to multiple systems in one operation using centralized Policy Configuration 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 208)

[How to Create SRM Policy](#) (see page 207)

[Apply Policy to Machines](#) (see page 257)

[Review Policy Application Progress](#) (see page 259)

[Configure and View Applied Policies](#) (see page 259)

[Agent Policy Dashboard Views](#) (see page 159)

[How to Monitor User-specific Metrics \(MIB Extensions\)](#) (see page 202)

[How to Monitor a Specific Windows Performance Registry Metric](#) (see page 204)

Agent Policy Dashboard Views

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

Policy Status Summary

Displays a pie chart and list showing the number of policies. A system can be in five different states:

Unconfigured

SystemEDGE agent is installed but no policy is configured.

Agent Installed

SystemEDGE agent is installed.

Configured

SystemEDGE agent is installed and a policy is configured.

Configuration Error

SystemEDGE is installed and a policy is configured but the last configuration failed.

Installed but not managed

SystemEDGE is installed but running in a mode that policy configuration cannot manage.

Policy Breakdown

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


Machines with Non-Standard Policies

Displays systems that contain nonstandard changes to an applied policy.

How to Apply Policy and Layered Templates to Servers

From the CA 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:

Apply Policy and Layered Templates to Servers



Follow these steps:

[Layered Templates Concept](#) (see page 162)

[Create a Policy](#) (see page 164)

[Create Templates for Server Workload](#) (see page 174)

[Apply Policy and Templates to Servers and Verify Settings](#) (see page 191)

[\(Optional\) Manage the Base Policy and Templates for One or More Servers](#) (see page 192)


[\(Optional\) Update the Policy or Templates](#) (see page 190)

[\(Optional\) Apply Policy and Template Updates to Servers and Verify Updates](#) (see page 194)

Layered Templates Concept

In an enterprise, the workload handled by a server or a server group varies. You can create multiple policies specific to the workload handled by a server or a server group. To assist in the creation of policies, templates are used to create application-specific monitors. The Base Policy and Layered Templates are combined to form a configuration file and applied to servers that you want to monitor. You can add or remove Layered Templates. Template updates can be applied directly to servers, without changing the base policy or reimporting the updated template into the Base Policy.

Example: Apply Base Policy and Templates to Servers



You can use Layered Templates in the following scenarios:

Disparate applications

Create a library of templates for each server running a different set of applications. You can directly apply the template updates to each server.

Dynamic environments

The workload of the servers changes frequently in dynamic environments. You can use Layered Templates to segregate the monitors in logical groups. Based on the workload changes, you can directly apply the logical groups to systems or removed from systems.

Shared servers

In an enterprise setup, servers are shared across multiple departments. Each department manages and monitors applications on the shared server. You can use Layered Templates to independently manage and apply templates to systems of each department.

Application maintenance

You can split monitoring into multiple templates. In a server, you can remove a template for an application not in use, without affecting the monitoring of the remaining system.

Out of the box templates

You can apply out of the box templates to managed nodes. Configure the policy with the template configuration on the managed nodes. The following table explains the templates that are available for different OS:

Platform	Name of the Operating System	Template
	All types of operating systems	CPU Utilization - Autowatch Swap Capacity
Windows	—	App Monitoring - CA eTrust Antivirus

Platform	Name of the Operating System	Template
UNIX	AIX	Process Crash
		System Errors
		System Processes
		User Activity
		Windows Services - Autowatch
UNIX	AIX	System Messages
		System Processes
		User Activity
UNIX	HPUnix	System Messages
		System Processes
		User Activity
UNIX	Linux	System Messages
		System Processes
		User Activity
UNIX	Solaris	System Messages
		System Processes
		User Activity

Create a Policy

Create a Base Policy to define a set of Monitors, MIB Extensions, Traps & Communities, and Control Settings to control the agent monitoring.

Common settings in Traps & Communities and Control Settings are available for policies only. If you use Layered Templates, common settings are specified in the Base Policy.

Follow these steps:

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

The Available Policies page appears.

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

The New SystemEDGE Policy dialog appears.

3. Enter a name and an optional description for the policy, the system type and whether to base it on an existing policy and click Ok.

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

4. Click Save Policy.

The policy is created and saved.

Note: You can also use the existing default policy as a Base Policy, if necessary.

More Information:

[Copy SystemEDGE Policy](#) (see page 209)

[Rename SystemEDGE Policy](#) (see page 209)

[Delete SystemEDGE Policy](#) (see page 210)

Define SystemEDGE Policy Control Settings

You can control the following agent behavior using the SystemEDGE policy control settings:

- Security settings
- SNMP settings
- MIB table population
- UNIX settings
- Performance monitoring settings

You can segregate these common control settings from specific server workload configurations by adding them to the Base Policy.

You can apply the control settings defined in the policy to all systems you want to monitor with this configuration.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Control Settings.

The Controls page appears.

4. (Optional) Click Use Defaults.

The default selections pane appears. You can change the default settings.

5. Configure the following control settings:

SNMP

Lets you define the following basic SNMP properties:

Bind Address

Specifies an interface to which the agent binds and listens for incoming SNMP requests. Valid addresses are IPv4 or IPv6 address.

Note: The corresponding default _port is specified during installation.

Bind Port

Specifies the trap port the agent binds to for sending SNMP traps. If no bind_address is specified, the agent binds to all available UDP addresses.

Default: Port selected by the system

IP Family

Specifies the agent communication method: IPv4 only, IPv6 only, or both. By default, the agent tries using IPv4 and then IPv6.

FIPS Mode

Specifies the agent to use FIPS-compliant encryption. Select Non-FIPS Mode to enable the CA eTrust Public Key Infrastructure libraries, and if this method fails, fall back to the internal minimum security solution. Select FIPS Co-existence Mode to enable FIPS-compliant encryption, and if this method fails, fall back to the CA eTrust Public Key Infrastructure Libraries. If they fail, select FIPS Only Mode to enable the RSA BSAFE Crypto-C Micro Edition FIPS-compliant libraries and perform no encryption.

Default: Non-FIPS Mode

Trap Source

Specifies the source address used to send traps. Valid addresses are IPv4, IPv6 address, or a host name.

Default: Host name of the agent

Security Settings

Lets you define the following security preferences:

Authentication Traps

Sends an authentication failure trap when the agent receives an SNMP message with a community name that the agent cannot recognize.

Default: Disabled

Process Sets

Permits access to processes and other software running on agent systems in the Process table and Running Software table. Allowing SNMP Sets on these tables can cause security issues.

Remote Shell Group

Permits management systems to instruct the agent remotely to run shell scripts and programs on the agent system through the Remote Shell group. The disclosure of this type of information can post a potential security risk.

Execution Action

Enables the execution of action commands with the monitoring tables when a threshold breach occurs. The capability to run action commands and scripts can be a security issue.

MIB Table Population

Populates the following tables in the Systems Management MIB:

- Process Table
- User Group Table
- Who Table
- Trap Community Table
- Monitor Mirror Table
- Aggregate Mirror Table
- Top Processes Table

Each table either contains sensitive information that you can expose in a MIB or nonessential information that you can disable to save disk space. The default settings enable population of all tables except for the process table.

Miscellaneous

Lets you define the following miscellaneous settings:

Allow agent to be Updated using SNMP

Permits agent updates using SNMP Sets (for example, removes write communities). If you permit SNMP Sets on the agent, any updates through this method cause a notification of an SNMP Set change. These updates also cause an exception when viewing policy details for the system.

Notify Manager of Configuration Updates

Enables the agent to send a notification to the manager for any SNMP Set request that the agent processes.

Warm Start Discovery

Enables an agent rediscovery of all devices after every warm start configuration update. If you manage a system with many devices, a discovery after every warm start can consume too much time and too many resources.

Use Perl Compatible Regular Expressions

Perl Compatible Regular Expressions (PCRE) enables you to specify i18n compatible regular expressions while defining monitors that support regular expressions. The examples of regular expressions are log file, process, process group, Windows services and Windows events. You can also use this option to create more complex regular expressions. This option is provided in SystemEDGE agent 5.1.0 and above versions.

Automatically Resolve Index Conflicts

Enables you to resolve Index conflicts. When you apply the layered templates to all systems, indexes are assigned to the monitors added in the template. If the assigned indexes conflict with existing indexes either within the base policy or another template, this option reassigns unique index values.

Note: Indexes contained within the base policy are always maintained in the delivered configuration. If this option is disabled, you cannot resolve conflicting indexes. However, when you apply layered templates to the systems, the conflicting indexes are displayed as errors on the layered templates that caused the conflicting indexes.

Historical Performance Monitoring

Lets you define the following settings for the Performance Cube AIM, which collects history information into Systems Performance cubes for historical performance management:

Collection Interval

Specifies how often to collect information from the History table into performance cubes.

Index Range Start

Specifies the beginning of the reserved range of indexes, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

Index Range End

Specifies the end of the reserved range of indexes, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

UNIX Control Settings

Lets you define the following settings for agents running on UNIX systems:

Sub-program Group

Specifies a group name other than root under which to run subprograms.

Sub-program User

Specifies a user name other than root under which to run subprograms.

Linux Freemem Include

Specifies whether to include system buffers, disk cached memory, or both in free memory calculation.

Query System Devices

Lets you enable querying of the following system device metrics:

- Serial device status
- Floppy disk status
- Disk size, capacity, description, and other properties (Probe Disks)
- NFS file system status
- HP-UX graphics status

Querying these metrics can cause issues with potential agent blocking. The default settings enable querying of only serial device status and NFS file system status.

6. Click Plugins.

The Plugins pane appears. This pane controls which AIMs to load with the agent.

7. Do one of the following:

- Select 'Load all available plugins' to load all AIMs available on the agent system.
- Select 'Load plugins selected in the table'.
- Click + (New) on the External Plugins toolbar to add an AIM to the External Plugins table.

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

AIM loading is configured.

8. Click Aggregate Monitors.

Configure aggregate monitors as described in [Configure Object Aggregation](#) (see page 170).

The control settings are defined.

9. Click Save Policy.

The policy is saved.

More Information:

[Configure Object Aggregation](#) (see page 216)

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 211)

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 175). To add MIB extensions to the template, see section [Define MIB Extensions](#) (see page 187).

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 176)
- [Create a Process Monitor](#) (see page 178)
- [Create a Log File Monitor](#) (see page 180)
- [Create a Windows Event Monitor](#) (see page 182)
- [Create a History Monitor](#) (see page 183)
- [Create a Process Group Monitor](#) (see page 185)

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

5. Click Save.

The monitor is loaded to the policy or the template.

More Information:

[Define a Threshold Monitor](#) (see page 227)

[Define a Process Monitor](#) (see page 229)

[Define a Log File Monitor](#) (see page 231)

[Define a Windows Event Monitor](#) (see page 233)

[Define a History Monitor](#) (see page 234)

[Define a Process Group Monitor](#) (see page 236)

Create a Threshold Monitor

Create a threshold monitor that lets the agent monitor the servers or the server groups against specified thresholds. The agent sends a trap when thresholds are breached.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Threshold.

The Threshold Monitors page appears.

5. Click + (New) on the Threshold Monitors toolbar.

The Threshold Monitor Details: New dialog appears.

6. Configure the following threshold settings:

Index

Defines the table index that you want to use.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class

Specifies the object class to monitor. The values refer to the available MIB tables.

Object Class Name

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

Object Attribute

Specifies the object attribute to monitor. The values refer to the available attributes of the table selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this threshold monitor.

Object Attribute Name

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

Object Instance

Specifies the object instance to monitor. This value, for example, .3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this threshold monitor. For some object classes, the name of the instance itself can be given (for example, C: instead of .3, or /var for a Unix machine).

Object Instance Name

Defines the object instance name to use for the object state model. Value is an arbitrary string, for example, SysVol_C.

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

Severity

Specifies the severity to use for the object state model.

Operator

Specifies the operator to use.

Value

Defines the value to use.

Sample Type

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Threshold Monitor settings are saved.

8. Click Save Template.

The Threshold Monitor is loaded to the template.

Create a Process Monitor

Create a process monitor that lets the agent monitor a process, service, or process table objects against specified thresholds. The agent sends a trap when thresholds are breached or the state of a process (running or stopped) changes.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Process.

The Process Monitors page appears.

5. Click + (New) on the Process Monitors toolbar.

The Process Monitor Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index that you want to use.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class Name

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

Object Attribute

Specifies the object attribute to monitor. The values define the available attributes for process monitoring.

Object Attribute Name

Defines the object attribute name to use for the object state model. Value is an arbitrary string, for example, MemUsedPercent.

Object Instance

Specifies the object instance to monitor. This is the regular expression (dependent from optional settings) to use for matching processes by name, or Windows services by name. Pattern should uniquely match a single process (service). Arguments can be included (see optional settings).

Object Instance Name

Specifies the object instance name to use for the object state model. Value is an arbitrary string, for example, ApacheServer.

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

Severity

Specifies the severity to use for the object state model.

Operator

Specifies the operator to use.

Value

Defines the value to use.

Sample Type

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Process Monitor settings are saved.

8. Click Save Template.

The Process Monitor is loaded to the Policy.

Create a Log File Monitor

Create a log file monitor that lets the agent monitor any UTF-8 encoded system or application log file by searching for strings specified as regular expressions. The agent sends a trap when a match occurs.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Log File.

The Log File Monitors page appears.

5. Click + (New) on the Log File Monitors toolbar.

The Log File Monitor Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index that you want to use.

Monitor Type

Specifies the monitor type that you want to use.

Platform

Specifies the platform.

Description

Defines an optional description.

Log File/Directory Name

Defines the path to the file or the directory to monitor.

Search Filter

Specifies the search filter.

Interval

Defines the evaluation interval for the monitor in minutes

Severity

Specifies the significance of the monitor on a match.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings page lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Log File Monitor settings are saved.

8. Click Save Template.

The Log File Monitor is loaded to the Policy.

Create a Windows Event Monitor

Create a windows event monitor that lets the agent monitor the Windows event log entries using different filters (event source). The agent sends a trap when a match occurs.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click Windows Event.

The Windows Event Monitors page appears.

5. Click + (New) on the Windows Event Monitors toolbar.

The Windows Event Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index that you want to use.

Platform

Specifies the platform.

Description

Defines an optional description.

Event Log

Specifies the event log to read.

Event Type

Specifies the event type to match.

Source Filter

Defines the source filter to use.

Description Filter

Defines the description filter to use.

Severity

Specifies the significance of the monitor on a match.

The Maintenance Window subtab lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Windows Event Monitor settings are saved.

8. Click Save Template.

The Windows Event Monitor is loaded to the Policy.

Create a History Monitor

Create a history monitor that lets the agent provide the historical data collection for manager-side baseline and trend analysis. The agent uses the metrics to provide a picture of average system performance during a specific time interval.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.

The Template List page appears.

2. Select the template in the Template List.

The Summary page for the template appears.

3. Click the Monitors tab.

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

4. Click History.

The History Monitors page appears.

5. Click + (New) on the History Monitors toolbar.

The Historical Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index that you want to use.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class

Specifies the object to monitor. The values refer to the available MIB table values.

Object Attribute

Specifies the object attribute to monitor. The values refer to the available attributes of the table selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this History entry.

Object Instance

Defines the object instance to monitor. This value (for example, 0.3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this History entry.

Interval

Defines the collection interval in a multiple of 30 seconds.

Buckets

Defines the number of samples to collect.

Add to Performance Cube check box

Specifies whether to collect performance cube data for this entry.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

7. Click Save
The History Monitor settings are saved.
8. Click Save Template.
The History Monitor is loaded to the Policy.

Create a Process Group Monitor

Create a process group monitor that lets the agent define a group of processes and monitors that group for changes. If the process group changes, the agent sends a trap.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates and appropriate subcategories.
The Template List page appears.
2. Select the template in the Template List.
The Summary page for the template appears.
3. Click the Monitors tab.
The Summary page appears with a list of monitors managed by the policy.
4. Click Process Group.
The History Monitors page appears.
5. Click + (New) on the Process Group Monitors toolbar.
The Process Group Details: New dialog appears.
6. Configure the following process settings:

Index

Defines the table index that you want to use.

Platform

Specifies the platform.

Description

Defines an optional description.

Process Name

Defines the process name. This is the regular expression (dependent from optional settings) to use for matching processes by name.

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

User Name

Defines the user name to match in addition to any process name regular expression.

Group Name

Defines the group name to match in addition to any process name regular expression.

Severity

Specifies the significance of the monitor on a group change

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings page lets you define the flags that can be used for the different monitor entries or history control entries.

7. Click Save

The Process Group Monitor settings are saved.

8. Click Save Template.

The Process Group Monitor is loaded to the Policy.

Define MIB Extensions

Defining MIB extensions provide functional benefits that are not available in local file manipulation. The policy configuration feature provides field names and the list of key properties such as object type.

When you configure a policy or a monitoring template, click the MIB Extensions tab to add the following objects:

- MIB Extensions
- Windows Performance
- Windows Registry

Note: To add MIB Extensions to a template or a policy, see [Add MIB Extensions to a Template or a Policy](#) (see page 187). MIB Extensions within templates are supported for the purposes of applying the MIB Extensions directly to monitored systems. MIB Extensions for use within policies should be created directly in the Policy itself.

Add MIB Extensions to a Template or a Policy

Define MIB extensions for a template or policy using the policy configuration feature.

Follow these steps:

1. Click the Resources tab, open the Configure pane, and expand Monitoring Templates or Policies.
2. From the Templates List or Available Policies page, click the template or policy name.
The Summary page appears.
3. Click the MIB Extensions tab.
The MIB Extensions page appears.
4. Define the MIB extension attribute using the following controls and click Add:

Index

Defines the attribute leaf number.

Type

Specifies the attribute type.

Extension Command

Defines the full path or the name (including parameters) of the script or binary to execute.

Access Rights

Specifies the attributes access rights.

5. Click the Windows Performance tab.

The Windows Performance pane appears.

6. Define the Windows Performance attributes using the following controls and click Add:

Index

Defines the attribute leaf number.

Type

Specifies the attribute type.

Object

Specifies the performance registry object.

Counter

Specifies the performance registry counter.

Instance

Defines the performance registry instance.

7. Click the Windows Registry tab.

The Windows Registry pane appears.

8. Define Windows Registry attribute using the following controls and click Add:

Index

Defines the attribute leaf number.

Type

Specifies the attribute type.

Key

Defines the registry key in HKEY_LOCAL_MACHINE.

Value

Defines the attribute value.

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

9. Click Save Template or Policy.

The configuration is saved.

(Optional) Reindex Monitors from Templates or a Policy

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

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

To reindex the monitors, consider the following:

- Verify that the monitors exist.

Follow these steps:

1. Click the Resources tab, open the Configure pane, and expand Monitoring Templates or Policies.
2. From the Templates List or Available Policies page, click the template or policy name.

The Summary page appears.

3. Click the Monitors tab.

The Summary page appears with a list of monitors.

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

The new base index dialog appears.

5. Enter a numeric value as the base index

Example:1000

6. Select Make indexes contiguous

Make indexes contiguous

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

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

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

Example: 1001, 1010, 1020, 1030 and so on.

7. Click Ok to confirm the reindex.

The monitors are reindexed.

Delete Monitors from Templates or a Policy

You can delete a monitor from a policy or template.

Follow these steps:

1. Click the Resources tab, open the Configure pane, and expand Monitoring Templates or Policies.
2. From the Templates List or Available Policies page, click the template or policy name.
The Summary page appears.
3. Click the Monitors tab.
The Summary page appears with a list of monitors managed by the policy.
4. Click the appropriate monitor tab and select one or multiple monitors you want to delete.
5. Click Action and select Delete.
A warning message appears.
6. Click Ok to confirm the deletion.
7. (Optional) Repeat the process for any additional monitors.
8. Click Save Policy.
The monitor is deleted from the Policy.

Note: You cannot delete a template, or the policy with the template which is used by a server or a server group.

(Optional) Update the Policy or Templates

If necessary, you can update the existing policy or templates by adding or deleting monitors from the policy or the template. The update procedures are similar to the creation process.

Follow these steps:

1. Add or delete monitors that are specific to the server workload. To add monitors to the template or the Policy, see [Add Monitors to the Template or the Policy](#) (see page 175). To delete monitors from the Policy, see [Delete Monitors from Templates or a Policy](#) (see page 190).
 2. [Define MIB Extensions](#) (see page 187).
 3. [Define SystemEDGE Policy Control Settings](#) (see page 165).
- The policy or templates are updated.

Apply Policy and Templates to Servers and Verify Settings

After you create the template, you can apply the policy, with the template directly to the servers or the server groups across the enterprise.

Follow these steps:

1. Select the policy in the Available Policies table or select a template from the Template List.

The Summary page for the policy or template appears.

2. Select the Managed Machines tab.

The list of managed machines appears.

3. Click Action and select Apply.

Tabs appear for selecting systems on which to apply the policy.

Update machines running this policy/template

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

Apply to machines not running this policy/templates

Lets you apply policy or a template to systems.

4. (Options for policies) Do one of the following options from the 'Update machines running this policy' tab:

- Select 'Update all machines using this policy' to deploy the policy on all machines currently running it. This option is useful if you have made configuration policy changes that you want to apply globally.
- Select 'Update selected groups of machines' to update only machines that meet any of the following criteria:

- Machines running an out-of-date version of the policy
- Machines where policy exceptions have been applied
- Machines running current version of the policy
- Machines with configuration errors for this policy

Policy exceptions occur when a user applies a point configuration change to an agent that is not represented in the applied policy.

- Select 'Advanced (manually select machines)' to add the machines manually in the Select Machines pane to which you want to reapply the policy.

5. (Options for templates) Select one of the following options from the 'Update machines running this template' tab:
Under Existing Machines, select one of the following options:
 - Update all machines with this template applied.
 - Update only those machines that do not have the latest changes of this template applied.
 - Update only those machines where the template has not been successfully applied.
 - Advanced (manually select machines)
 - Remove this template from machines.
6. (Optional) Select systems from the 'Apply to Machines not running this policy/template' tab for applying the policy or template.
7. Click Apply Policy or Apply Template.
The application is initiated.
8. Verify if the servers behave as expected. If necessary, you can update and apply the updated policies and templates.

(Optional) Manage the Base Policy and Templates for One or More Servers

Manage the templates and the base policy for a single 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 175). To delete monitors from the Policy, see [Delete Monitors from Templates or a Policy](#) (see page 190).
2. [Define MIB Extensions](#) (see page 187).
3. [Define SystemEDGE Policy Control Settings](#) (see page 165).

The policy or templates are updated.

(Optional) Apply Policy and Template Updates to Servers and Verify Updates

After you update the template, apply the template updates directly to servers or server groups across the enterprise.

Follow these steps:

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

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

2. Select the Template Name.

The Summary page appears with the Template information.

3. Click Action and select Apply.

Tabs appear for selecting machines on which to apply the monitoring template. The 'Update machines running this template' tab lets you apply the monitoring template to machines that are already using the template. The 'Apply to Machines not running this template' tab lets you apply the monitoring template to machines without using any template.

4. (Optional) Under Existing Machines, select machines from the 'Update machines running this template' tab options.
5. (Optional) Under Selected Machines, select the machines to which the template is re-applied.
6. (Optional) Select machines from the 'Apply to the Machines not running this Template' tab to apply the template.
7. Click Apply.

The template application is initiated and the view Status link appears.

8. Click View Status link to verify whether the SystemEDGE monitoring template updates are applied to servers.

The page appears with the list of servers to which the SystemEDGE monitoring template updates are applied.

The Layered Template Updates have successfully been applied to the servers or the server groups.

9. Verify if the servers behave as expected. If necessary, you can update and apply the updated policies and templates again.


How to Create and Apply an Autowatcher to a System

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

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

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

How to Create and Apply an Autowatcher to a System



Follow these steps:

[Review Requirements](#) (see page 196)

[How Autowatchers Work](#) (see page 196)

[Apply Predefined Autowatchers](#) (see page 199)

[Create and Apply an Autowatcher to a System](#) (see page 200)

[Verify the Autowatcher](#) (see page 201)

Review Requirements

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

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

More information:

[Apply Predefined Autowatchers](#) (see page 199)

[Create and Apply an Autowatcher to a System](#) (see page 200)

How Autowatchers Work

Autowatchers run periodic discovery processes using regular expressions as patterns to match the names of resources for which the Autowatchers create monitors.

Autowatchers enable SystemEDGE to create automatically monitors for new resources when they come online. Autowatchers create monitors in a reserved range of indexes (1000000 - 1999999).

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


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

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

You can use the following Autowatcher Types:

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

Process Workflow of an Autowatcher



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

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

More information:

[Generic Autowatchers](#) (see page 198)

[Process and Service Autowatchers](#) (see page 198)

Generic Autowatchers

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

The following list provides you some examples of Generic Autowatchers:

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

More information:

[How Autowatchers Work](#) (see page 196)

Process and Service Autowatchers

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

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

A Process Autowatchers create process monitors in two ways:

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

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

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

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

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

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

More information:

[How Autowatchers Work](#) (see page 196)

Apply Predefined Autowatchers

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

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

Verify, if you can use predefined Autowatchers.

Follow these steps:

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

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

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

The Autowatchers Details pane opens.

3. Click Action, Apply.

The machine selection page opens.

4. Select the appropriate systems and click Apply.

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

SystemEDGE creates monitors automatically based on the Autowatcher settings.

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

More information:

[Create and Apply an Autowatcher to a System](#) (see page 200)

Create and Apply an Autowatcher to a System

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

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies or Monitoring Templates, and click SystemEDGE.
The Available Policies pane or Template List opens.
2. Open a policy or template and click Autowatchers.
The Generic Autowatchers pane opens.
3. Select the Process/Service tab if you want to add Process or Service Autowatchers.
4. Click  (Add) on the toolbar.
The Autowatcher Details pane opens.
5. Specify the required values and click Save.
The Autowatcher is saved.
6. Click Action, Apply.
The machine selection page opens.

7. Select the appropriate systems and click Apply.

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

SystemEDGE creates monitors automatically based on the Autowatcher settings.

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

More information:

[Verify the Autowatcher](#) (see page 201)

Verify the Autowatcher

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

Follow these steps:

1. Click the Resources tab.

The Resources page appears.

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

The discovered and managed resources in the data center appear.

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

The Quick Start tasks for the selected resource appears.

4. Click the Configuration tab.

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

More information:

[How Autowatchers Work](#) (see page 196)

How to Monitor User-specific Metrics (MIB Extensions)

This step-by-step example describes how to monitor a user-specific metric.

How to monitor user-specific metrics (MIB extensions)

1. Create a program that returns the data required. For example, a simple DOS batch script on the agent system to return some fixed data.

```
@echo off  
echo 99
```

2. Open a text editor and store these two lines in data.bat on the C: drive.
3. Create an MIB extension that references this batch file.

- a. From the user interface, click *Policy*, open *Configuration* in the navigation pane, expand the Policy tree, and open a SystemEDGE policy.

The policy details appear in the right pane.

- b. Click the MIB Extensions tab.

The MIB Extensions pane opens.

- c. Add the following data into the fields:

Index: 1 (if it is the first MIB extension)

Type: integer

Extension Command: C:\data.bat

Access Rights: Read Only

- d. Click *Add*.

The MIB Extension is added to the Policy.

- e. Click *Save Policy*.

The policy is saved.

4. Create a threshold monitor to check the value of the new monitor.

- a. Click *Monitors* and then *Thresholds*.

The Threshold Monitor Details Edit pane appears.

- b. Add the following data into the fields:

Index: (automatically added)

Platform: OS Independent

Object Class: extensionGroup [Extended mib from adding new scalar variables]

Object Attribute: 1

Object Instance Name: MyData

Interval: 60

Severity: Major Alarm

Operator: greater than or equal to

Value: 50

Scale: 1

Sample Type: absolute value

- c. Click *Save*.

The policy is saved. A 'major' alarm with threshold '50' is added. This threshold will be breached immediately as the script created previously always returns the value '99'.

- d. Click *Action* and then *Apply*, to apply the policy to a computer.

The Selected Machines pane appears.

- e. Verify that the selected machines are correct and click *Apply*.

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

Click *Return to Policy*.

The policy details pane appears.

Once the agent is configured, you can view the state of this threshold monitor from the Resources tab. You can see that the "major" threshold has been breached.

How to Monitor a Specific Windows Performance Registry Metric

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

How to monitor user-specific metrics (MIB extensions):

1. Create a MIB extension for a Windows Performance Registry metric.
 - a. From the user interface, click the Resources tab, open Configuration pane, expand the Policy tree, and click an appropriate subcategory.
The policy details appear in the right pane.
 - b. Click the *MIB Extensions* tab.
The MIB Extensions pane opens.
 - c. Click *Windows Performance*.
The Windows Performance Defined Extensions pane appears.
 - d. In the Add the following data into the fields:
Example:
Index: 1 (If the extension is the first one).
Type: integer
Object: System
Counter: Processes (Provides the total number of running processes).
The System metrics have no 'instance' so this field is left blank.
Note: You can specify custom entries for Object and Counter while creating a policy. The same metrics are saved for future use while creating another policy.
 - e. Click *Add*.
The MIB Extension is added to the Policy.
 - f. Click *Save Policy*.
The policy is saved.
2. Create a threshold monitor to check the value of the new monitor.
 - a. Click *Monitors* and then *Thresholds*.
The Threshold Monitor Details Edit pane appears.
 - b. Click + (New) to create a monitor.
The Threshold Monitor Details: New dialog appears.

- c. Configure the following threshold settings:

Index

Defines the table index that you want to use.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class

Specifies the object class to monitor. The values refer to the available MIB tables.

Object Class Name

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

Object Attribute

Specifies the object attribute to monitor. The values refer to the available attributes of the table; selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this threshold monitor.

Object Attribute Name

Defines the object attribute name to use for the object state model as an arbitrary string, for example, PercentUsed.

Object Instance

Specifies the object instance to monitor. This value, for example, .3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this threshold monitor. For some object classes, the name of the instance itself can be given (for example, C: instead of .3, or /var for a Unix machine).

Object Instance Name

Defines the object instance name to use for the object state model. Value is an arbitrary string, for example, SysVol_C.

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

Severity

Specifies the severity to use for the object state model.

Operator

Specifies the operator to use.

Value

Defines the value to use.

Sample Type

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

- d. Click *Save*.

The monitor is added to the policy.

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

The Selected Machines pane appears.

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

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

- b. Click *Return to Policy*.

The policy details pane appears.

Once the agent is configured, you can view the state of this threshold monitor from the Resources tab under Explore, Summary.

How to Create SRM Policy

You create SRM policy to define tests to perform, thresholds to monitor, configuration preferences, and other settings that control how the agent runs and what it monitors. Once you create a policy, you can apply it to any number of systems running SystemEDGE agents with the SRM AIM in managed mode. Policy lets you perform all configuration operations that you can manually configure locally with the benefit of a consolidated interface, pick lists, and dynamic deployment to remote systems.

The following process describes how to create SRM policy:

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

The Service Response pane appears.

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

The New Service Response Monitoring Policy dialog appears.

3. Enter a name and description for the policy and whether to base it on an existing policy and click OK.

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

4. Define tests to include.
5. Define test thresholds.
6. [Define control settings](#) (see page 249).
7. Click Save Policy.

The policy is saved.

Discovering the Agents

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

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

Follow these steps:

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

A confirmation dialog opens.

3. Click OK.

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

Common Usage of Policy Configuration Functions

This section describes the common Policy Configuration functions.

How to Create SystemEDGE Policy

You create SystemEDGE policy to define a set of monitors, AIMs to load, configuration preferences, and other settings that control how the agent runs and what it monitors. Once you create a policy, you can apply it to any number of systems running SystemEDGE agents in managed mode. Policy lets you perform all configuration operations that you can manually configure locally with the benefit of a consolidated interface, pick lists, and dynamic deployment to remote systems.

The following process describes how to create SystemEDGE policy:

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

The SystemEDGE pane appears.

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

The New SystemEDGE Policy dialog appears.

3. Enter a name and description for the policy and whether to base it on an existing policy and click OK.

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

4. Define monitors to include.
5. [Define control settings](#) (see page 211).
6. [Define SNMP settings](#) (see page 171).
7. Define MIB extensions.
8. Click Save Policy.
The policy is saved.

Copy SystemEDGE Policy

You can copy an existing SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy you want to copy in the Available Policies table, click Actions and select Copy. You can also right-click the policy in the Configure pane and select Copy.
The Copy dialog appears.
3. Enter a new name for the policy and click Ok.
The policy is copied and a configuration screen appears in the right pane.
4. Click Save Policy.
The policy is saved.

Rename SystemEDGE Policy

You can rename an existing SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy you want to rename in the Available Policies table, click Actions and select Rename. You can also right-click the policy in the Configure pane and select Rename.
The Rename dialog appears.
Note: If the policy is in use, an error message is displayed indicating that the policy cannot be renamed.

3. Enter a new name for the policy and click Ok.
A confirmation message appears notifying you that the policy is renamed.
4. Click Save Policy.
The policy is saved.

Delete SystemEDGE Policy

You can delete an existing SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy you want to delete in the Available Policies table, click Actions, and select Delete. You can also right-click the policy in the Configure pane and select Delete.

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

A warning message appears.

3. Click Ok to confirm the deletion.
A confirmation message appears. The policy is deleted.

Import a SystemEDGE Configuration to a Policy

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

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Click + (New) on the Available Policies toolbar.
The New SystemEDGE Policy dialog appears.

3. Click Import.

The SystemEDGE Agent Machines window appears.

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

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

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

The policy is saved.

Define SystemEDGE Policy Control Settings

You can control the following agent behavior using the SystemEDGE policy control settings:

- Security settings
- SNMP settings
- MIB table population
- UNIX settings
- Performance monitoring settings

You can apply the control settings defined in the policy to all machines.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Control Settings.

The Controls page appears.

4. (Optional) Click Use Defaults.

The default selections pane appears. You can change the default settings.

5. Configure the following control settings:

SNMP

Lets you define the following basic SNMP properties:

Bind Address

Specifies an interface to which the agent binds and listens for incoming SNMP requests. Valid addresses are IPv4 or IPv6 address.

Note: The corresponding default _port is specified during installation.

Bind Port

Specifies the trap port the agent binds to for sending SNMP traps. If no bind_address is specified the agent binds to all available UDP addresses.

Default: Port selected by the system

IP Family

Specifies the agent communication method: IPv4 only, IPv6 only, or both. By default, the agent tries using IPv4 and then IPv6.

FIPS Mode

Specifies the agent to use FIPS-compliant encryption. Select Non-FIPS Mode to enable the CA eTrust Public Key Infrastructure libraries, and if this method fails, fall back to the internal minimum security solution. Select FIPS Co-existence Mode to enable FIPS-compliant encryption, and if this method fails, fall back to the CA eTrust Public Key Infrastructure Libraries. Select FIPS Only Mode to enable the RSA BSAFE Crypto-C Micro Edition FIPS-compliant libraries and perform no encryption if they fail.

Default: Non-FIPS Mode

Trap Source

Specifies the source address used to send traps. Valid addresses are IPv4, IPv6 address, or a host name.

Default: Host name of the agent

Security Settings

Lets you define the following security preferences:

Authentication Traps

Sends an authentication failure trap when the agent receives an SNMP message with a community name that the agent cannot recognize.

Default: Disabled

Process Sets

Permits access to processes and other software running on agent systems in the Process table and Running Software table. Allowing SNMP Sets on these tables can cause security issues.

Remote Shell Group

Permits management systems to remotely instruct the agent to run shell scripts and programs on the agent system through the Remote Shell group. The disclosure of this type of information can post a potential security risk.

Execution Action

Enables the execution of action commands with the monitoring tables when a threshold breach occurs. The capability to run action commands and scripts can be a security issue.

MIB Table Population

Populates the following tables in the Systems Management MIB:

- Process Table
- User Group Table
- Who Table
- Trap Community Table
- Monitor Mirror Table
- Aggregate Mirror Table
- Top Processes Table

Each table either contains sensitive information that you can expose in a MIB or non-essential information that you can disable to save disk space. The default settings enable population of all tables except for the process table.

Miscellaneous

Lets you define the following miscellaneous settings:

Allow agent to be Updated using SNMP

Permits agent updates using SNMP Sets (for example, removes write communities). If you permit SNMP Sets on the agent, any updates through this method cause a notification of an SNMP Set change and also an exception when viewing policy details for the system.

Notify Manager of Configuration Updates

Enables the agent to send a notification to the manager for any SNMP Set request that the agent processes.

Warm Start Discovery

Enables an agent rediscovery of all devices after every warm start configuration update. If you manage a system with many devices, a discovery after every warm start can consume too much time and too many resources.

Use Perl Compatible Regular Expressions

Perl Compatible Regular Expressions (PCRE) enables you to specify i18n compatible regular expressions while defining monitors that support regular expressions. The examples of regular expressions are log file, process, process group, Windows services and Windows events. You can also use this option to create more complex regular expressions. This option is provided in SystemEDGE agent 5.1.0 and above versions.

Automatically Resolve Index Conflicts

Enables you to resolve Index conflicts. When you apply the layered templates to all machines, indexes are assigned to the monitors added in the template. If the assigned indexes conflict with existing indexes either within the base policy or another template, this option reassigns unique index values.

Note: Indexes contained within the base policy are always maintained in the delivered configuration. If this option is disabled, you cannot resolve conflicting indexes. However, when you apply layered templates to the machines, the conflicting indexes are displayed as errors on the layered templates that caused the conflicting indexes.

Historical Performance Monitoring

Lets you define the following settings for the Performance Cube AIM, which collects history information into Systems Performance cubes for historical performance management:

Collection Interval

Specifies how often to collect information from the History table into performance cubes.

Index Range Start

Specifies the beginning of the reserved range of indices, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

Index Range End

Specifies the end of the reserved range of indices, where the agent per default creates history control entries for collection of performance cube data. This reserved range is used, for example, if SRM (Service Response Monitoring) is configured to collect performance data.

UNIX Control Settings

Lets you define the following settings for agents running on UNIX systems:

Sub-program Group

Specifies a group name other than root under which to run subprograms.

Sub-program User

Specifies a user name other than root under which to run subprograms.

Linux Freemem Include

Specifies whether to include system buffers, disk cached memory, or both in free memory calculation.

Query System Devices

Lets you enable querying of the following system device metrics:

- Serial device status
- Floppy disk status
- Disk size, capacity, description, and other properties (Probe Disks)
- NFS file system status
- HP-UX graphics status

Querying these metrics can cause issues with potential agent blocking. The default settings enable querying of only serial device status and NFS file system status.

6. Click Plugins.

The Plugins pane appears. This pane controls which AIMs to load with the agent.

7. Do one of the following:

- Select 'Load all available plugins' to load all AIMs available on the agent system.
- Select 'Load plugins selected in the table'.
- Click + (New) on the External Plugins toolbar to add an AIM to the External Plugins table.

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

AIM loading is configured.

8. Click Aggregate Monitors.

Configure aggregate monitors as described in [Configure Object Aggregation](#) (see page 216).

The control settings are defined.

9. Click Save Policy.

The policy is saved.

More Information:

[Configure Object Aggregation](#) (see page 216)

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 211)

Define New SystemEDGE Monitoring Template

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

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

Follow these steps:

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

The SystemEDGE page appears

2. Click + (New) on the Template List toolbar.

The New SystemEDGE Monitoring Template dialog appears.

3. Enter a name and an optional description for the template, the system type and whether to base it on an existing template and click Ok.

The template is created, and the Summary page appears. To add a monitor to the template, see the section [Add a Monitor To SystemEDGE Policy](#). (see page 226)

4. Click Save Template.

The template is saved.

More Information:

[Layered Templates](#) (see page 219)

[Import a Monitoring Template to SystemEDGE Policy](#) (see page 221)

[Copy SystemEDGE Monitoring Template](#) (see page 222)

[Modify SystemEDGE Monitoring Template](#) (see page 222)

[Rename SystemEDGE Monitoring Template](#) (see page 223)

[Delete SystemEDGE Monitoring Template](#) (see page 223)

[Review Monitoring Template Application Progress](#) (see page 224)

[Apply Templates to Machines](#) (see page 224)

[Layered Templates](#) (see page 219)

[Modify SystemEDGE Monitoring Template](#) (see page 222)

[Rename SystemEDGE Monitoring Template](#) (see page 223)


[Review Monitoring Template Application Progress](#) (see page 224)

[Apply Templates to Machines](#) (see page 224)

Layered Templates

In an enterprise, the workload handled by a server or a server group varies. You can create multiple policies specific to the workload handled by a server or a server group. To assist in the creation of policies, templates are used to create application-specific monitors. The Base Policy and Layered Templates are combined to form a configuration file and applied to servers that you want to monitor. You can add or remove Layered Templates. Template updates can be applied directly to servers, without changing the base policy or reimporting the updated template into the Base Policy.

Example: Apply Base Policy and Templates to Servers



You can use Layered Templates in the following scenarios:

Disparate applications

Create a library of templates for each server running a different set of applications. You can directly apply the template updates to each server.

Dynamic environments

The workload of the servers changes frequently in dynamic environments. You can use Layered Templates to segregate the monitors in logical groups. Based on the workload changes, you can directly apply the logical groups to systems or removed from systems.

Shared servers

In an enterprise setup, servers are shared across multiple departments. Each department manages and monitors applications on the shared server. You can use Layered Templates to independently manage and apply templates to systems of each department.

Application maintenance

You can split monitoring into multiple templates. In a server, you can remove a template for an application not in use, without affecting the monitoring of the remaining system.

Out of the box templates

You can apply out of the box templates to managed nodes. Configure the policy with the template configuration on the managed nodes. The following table explains the templates that are available for different OS:

Platform	Name of the Operating System	Template
	All types of operating systems	CPU Utilization - Autowatch Swap Capacity
Windows	—	App Monitoring - CA eTrust Antivirus Process Crash System Errors System Processes User Activity Windows Services - Autowatch
UNIX	AIX	System Messages System Processes

Platform	Name of the Operating System	Template
	HPUnix	User Activity
		System Messages
		System Processes
		User Activity
	Linux	System Messages
		System Processes
		User Activity
	Solaris	System Messages
		System Processes
		User Activity

Import a Monitoring Template to SystemEDGE Policy

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

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click the Monitors tab.
The Summary page appears with a list of monitors managed by the policy.
4. Click Action and select Import.
The Import Template Wizard appears.

5. Select the System Type and the monitoring template you want to import from the drop-down lists.
6. (Optional) Define a new base index for each of the imported monitors.
7. Select a Conflict Resolution Option from the drop-down list and click Next.
The Resolve Conflict page appears.
8. Review any monitor conflicts and make adjustments to the indexes, then click Next.
The Summary page appears.
9. Review the monitors that will be imported, then click Finish to complete the import process.
10. Click Save Policy.
The policy is saved.

Copy SystemEDGE Monitoring Template

You can copy an existing SystemEDGE monitoring template.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Monitoring Templates, then click SystemEDGE.
The summary page displays a list of the SystemEDGE Monitoring Templates.
2. Select the monitoring template you want to copy, click Action and select Copy. You can also right-click the monitoring template in the Configure pane and select Copy.
The Copy dialog appears.
3. Enter a new name for the monitoring template and click Ok.
The monitoring template is copied and a configuration screen appears in the right pane.

Modify SystemEDGE Monitoring Template

You can modify the SystemEDGE monitoring template.

Follow these steps:

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

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

2. Select the Template Name.

The Summary page appears with the Template information.

3. Click the Monitors tab.

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

4. Click the appropriate monitor tab and select the monitor you want to modify.

The Edit dialog appears.

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

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

7. Click Save.

The Monitoring Template is saved.

Rename SystemEDGE Monitoring Template

You can rename an existing SystemEDGE monitoring template.

Follow these steps:

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

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

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

The Rename dialog appears.

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

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

Delete SystemEDGE Monitoring Template

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

Follow these steps:

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

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

2. Click the Managed Machines tab

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

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

A confirmation message appears.

4. Click Ok to confirm the deletion.

The monitoring template is deleted.

Review Monitoring Template Application Progress

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

Follow these steps:

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

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

2. Select the Template Name.

The Summary page appears with the Template information.

3. Click the Managed Machines tab.

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

4. (Optional) Click View Configuration.

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

Apply Templates to Machines

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

Follow these steps:

1. Click the Resources pane, open the Configure pane, expand Monitoring Templates, and then select SystemEDGE.

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

2. Select the Template Name.

The summary page appears with the Template information.

3. Click Action and select Apply.

Tabs appear for selecting machines on which to apply the monitoring template. The 'Update machines running this template' tab lets you apply the monitoring template to machines that are already using the template. The 'Apply to Machines not running this template' tab lets you apply the monitoring template to machines without using any template.

4. (Optional) Under Existing Machines, select one of the following options:
 - Update all machines with this template applied.
 - Update only those machines that do not have the latest changes of this template applied.
 - Update only those machines where the template has not been successfully applied.
 - Advanced (manually select machines)
 - Remove this template from machines.
5. (Optional) Under Selected Machines, select the machines to which the template is reapplied.
6. (Optional) Select machines from the 'Apply to the Machines not running this Template' tab to apply the template.
7. Click Apply.

The template application is initiated.

Import a SystemEDGE Configuration to a Template

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

Follow these steps:

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

The Available SystemEDGE Monitoring Templates page appears.

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

The New SystemEDGE Monitoring Template dialog appears.

3. Click Import.

The SystemEDGE Agent Machines window appears.

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

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

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

The template is saved.

Add a Monitor To SystemEDGE Policy

You can add a monitor to a SystemEDGE policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Monitors and select the monitor you want to add.

- [Define a Threshold Monitor](#) (see page 227)
- [Define a Process Monitor](#) (see page 229)
- [Define a Log File Monitor](#) (see page 231)
- [Define a Windows Event Monitor](#) (see page 233)
- [Define a History Monitor](#) (see page 234)
- [Define a Process Group Monitor](#) (see page 236)

4. (Optional) Repeat the process for any additional monitors
5. Click Save Policy.

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

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

More Information:

[Define a Threshold Monitor](#) (see page 227)

[Define a Process Monitor](#) (see page 229)

[Define a Log File Monitor](#) (see page 231)

[Define a Windows Event Monitor](#) (see page 233)

[Define a History Monitor](#) (see page 234)

[Define a Process Group Monitor](#) (see page 236)

Define a Threshold Monitor

You can define threshold settings for the SystemEDGE policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Monitors.

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

4. Click Threshold.

The Threshold Monitors page appears.

5. Click + (New) on the Threshold Monitors toolbar.

The Threshold Monitor Details: New dialog appears.

6. Configure the following threshold settings:

Index

Defines the table index to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class

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

Object Class Name

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

Object Attribute

Specifies the object attribute to monitor. The values in the drop-down list refer to the available attributes of the table selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this threshold monitor.

Object Attribute Name

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

Object Instance

Specifies the object instance to monitor. This value, for example, .3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this threshold monitor. For some object classes, the name of the instance itself can be given (for example, C: instead of .3, or /var for a Unix machine).

Object Instance Name

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

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

Severity

Specifies the severity to use for the object state model.

Operator

Specifies the operator to use.

Value

Defines the value to use.

Sample Type

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

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

7. Click Save
The Threshold Monitor settings are saved.
8. Click Save Policy.
The Threshold Monitor is loaded to the Policy.

Define a Process Monitor

You can define process settings for the SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click Monitors.
The Summary page appears with a list of monitors managed by the policy.
4. Click Process.
The Process Monitors page appears.
5. Click + (New) on the Process Monitors toolbar.
The Process Monitor Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class Name

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

Object Attribute

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

Object Attribute Name

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

Object Instance

Specifies the object instance to monitor. This is the regular expression (dependent from optional settings) to use for matching processes by name, or Windows services by name. Pattern should uniquely match a single process (service). Arguments can be included (see optional settings).

Object Instance Name

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

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

The Threshold Configuration page lets you define the following settings:

Severity

Specifies the severity to use for the object state model.

Operator

Specifies the operator to use.

Value

Defines the value to use.

Sample Type

Specifies the sample type to use.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

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

7. Click Save

The Process Monitor settings are saved.

8. Click Save Policy.

The Process Monitor is loaded to the Policy.

Define a Log File Monitor

You can define log file settings for the SystemEDGE policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Monitors.

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

4. Click Log File.

The Log File Monitors page appears.

5. Click + (New) on the Log File Monitors toolbar.

The Log File Monitor Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index to be used.

Monitor Type

Specifies the monitor type to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Log File/Directory Name

Defines the path to the file or the directory to monitor.

Search Filter

Specifies the search filter.

Interval

Defines the evaluation interval for the monitor in minutes

Severity

Specifies the significance of the monitor on a match.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings page lets you define the flags that can be used for the different monitor entries or history control entries.

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

7. Click Save

The Log File Monitor settings are saved.

8. Click Save Policy.

The Log File Monitor is loaded to the Policy.

Define a Windows Event Monitor

You can define Windows event settings for the SystemEDGE policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Monitors.

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

4. Click Windows Event.

The Windows Event Monitors page appears.

5. Click + (New) on the Windows Event Monitors toolbar.

The Windows Event Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Event Log

Specifies the event log to read.

Event Type

Specifies the event type to match.

Source Filter

Defines the source filter to use.

Description Filter

Defines the description filter to use.

Severity

Specifies the significance of the monitor on a match.

The Maintenance Window subtab lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive.

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings subtab lets you define the flags that can be used for the different monitor entries or history control entries.

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

7. Click Save
The Windows Event Monitor settings are saved.
8. Click Save Policy.
The Windows Event Monitor is loaded to the Policy.

Define a History Monitor

You can define history settings for the SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click Monitors.
The Summary page appears with a list of monitors managed by the policy.
4. Click History.
The History Monitors page appears.
5. Click + (New) on the History Monitors toolbar.
The Historical Details: New dialog appears.

6. Configure the following process settings:

Index

Defines the table index to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Object Class

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

Object Attribute

Specifies the object attribute to monitor. The values in the drop-down list refer to the available attributes of the table selected as Object Class. The attribute (for example, devCapacity = 1.3.6.1.4.1.546.1.1.1.7.1.14) specifies the initial part of the MIB object (OID) to monitor with this History entry.

Object Instance

Defines the object instance to monitor. This value (for example, .3 to monitor the third row in the device table (devTable) specifies the index part of the MIB object (OID) to monitor with this History entry.

Interval

Defines the collection interval in a multiple of 30 seconds.

Buckets

Defines the number of samples to collect.

Add to Performance Cube check box

Specifies whether to collect performance cube data for this entry.

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

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

7. Click Save
The History Monitor settings are saved.
8. Click Save Policy.
The History Monitor is loaded to the Policy.

Define a Process Group Monitor

You can define process group settings for the SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click Monitors.
The Summary page appears with a list of monitors managed by the policy.
4. Click Process Group.
The History Monitors page appears.
5. Click + (New) on the Process Group Monitors toolbar.
The Process Group Details: New dialog appears.
6. Configure the following process settings:

Index

Defines the table index to be used.

Platform

Specifies the platform.

Description

Defines an optional description.

Process Name

Defines the process name. This the regular expression (dependent from optional settings) to use for matching processes by name.

Interval

Defines the evaluation interval for the monitor in a multiple of 30 seconds.

User Name

Defines the user name to match in addition to any process name regular expression.

Group Name

Defines the group name to match in addition to any process name regular expression.

Severity

Specifies the significance of the monitor on a group change

The Maintenance Window page lets you define the following settings:

State

Specifies if the monitor maintenance entry is active or inactive

Start Time

Defines the start time when the monitor is switched off and the maintenance window begins.

Stop Time

Defines the stop time when the monitor is switched on again and the maintenance window ends.

The Optional Settings page lets you define the flags that can be used for the different monitor entries or history control entries.

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

7. Click Save
The Process Group Monitor settings are saved.
8. Click Save Policy.
The Process Group Monitor is loaded to the Policy.

View Monitors Within a SystemEDGE Policy

You can view the monitors contained within a SystemEDGE policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Monitors tab.

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

More Information

[Modify a Monitor Within SystemEDGE Policy](#) (see page 239)

[Delete a Monitor from SystemEDGE Policy](#) (see page 239)

Copy a Monitor Within SystemEDGE Policy

You can copy a monitor within a SystemEDGE policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Monitors tab.

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

4. Click the appropriate monitor tab and select the monitor you want to copy.

5. Click Action and select Copy.

The Edit dialog appears.

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

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

8. Click Save Policy.

The policy is saved.

Modify a Monitor Within SystemEDGE Policy

You can modify a monitor within a SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click the Monitors tab.
The Summary page appears with a list of monitors managed by the policy.
4. Click the appropriate monitor tab and select the monitor you want to modify.
5. Click Action and select Modify.
The Edit dialog appears.
6. Modify the settings according to your needs and click Save.
7. (Optional) Repeat the process for any additional monitors.
8. Click Save Policy.
The policy is saved.

Delete a Monitor from SystemEDGE Policy

You can delete a monitor from a SystemEDGE policy.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click the Monitors tab.
The Summary page appears with a list of monitors managed by the policy.
4. Click the appropriate monitor tab and select the monitor you want to delete.
5. Click Action and select Delete.
A warning message appears.

6. Click Ok to confirm the deletion.
7. (Optional) Repeat the process for any additional monitors.
8. Click Save Policy.

The policy is saved.

Modify Existing Template in SystemEDGE Policy

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

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click the Monitors tab.
The Summary page appears with a list of monitors managed by the policy.
4. Click Action and select Import.
The Import Template Wizard appears.
5. Update the monitoring template with the required information.
6. Select the System Type and the updated monitoring template to import from the drop-down lists.
7. (Optional) Define a new base index for each of the imported monitors.
8. Select the Conflict Resolution Option as "Replace existing monitors with imported entities" and click Next.
The Resolve Conflict page appears.
9. Review any monitor conflicts and make adjustments to the indexes, then click Next.
The Summary page appears.
10. Review the monitors that will be imported, then click Finish to complete the import process.
11. Click Save Policy.
The policy is saved.
12. Select Apply from Action drop-down list.
The saved policy will be applied to desired machines.

Define New SRM Policy

You can create SRM policy to define tests to perform, thresholds to monitor, configuration preferences, and other settings that control how the agent runs and what it monitors.

Follow these steps:

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

The Available Policies page appears.

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

The New Service Response Monitoring Policy dialog appears.

3. Enter a name and an optional description for the policy, the system type and whether to base it on an existing policy and click OK.

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

4. Click Save Policy.

The policy is saved.

More Information:

[Copy SRM Policy](#) (see page 241)

[Rename SRM Policy](#) (see page 242)

[Delete SRM Policy](#) (see page 242)

Copy SRM Policy

You can copy an existing SRM policy.

Follow these steps:

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

The Available Policies page appears.

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

The Copy dialog appears.

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

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

4. Click Save Policy.
The policy is saved.

Rename SRM Policy

You can rename an existing SRM policy.

Follow these steps:

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

The Available Policies page appears.

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

The Rename dialog appears.

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

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

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

4. Click Save Policy.
The policy is saved.

Delete SRM Policy

You can delete an existing SRM policy.

Follow these steps:

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

The Available Policies page appears.

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

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

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

4. Click Save Policy.
The policy is saved.

Add a Test to SRM Policy

You can add a test to an SRM policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click Test and click + (New) on the Test Monitors toolbar.

The New dialog appears.

4. Specify a unique name for the test in the Test name field. The name must be 64 characters or less. Test names are case sensitive.
5. (Optional) Enter a Description for the test.
6. (Optional) Define a test class.
7. Click the test type you want in the Test Type list:

Active Directory

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

Custom

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

DHCP

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

DNS

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

File I/O

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

FTP and TFTP

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

HTTP and HTTPS

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

LDAP

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

NIS

Verifies that NIS map requests are being processed.

NNTP

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

Ping

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

Email

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

SNMP

Verifies that SNMP agents are responding to SNMPv1 GET requests.

SQL Query

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

TCP

Verifies that systems are listening for and processing connection requests.

Virtual User

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

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

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

Note: For more information, see the *SRM User Guide*

11. (Optional) Select the Keep Historical data check box
12. Click Save to add the test to the policy.

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

The policy is saved.

More Information:

[Modify SRM Test](#) (see page 245)

[Add a Threshold Definition To SRM Policy](#) (see page 247)

[Modify SRM Threshold Definition](#) (see page 248)

[Define SRM Control Settings](#) (see page 249)

Modify SRM Test

You can modify an existing SRM test.

Follow these steps:

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

The Available Policies page appears.
2. Select the policy containing the test you want to modify in the Available Policies table.

The Summary page appears.

3. Click the Test tab.
The Test Monitors page appears.
4. Select the test you want to modify, click Actions, and select Modify.
The Edit dialog appears.
5. Modify the test according to your needs and click Save.
The test is updated.
6. Click Save Policy.
The policy is saved.

Copy SRM Test

You can copy an existing SRM test.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click Service Response.
The Available Policies page appears.
2. Select the policy containing the test you want to copy in the Available Policies table
The Summary page appears.
3. Click the Test tab.
The Test Monitors page appears.
4. Select the test you want to copy, click Actions, and select Copy.
A copy dialog appears.
5. Enter the SRM test name.
The SRM test is copied.

Delete SRM Test

You can delete an existing SRM test.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click Service Response.
The Available Policies page appears.
2. Select the policy containing the test you want to delete in the Available Policies table
The Summary page appears.

3. Click the Test tab.

The Test Monitors page appears.

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

5. Confirm your action.

The SRM test is deleted.

Add a Threshold Definition To SRM Policy

You can add a threshold definition to an SRM policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Threshold tab and Click + (New) on the Threshold Monitors toolbar.

The Threshold Monitor Details dialog appears.

4. Configure the following threshold monitor settings:

Name

Defines the threshold monitor name.

Attribute

Specifies the attribute to use.

Operator

Specifies the operator to use.

Warning Value

Defines the warning value to use.

Minor Value

Defines the minor value to use.

Major Value

Defines the major value to use.

Critical Value

Defines the critical value to use.

Fatal Value

Defines the fatal value to use.

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

The threshold definition is saved.

6. Click Save Policy.

The policy is saved.

Modify SRM Threshold Definition

You can modify an existing SRM threshold definition.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy containing the threshold definition you want to modify in the Available Policies table

The Summary page appears.

3. Click the Threshold tab.

The Threshold Monitors page appears.

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

The Edit dialog appears.

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

The threshold definition is updated.

6. Click Save Policy.

The policy is saved.

Define SRM Control Settings

SRM control settings define various aspects of AIM behavior that you typically control in the svcrsp.cf file, including the following:

- Security settings
- Log level
- Index reservations

Control settings defined in SRM policy are applied to all machines to which you apply the policy.

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Control Settings tab on the policy page.

The Control pane appears.

4. Configure the following control settings:

Maximum Number of Threads

Specifies the number of threads the jcollector should use to perform tests

Log Level

Specifies the log level of the SRM AIM. Default is Warning.

Allow External Scripts

Specifies if execution of external scripts is allowed.

Allow Execution of File I/O Tests

Specifies if execution of file I/O tests is allowed.

Allow Untrusted SSL Certificates

Specifies if SSL tests with sites that do not have trusted SSL certificates is allowed.

Java bin Location

Defines the location of the Java executable.

Note: Specify the complete path and binary on AIX.

Override CLASSPATH in Environment

Defines extra classes to load. Overrides CLASSPATH in environment if defined.

No Collector

Specifies if SystemEDGE should start jcollector.

Bypass JRE Internal Cache

Specifies whether to bypass JRE internal cache.

No TOS for IPv4 (HP-UX)

Specifies whether to disable TOS.

Shared Memory Name

Defines the ID for the shared memory.

Reserved Test Indices

Defines reserved range of test indexes.

The control settings are defined.

5. Click Save Policy.

The policy is saved.

Define New SRM Test Definition Template

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

To define a new SRM test definition template

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

The Service Response page appears

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

The New Service Response Test Definition Template dialog appears.

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

The test definition template is created, and the Summary page appears. To add a test to the template, see the section [Add a Test to SRM Policy](#) (see page 243).

4. Click Save Template.

The template is saved.

More Information:

[Import a Test Definition Template into SRM Policy](#) (see page 251)

[Modify SRM Test Definition Template](#) (see page 252)

[Copy SRM Test Definition Template](#) (see page 252)

[Rename SRM Test Definition Template](#) (see page 252)

[Delete SRM Test Definition Template](#) (see page 253)

Import a Test Definition Template into SRM Policy

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

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, and click SystemEDGE.
The Available Policies page appears.
2. Select the policy in the Available Policies table.
The Summary page for the policy appears.
3. Click the Test tab.
The Summary page appears with a list of test monitors managed by the policy.
4. Click Action and select Import.
The Import Template Wizard appears.
5. Select the test template you want to import from the drop-down lists.
6. (Optional) Define a new base index for each of the imported test definitions.
7. Select a Conflict Resolution Option from the drop-down list and click Next.
The Resolve Conflict page appears.
8. Review any test definition conflicts and make adjustments to the indexes, uncheck any test definitions that should not be imported, then click Next.
The Summary page appears.
9. Review the test definitions that will be imported, then click Finish to complete the import process.
10. Click Save Policy.
The policy is saved.

Modify SRM Test Definition Template

You can modify an SRM test definition template.

To modify an SRM test definition template

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

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

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

The Summary page appears for the test template.

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

The Edit dialog appears.

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

5. Click Save Template.

The template is saved.

Copy SRM Test Definition Template

You can copy an SRM test definition template.

To copy an SRM test definition template

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

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

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

The Summary page appears for the test template.

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

The Copy dialog appears.

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

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

Rename SRM Test Definition Template

You can rename an SRM test definition template.

To rename an SRM test definition template

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

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

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

The Summary page appears for the test template.

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

The Rename dialog appears.

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

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

Delete SRM Test Definition Template

You can delete an SRM test definition template.

To delete an SRM test definition template

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

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

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

The Summary page appears for the test template.

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

A warning message appears.

4. Click Ok to confirm the deletion.

A confirmation message appears. The test template is deleted.

Define New SRM Threshold Definition Template

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

To define a new SRM threshold definition template

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

The Service Response page appears

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

The New Service Response Threshold Definition Template dialog appears.

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

The threshold definition template is created, and the Summary page appears. To add a threshold definition to the template, see the section [Add a Threshold definition To SRM Policy](#). (see page 247)

4. Click Save Template.

The template is saved.

More Information:

[Import a Threshold Definition Template into SRM Policy](#) (see page 254)

[Modify SRM Threshold Definition Template](#) (see page 255)

[Copy SRM Threshold Definition Template](#) (see page 255)

[Rename SRM Threshold Definition Template](#) (see page 256)

[Delete SRM Threshold Definition Template](#) (see page 256)

Import a Threshold Definition Template into SRM Policy

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

Follow these steps:

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

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Threshold tab.

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

4. Click Action and select Import.

The Import Template Wizard appears.

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

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

The Resolve Conflict page appears.

7. Review any threshold definition conflicts, make adjustments to the threshold definition name and uncheck any threshold definitions that should not be imported, then click Next.

The Summary page appears.

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

The policy is saved.

Modify SRM Threshold Definition Template

You can modify an SRM threshold definition template.

To modify an SRM threshold definition template

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

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

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

The Summary page appears for the test template.

3. Click Thresholds, select the threshold monitor you want to modify, click Action, and select Modify.

The Threshold Monitor Details dialog appears.

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

5. Click Save Template.

The template is saved.

Copy SRM Threshold Definition Template

You can copy an SRM threshold definition template.

To copy an SRM threshold definition template

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

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

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

The Summary page appears for the threshold template.

3. Click the Threshold tab, select the threshold monitor you want to copy, click Action, and select Copy. You can also right-click the threshold template in the Configure pane and select Copy.

The Copy dialog appears.

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

The threshold definition template is copied and appears in the Threshold Template list.

Rename SRM Threshold Definition Template

You can rename an SRM threshold definition template.

To rename an SRM threshold definition template

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

The Threshold Template List appears with a list of threshold templates.

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

The Summary page appears for the threshold template.

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

The Rename dialog appears.

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

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

Delete SRM Threshold Definition Template

You can delete an SRM threshold definition template.

To delete an SRM threshold definition template

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

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

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

The Summary page appears for the threshold template.

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

A warning message appears.

4. Click Ok to confirm the deletion.

A confirmation message appears. The threshold template is deleted.

Import an Existing SRM Configuration

After upgrading an existing Service Availability (SA) 2.0 AIM to SRM 3.1.0, you can import the previous SA 2.0 configuration and convert it to an SRM 3.1.0 policy.

Follow these steps:

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

The Available Policies page appears.

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

The New service Response Policy dialog appears.

3. Click Import, select the machine you want to import the policy from in the list, and click Ok.

4. Enter a name and an optional description, and click Ok to complete the import process.

5. Click Save Policy.

The policy is saved.

Apply Policy to Machines

After you create configuration policy, apply it to machines across the enterprise. When you apply configuration policy, CA Server Automation pushes a compiled configuration file containing all policy settings to all specified agent machines. The new policy is implemented after an automatic agent warm start.

If one of the following cases occurs, you can reapply the policy to machines:

- You updated the policy.
- You received a notification that the configuration on an agent machine has changed.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Select the policy that you want to apply.

Policy details appear in the right pane.

4. Click Action and select Apply.

Tabs appear for selecting machines on which to apply the policy. The 'Update machines running this policy' tab lets you apply the policy to machines that are already running the policy. The 'Apply to Machines not running this Policy' tab lets you apply policies to machines without any policy or using a different policy.

5. (Optional) Do one of the following options from the 'Update machines running this policy' tab:

- Select 'Update all machines using this policy' to deploy the policy on all machines currently running it. This option is useful if you have made configuration policy changes that you want to apply globally.
- Select 'Update selected groups of machines' to update only machines that meet any of the following criteria:

- Machines running an out-of-date version of the policy
- Machines where policy exceptions have been applied
- Machines running current version of the policy
- Machines with configuration errors for this policy

Select any of these options. Policy exceptions occur when a user applies a point configuration change to an agent that is not represented in the applied policy.

- Select 'Advanced (manually select machines)' to add the machines manually in the Select Machines pane to which you want to reapply the policy.

6. (Optional) Select machines from the 'Apply to Machines not running this Policy' tab to which to apply the policy.

7. Click Apply Policy.

The policy application is initiated.

Review Policy Application Progress

You can review the progress of policy application operations in detail for each individual policy.

Follow these steps:

1. Select the Resources tab, open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Managed Machines tab.

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

4. (Optional) Click View Exceptions.

The Policy Exceptions pane appears and lets you view SNMP sets that have been applied to the system since the policy was last applied.

Note: This screen is only available for SystemEDGE Policies.

5. (Optional) Click View Configuration.

The Policy Configuration page appears and lets you view the configuration file delivered for the agent.

6. (Optional) Click View Errors.

The Policy Errors pane appears. If the policy could not be successfully applied, you can view a list of errors returned by the agent when the policy was rejected.

Configure and View Applied Policies

The Policy feature lets you manage the policies and templates applied to an individual server, a server group, or a service. You can perform the following operations:

- Update Policies and Templates
- View exceptions since the last policy or template application.
- View policy configuration
- View policy errors
- Bulk update Policies
- Delete Templates

Follow these steps:

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Select a system or a service. Click the Resources page, and then Monitoring Software.
The Machine Details page appears.
3. Click Policies.
The Policies page displays a list of Policies of SystemEDGE and SRM, and SystemEDGE Templates.
Note: The Filter displays a list of layered templates in Pending, Delivered (successful), Configured, and Failed states.
4. You can bulk update policies and templates. In the Policies and Templates table, select the policy or template you want to bulk update, click Actions, and select one of the following options:
 - Bulk update SystemEDGE Policies.
 - Bulk update Service Response Policies.
 - Bulk update SystemEDGE Templates.
 - Bulk remove Templates

Note: If the policy is being applied to a single server, you are prompted for the policy name.

Bulk update for Policies:

When the selected policy is applied to a service group, you have an option to select the machine to apply the policy.

Bulk update for Templates:

A dialog provides an option to select the templates from the Available Templates. After you select the templates, click one of the following options:

Replace existing configurations with the selected templates

Removes the existing templates that are applied to all machines and applies the selected templates to all machines.

Append the selected templates to existing configurations

Adds the selected templates. If any of the templates selected are already applied as part of a machine configuration, then those templates are reapplied.

Bulk remove Templates:

Removes the existing templates that are applied to machines.

5. Click Apply Policy to apply the policy or template to the machines.

On the Policies page, you can view the progress of the policy or template being applied to the machines.

6. (Optional) Click View Configuration Icon.

The Policy Configuration page appears. For a machine with a template, it displays the Policies and Templates, and SystemEDGE Configuration file. For a machine with a Service Response Monitor, it additionally displays the Service Response Monitor Configuration file.

7. Click Save Policy.

The policy is saved.

Revert a Policy Back To an Earlier Version

You can revert a policy back to an earlier version.

Follow these steps:

1. Click the Resources tab, open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

The Available Policies page appears.

2. Select the policy in the Available Policies table.

The Summary page for the policy appears.

3. Click the Versions tab.

The Versions page appears.

4. Select the version you want to restore to in the table and click Make Current.

A message appears. Click Ok. A new version of the policy is created and the summary page appears.

5. (Optional) You can make a new copy of the version. Select the version in the Available Policies table and click Copy.

The Copy dialog appears.

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

The policy is copied and added to the policy tree in the Configure pane. The summary page for the new copy appears.

7. Click Save Policy.

The policy is saved.

Specify Default Policy for New Instances

You can set a single default policy for all new discovered instances. The policy will be delivered if a policy was not specified during installation or deployment of SystemEDGE or SRM, or if the specified policy is not available.

To specify default policy

1. Open the Configure pane, expand Policies, then select SystemEDGE or Service Response.

The Available Policies page appears.


2. In the Default Policy section, select the policy you want to use from the Default Policy drop-down list and click Apply

The default policy is applied.

How to Change the Configuration Mode for SystemEDGE

In some cases, it can be necessary to change the configuration mode of SystemEDGE. The following diagram provides an overview of the required actions to change the configuration mode.

How to Change the Configuration Mode of SystemEDGE



Follow these steps:

[Review Requirements](#) (see page 264)

[Review Managed Mode and Unmanaged Mode Details](#) (see page 264)

[Verify the Current Configuration Mode of SystemEDGE](#) (see page 265)

[How to Change SystemEDGE from Managed Mode to Unmanaged Mode](#) (see page 267)

[How to Change SystemEDGE from Unmanaged Mode to Managed Mode](#) (see page 270)

[Verify the SystemEDGE Configuration Mode](#) (see page 272)

Review Requirements

Review the following requirements before you change the configuration mode for SystemEDGE.

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You have a basic understanding of CA SystemEDGE.
- You can access a CA Server Automation manager installation that includes the Monitoring Agent (CA SystemEDGE).
- You can access the monitoring agents (CA SystemEDGE) on managed nodes.
- You can access the CA Server Automation user interface.
- CA Server Automation has discovered all relevant systems.

Review Managed Mode and Unmanaged Mode Details

Consider the following terminology: In this scenario, the terms unmanaged mode and managed mode are used in the context of SystemEDGE configuration.

Unmanaged mode

The SystemEDGE configuration on a particular server is not managed through CA Server Automation Policy Configuration. You can edit the sysedge.cf file to modify the configuration.

Managed mode

The SystemEDGE configuration on a particular server is managed through CA Server Automation Policy Configuration. You specify the SystemEDGE configuration in Policy Configuration on the CA Server Automation manager and distribute it to the appropriate servers in the network. If you edit the sysedge.cf file locally, CA Server Automation overwrites your changes with the next policy distribution.

Consider the following cases which have an influence on the configuration mode of SystemEDGE:

- If you run a typical SystemEDGE installation from the product media, SystemEDGE is configured to run in unmanaged mode after the installation.
- If you run a custom SystemEDGE installation from the product media, you can specify a manager system to use for the managed mode. If you have specified a manager system and CA Server Automation discovers SystemEDGE after the installation, SystemEDGE registers with Policy configuration automatically and SystemEDGE runs in managed mode.
- If you use Remote Deployment to install SystemEDGE on remote systems, you can specify the configuration mode for SystemEDGE in the deployment job. The default value is managed mode.

Important! The Explore pane shows the Managed and Unmanaged folders which list the discovered servers that are polled (managed) or not polled (unmanaged) by CA Server Automation. This property is different from the managed or unmanaged mode of the SystemEDGE configuration. The managed or unmanaged status of a server in the Explore pane has no influence on the configuration mode of SystemEDGE. Specific entries in the SystemEDGE configuration file indicate the configuration mode of SystemEDGE.

Verify the Current Configuration Mode of SystemEDGE

The following steps describe a method to determine the configuration mode of SystemEDGE.

The following terminology is used throughout these use cases:

Static sysedge.cf file

Identifies the file laid down by the installer, and is located in the *Installed_Dir*\SystemEDGE\config directory.

Default:

Windows: C:\Program Files\CA\SystemEDGE\config

UNIX/Linux: /opt/CA/SystemEDGE/config

Dynamic sysedge.cf file

Identifies the ongoing SystemEDGE configuration file, and is located under the *Data_Dir*\port<number> directory.

Default:

Windows: C:\Users\Public\CA\SystemEDGE\port161

UNIX/Linux: /opt/CA/SystemEDGE/config/port161

Follow these steps:

1. Log in the server on which SystemEDGE runs for which you want to determine the configuration mode.
2. Change to the 'data' directory of SystemEDGE and open the port<number> directory. On Windows, you can open the sysedge.cf file in the data directory from the SystemEDGE Control Panel.

Note: The dynamic sysedge.cf file in the 'data' directory is different from the static sysedge.cf file in the 'config' directory.

3. Open the dynamic sysedge.cf file in the port<number> directory.

If SystemEDGE runs in managed mode, the first line specifies a control value (ctrl_value).

Example:

```
ctrl_value 0x9e30d00e
```

```
# Generated file - DO NOT EDIT
#
# Configuration file generated on 2012:03:27 05:37
#
# Generated from default.generic.0.prof
#
version 5.7
```

If SystemEDGE runs in unmanaged mode, the first line specifies the version.


Example:

```
version 5.7
```

How to Change SystemEDGE from Managed Mode to Unmanaged Mode

The following diagram provides an overview of the required actions to change to unmanaged mode.

How to Change SystemEDGE from Managed Mode to Unmanaged Mode



Follow these steps:

[Remove Managed Mode Information from the SystemEDGE Configuration](#) (see page 267)

[Remove Managed Mode Information from the Manager](#) (see page 268)

[Discover the System Running SystemEDGE in Unmanaged Mode](#) (see page 269)

Remove Managed Mode Information from the SystemEDGE Configuration

The following procedure describes how to remove managed mode information from SystemEDGE configuration on a particular server.

Follow these steps:

1. Log in the server on which you want to change the configuration mode for SystemEDGE.
2. Create the following backup directories at a convenient location:

data.backup
config.backup

3. Stop SystemEDGE, using the normal mechanism.
4. Navigate to the 'data' directory of SystemEDGE and open the port<number> directory. The default directory is port161.
The content of the directory is listed.
5. Move the following files to your data.backup directory so that they no longer appear in the port<number> directory:
`.sysedge.id`
`sysedge.cf`
6. Change to the 'config' directory of SystemEDGE.
The content of the directory is listed.
7. Copy the following file to your config.backup directory:
`sysedge.cf`
8. Navigate to the 'config' directory, open the sysedge.cf file in a text editor, and scroll down to the bottom of the file.
9. Delete the following line:
`manager_name <hostname of the manager>`
10. Save the file and start SystemEDGE.
SystemEDGE creates a sysedge.cf file in the 'data' directory without any managed mode information.

Remove Managed Mode Information from the Manager

The following procedure describes how to remove managed mode information from SystemEDGE configuration on a particular server.

Follow these steps:

1. Log in the CA Server Automation user interface and change to Management.
The Resources tab opens and shows the Explore pane.
2. Enter the name of the server on which you have modified the SystemEDGE configuration into the Search field, and click  (Search).
The Search window opens and lists the search results.
3. Click one of the search results.
The resources page for that particular server opens and shows the Quick Start panel.
4. Click Delete from System.
The server disappears from the Explore pane. All server-related objects are deleted on the manager including managed mode information.

Discover the System Running SystemEDGE in Unmanaged Mode

The following procedure describes how to rediscover the server that runs in unmanaged mode.

Follow these steps:

1. Log in the CA Server Automation user interface and change to Management.

The Resources tab opens and shows the Explore pane.

2. Right-click Data Center, select Management, Discover, Server.

The Discover Window opens.

3. Enter the name of the server that you have deleted in the previous procedure and click Finish.

CA Server Automation discovers the server on which SystemEDGE runs in unmanaged mode.

After CA Server Automation has completed Discovery, SystemEDGE on the discovered server has not been registered in Policy Configuration. SystemEDGE runs in unmanaged mode.

4. Double-click the server name in the Explore pane.

The resources page for that server opens.


5. Change to the Summary tab to verify that CA Server Automation has correctly discovered the server. If necessary, you can select a different SNMP community which CA Server Automation uses for monitoring the server.

You can now configure SystemEDGE on that server through editing the sysedge.cf configuration file.

How to Change SystemEDGE from Unmanaged Mode to Managed Mode

The following diagram provides an overview of the required actions to change to managed mode.

How to Change SystemEDGE from Unmanaged Mode to Managed Mode



Follow these steps:

[Remove Unmanaged Mode Information from the SystemEDGE Configuration](#) (see page 270)

[Remove Unmanaged Mode Information from the Manager](#) (see page 271)

[Discover the System to Run SystemEDGE in Managed Mode](#) (see page 272)

Remove Unmanaged Mode Information from the SystemEDGE Configuration

The following procedure describes how to remove unmanaged mode information from SystemEDGE configuration on a particular server and to prepare the server to change to managed mode.

Follow these steps:

1. Log in the server on which you want to change the configuration mode for SystemEDGE.
2. Create the following backup directories at a convenient location:
data.backup
config.backup
3. Stop SystemEDGE, using the normal mechanism.

4. Navigate to the 'data' directory of SystemEDGE and open the port<number> directory. The default directory is port161.

The content of the directory is listed.

5. Move the following file to your data.backup directory so that it no longer appears in the port<number> directory:

sysedge.cf

6. Change to the 'config' directory of SystemEDGE.

The content of the directory is listed.

7. Copy the following file to your config.backup directory:

sysedge.cf

8. Navigate to the 'config' directory, open the sysedge.cf file in a text editor, and scroll down to the bottom of the file.

9. Add the following line:

manager_name <hostname of the manager>

10. Save the file and start SystemEDGE.

SystemEDGE creates a sysedge.cf file in the 'data' directory.

Remove Unmanaged Mode Information from the Manager

The following procedure describes how to remove unmanaged mode information from SystemEDGE configuration on a particular server.

Follow these steps:

1. Log in the CA Server Automation user interface and change to Management.

The Resources tab opens and shows the Explore pane.

2. Enter the name of the server on which you have modified the SystemEDGE configuration into the Search field, and click  (Search).

The Search window opens and lists the search results.

3. Click one of the search results.

The resources page for that particular server opens and shows the Quick Start panel.

4. Click Delete from System.

The server disappears from the Explore pane. All server-related objects are deleted on the manager.

Discover the System to Run SystemEDGE in Managed Mode

The following procedure describes how to rediscover the server that causes SystemEDGE to run in managed mode.

Follow these steps:

1. Log in the CA Server Automation user interface and change to Management.
The Resources tab opens and shows the Explore pane.
2. Right-click Data Center, select Management, Discover, Server.
The Discover Window opens.
3. Enter the name of the server that you have deleted in the previous procedure and click Finish.
CA Server Automation discovers the server.
After CA Server Automation has completed Discovery, SystemEDGE on the discovered server has been registered in Policy Configuration. SystemEDGE runs in managed mode.
4. Double-click the server name in the Explore pane.
The resources page for that server opens.
5. Change to the Summary tab to verify that CA Server Automation has correctly discovered the server. If necessary, you can select a different SNMP community which CA Server Automation uses for monitoring the server.

Verify the SystemEDGE Configuration Mode

Basically, you can repeat the procedure in [Verify the Current Configuration Mode](#) (see page 265).

If SystemEDGE runs in unmanaged mode, the first line of the dynamic sysedge.cf file in the 'data' directory specifies the version.

Example

```
version 5.7
```


If SystemEDGE runs in managed mode, the first line specifies a control value (ctrl_value).

Example

```
ctrl_value 0x9e30d00e

# Generated file - DO NOT EDIT
#
# Configuration file generated on 2012:03:27 05:37
#
# Generated from default.generic.0.prof
#
version 5.7
```

During the discovery process, additional meta information has been added to the end of the dynamic sysedge.cf file.

Example

```
template data_directory <path>
data_directory "C:\Users\Public\CA\SystemEDGE\"
template default_port CA Portal
default_port 161
template manager_name <name>
manager_name manager_server.mycompany.com
template manager_policy_name <policy>
manager_policy_name default.generic
template manager_policy_version <version>
manager_policy_version 1
```

More information:

[Verify the Current Configuration Mode of SystemEDGE](#) (see page 265)

Chapter 6: Managing Virtual Environments

This section contains the following topics:

[VMware vSphere and vCenter Server](#) (see page 275)

[VMware vCloud](#) (see page 335)

[IBM PowerVM \(LPAR\)](#) (see page 354)

[Citrix XenServer](#) (see page 388)

[Solaris Zones](#) (see page 406)

[Microsoft Hyper-V Server](#) (see page 424)

[Cisco UCS](#) (see page 446)

VMware vSphere and vCenter Server

CA Server Automation manages VMware vSphere and vCenter Server virtual environments. The vCenter Server is the central component which CA Server Automation and the vCenter AIM use to access the vSphere environment. SystemEDGE and the vCenter AIM run on the CA Server Automation manager server or on an arbitrary Windows server.

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

The vCenter AIM communicates with one or more remote vCenter Server instances through web-services. The AIM communicates with the manager through SNMP. If more than one vCenter AIM is available to manage a vCenter Server, you can specify your preferred vCenter AIM during configuration or you can let the manager choose it on its own.

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

Monitored vSphere and vCenter Server Resources

The vCenter AIM detects the logical and physical relationships between the components in a vSphere environment. The AIM provides a view of the entire virtualized environment and manages the following resource types and properties:

vCenter Server

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

Datacenter

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.

VMware Cluster/High Availability/Fault Tolerance

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

Resource Pool

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

vApp

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

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 other disk media.

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.

vNetwork Distributed Switch

Abstracts the configuration of virtual switches from the host to the datacenter level. A vNetwork Distributed Switch operates as a single virtual switch that spans across all hosts in a datacenter which are associated with that switch. vNetwork Distributed Switches consist of distributed port groups which are similarly configured to port groups on standard switches, but extend across multiple hosts. These properties allow virtual machines to maintain a consistent network configuration as they migrate among multiple hosts.

vNetwork Standard Switch

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

How to Configure the vCenter Server Management Components

The following diagram provides an overview of the required actions. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the vCenter Server Management Components



Follow these steps:

[Review Requirements](#) (see page 279)

[Review Interactions Between vCenter Server Management Components](#) (see page 280)

[Add a New vCenter Server Connection to the Manager](#) (see page 282)

[Troubleshoot the vCenter Server Connection](#) (see page 283)

[Add the AIM Instance for the vCenter Server](#) (see page 286)

[Troubleshoot the vCenter AIM Instance Connection](#) (see page 287)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 292)

Review Requirements

Review the following requirements before you start configuring the vCenter Server management components of CA Server Automation:

- 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 know which protocol (HTTP or HTTPS) and port to use for accessing the vCenter Server of the vSphere environment through web services. Default: HTTPS, Port 443
- You verified that the new vSphere environment and its vCenter Server are running properly.
- If the VMware PMM and vCenter AIM are installed on different systems, you have verified that the SNMP settings on these systems are consistent. Read and write community strings and the SNMP port number must be identical.
- You verified that the CA Server Automation manager discovered remote vCenter AIM Servers that you want to use.

More information:

[Review Interactions Between vCenter Server Management Components](#) (see page 280)

[Add a New vCenter Server Connection to the Manager](#) (see page 282)

[Add the AIM Instance for the vCenter Server](#) (see page 286)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 292)

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 other virtual resources. Based on thresholds, SystemEDGE and the vCenter AIM determine the status of a monitored resource and propagate this information to the CA 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.

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 that are automatically controlled by thresholds, status, and values that are 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 that are represented by four vCenter Servers. In general, the vCenter PMM and each vCenter AIM with its multi-instance support can connect to multiple vCenter Servers. The number of connections shown in the diagram 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.

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

More information:

[Add a New vCenter Server Connection to the Manager](#) (see page 282)

[Add the AIM Instance for the vCenter Server](#) (see page 286)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 292)

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), specify the preferred AIM, enable Managed Status (checkbox), and click OK.

If the network connection has been established successfully, the vCenter Server is added to the top right vCenter Servers pane with a green status icon. CA 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 283).

More information:

[Add the AIM Instance for the vCenter Server](#) (see page 286)


[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 292)

[Troubleshoot the vCenter Server Connection](#) (see page 283)

Troubleshoot the vCenter Server Connection

The vCenter Server connection has failed. Follow the troubleshooting information indicated in the following diagram:

How to Troubleshoot the vCenter Server Connection



Follow these steps:

[vCenter Server Connection Failed](#) (see page 284)

[Add the AIM Instance for the vCenter Server](#) (see page 286)

[Verify the vCenter Server Folder Appearance in the Resources Tree](#) (see page 292)

vCenter Server Connection Failed

Symptom:

After I have added a vCenter Server connection under Administration, Configuration, the validation of the connection to the vCenter Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used vCenter Server connection data (server name, user, password, protocol, port) is still valid. If necessary, update the connection data.
- Verify, if the vCenter Server system is running and accessible.
- Verify, if the VMware Management Service on the vCenter Server system is running properly.

To update the vCenter Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.

The New or Edit vCenter Server dialog appears.

2. Add the valid server name, user, password, protocol, 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 other virtual components.

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 292)

[Troubleshoot the vCenter AIM Instance Connection](#) (see page 287)


Troubleshoot the vCenter AIM Instance Connection

The vCenter AIM Connection is in not-ready status. One of the following status icons appears:

-  Discovery in progress - Wait until the platform manager synchronizes all data.
-  Error - Unable to connect to the AIM. Check the network configuration.
-  No Polling - The CA Server Automation manager does not poll this AIM instance.
-  Disabled - This instance is not managed.

Follow the troubleshooting information indicated in the following diagram:

How to Troubleshoot the vCenter AIM Instance Connection



More information:

[vCenter AIM Instance Status Icon Shows Discovery in Progress](#) (see page 289)

[vCenter AIM Instance Status Icon Shows Error](#) (see page 289)

[vCenter AIM Instance Status Icon Shows No Polling](#) (see page 290)

[vCenter AIM Instance Status Icon Shows Disabled](#) (see page 291)

vCenter AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Discovery in Progress).

Solution:

Wait until the discovery process of the vSphere environment has completed. The discovery duration depends on the amount of managed objects that are related to virtual and physical resources in vSphere. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Server Automation adds a vCenter Server folder to the Resources tree. Then you can start managing vSphere and its entire virtual infrastructure.

vCenter AIM Instance Status Icon Shows Error

Symptom:

After I add a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the vCenter AIM:

- Verify if the vCenter AIM Server is accessible.
- Verify if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the vCenter AIM server system is accessible:

1. Open a command prompt on the CA 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 host 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.

vCenter AIM Instance Status Icon Shows No Polling

Symptom:

After I add a vCenter AIM instance for a vCenter Server under Administration, Configuration, the status icon shows  (no polling).

Solution:

No specific actions are required for the associated instance. This icon informs you that the CA Server Automation manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular vCenter Server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

vCenter AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation discovers vCenter AIM instances in the network, the status icons of several instances show  (Disabled). This vCenter AIM instance is not managed.

This status appears if CA 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.

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.

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.

Device Management for VMs

Device management includes the following:

- Adding and removing vDisks
- Adding and removing vNICs

Add or Remove Virtual Disk

You can dynamically add or remove virtual disks from a VM. The following disks can be added:

- A new disk from the same or another datastore
- An existing disk from the datastore
- Adding an existing disk from another datastore

To add a virtual disk

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Find and right-click a virtual machine on the Explore pane, or click the Quick Start tab, and select Add New Disk.
The Add New Disk dialog appears.

4. Enter the new disk details according to your needs.
A message prompts for confirmation.
5. Click Ok.
A message appears confirming that the new disk is added.

To remove a virtual disk

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Find and right-click a virtual machine on the Explore pane, or click the Quick Start tab, and select Delete Disk.
The Delete Disk dialog appears.
4. Select the hard drive and whether to delete data.
A message prompts for confirmation.
5. Click Ok.
A message appears confirming that the disk is deleted.

Add or Remove Virtual Network Interface

You can dynamically add or delete a virtual network interface from an existing VM.

To add a virtual network interface

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Find and right-click a virtual machine on the Explore pane, or click the Quick Start tab, and select Add New Virtual Network Interface.
The Add New Virtual Network Interface dialog appears.
4. Enter the new network interface details.
A message prompts for confirmation.
5. Click Ok.
A message appears confirming that the new card is added.

To remove a virtual network interface

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Find and right-click a virtual machine on the Explore pane, or click the Quick Start tab, and select Delete Virtual Network Interface.
The Delete New Virtual Network Interface dialog appears.
4. Select the NIC you want to delete.
A message prompts for confirmation.
5. Click Ok.
A message appears confirming that network interface is deleted.

Fault Tolerance for Virtual Machines

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

Fault Tolerance Properties of Virtual Machines

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.

Monitor 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.

To monitor fault tolerance properties

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Expand the VMware vCenter Server folder and the ESX server object.
A list of VMs appears.
4. (Optional) Select the ESX host.

The following FT attributes appear in the Summary tab:

Fault Tolerance

Identifies whether the host has the fault tolerance enabled.

Fault Tolerance version

Identifies the version of Fault Tolerance running on the host.

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

Total Primary VMs

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

Total Secondary VMs

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

Powered on Primary VMs

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

Powered on Secondary VMs

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

5. (Optional) Select the VM.

The following FT properties display in the Summary tab.

Fault Tolerance Status.

Indicates the VM fault tolerance status.

Not Fault Tolerant

Indicates that the VM is not fault tolerant.

Protected

Indicates that the VM is fault tolerant and protected.

Not Protected (Starting)

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

Not Protected (Need Secondary VM)

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

Not Protected (Disabled)

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

Not Protected (VM Not Running)

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

Secondary VM Location

Identifies the secondary host location.

Manage Fault Tolerance

You can control the fault tolerance properties of the VMs.

To manage fault tolerance properties of VMs

1. Select a VM in the Explore pane.
The General Information pane appears on the right side, displaying the Fault Tolerance Status of the VM.
2. Right-click a VM, select Management, and select one action from the drop-down menu. The following actions for managing fault tolerance are available:
 - Turn Off Fault Tolerance
 - Enable Fault Tolerance
 - Disable Fault Tolerance
 - Migrate Secondary VM
3. Provide information and or confirmation for a selected action.
Confirmation message appears.

Hot-plug Support for VMs

CA 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*.

Dynamically Add or Remove vCPU

You can dynamically add or remove CPU to VMs that have been provisioned. If hot plug is enabled for the VM, you can dynamically add vCPUs during runtime.

CA Server Automation verifies the following VM properties:

- ESX license (ESX Level)
- Maximum supported vCPUs (ESX Level)
- Hot plug enabled (VM Level)

Examples

- If ESX license allows 8 CPUs (Enterprise Plus) AND Max Support vCPUs is 8 AND Hot Plug is DISABLED then you can add: 1, 2, 4, 8 CPUs
- If ESX license allows 8 CPUs (Enterprise Plus) AND Max Support vCPUs is 8 AND Hot Plug is ENABLED then you can add: 1, 2, 3, 4, 5, 6, 7, 8 CPUs

To add vCPU

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Find and right-click a virtual machine on the Explore pane, or click the Quick Start tab, and select Add/Remove vCPU.
The Modify vCPU dialog appears.
4. Adjust the number of CPUs according to your needs.
A message prompts for confirmation.
5. Click Ok.
A message appears confirming the modification.

Note: To remove vCPU, the virtual machine must be turned off.

Dynamically Add or Remove Memory

You can dynamically add or remove memory to VMs that have been provisioned. If hot plug is enabled for the VM, you can dynamically add memory during runtime.

To add memory

1. Click Resources.
The Resources page appears.
 2. Open the Explore pane.
Available groups, services, and systems appear.
 3. Find and right-click a virtual machine on the Explore pane, or click the Quick Start tab, and select Add/Remove Memory.
The Modify Virtual Memory dialog appears.
 4. Adjust the memory according to your needs.
A message prompts for confirmation.
 5. Click Ok.
A message appears confirming the modification.
- Note:** To remove memory, the VM must be turned off.

Logical Volumes in Virtual Machines

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

Resource Allocation

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

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

Resource Allocation Shares

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

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

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

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

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

Resource Allocation Reservation

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

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

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

Resource Allocation Limit

A limit specifies the maximum 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.

Edit VM CPU and Memory Allocation

You can edit the number of CPU and memory shares allocated to a virtual machine to adjust its allocated resources. When you add resources, the appropriate amount of unassigned memory or CPU shares must be available for the operation to succeed. If values exist for the minimum and maximum allowed memory or CPU shares, any resource allocation change must stay within these limits.

You can also create and schedule policy with specific VM resource allocation actions.

To edit VM CPU and memory allocation

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Configuration, Resource Allocation.
The Resource Allocation section appears.
3. Adjust the number of CPU and memory shares allocated to the virtual machine and click Save for each value that you edit.
A confirmation message appears.


How to Use Policy Actions to Identify Performance Issues

This scenario provides information about how a system administrator can identify and dynamically address performance issues. This information is meant to help System Administrators to optimize the allocation of resource shares of their managed vCenter environments.

The policy actions identify VM resources and dynamically adjust the allocation of CPU shares. *Shares* determine which VM gets resources when there is competition for resources among VMs. Using shares allows dynamic allocation of CPU resources. Each VM is allocated a specified number of shares. The allocation is dynamically changed based on the current usage of CPU resources on the ESX Server host.

If CPU usage of any VM is over 50 percent, allocation of CPU shares increases dynamically. If CPU usage is less than 30 percent, the CPU shares allocation decreases dynamically. The policy component not only identifies the problematic virtual machines but ensures dynamic actions that sustain business continuity. Using policy actions ensures that resources are allocated to virtual machines that are in need and deallocated when the need is gone.

How to Use Policy Actions to Identify Performance Issues



To identify and address performance issues using policy actions, follow these steps:

1. [Create an action for CPU metric.](#) (see page 306)
2. If CPU usage is more than 50 percent, [create a rule for CPU metric to increase allocation.](#) (see page 307)
3. If CPU usage is less than 30 percent, [create a rule for CPU metric to decrease allocation.](#) (see page 307)

Create an Action for CPU Metric

Policy provides the creation of rules and actions that can be used to create policies for the automated management of systems. Custom actions can be created for actions not included in the default library.

Follow these steps:

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Click Policy, then click Actions.
The Actions page appears.
4. Click '+' on the upper right side bar to add an action.
5. Enter the name of the Action.
6. Select Resource Configuration from the Category drop-down list.
7. Select Configure Shares from the Type drop-down list.
8. In the VC Server field, leave the entry as "%VCServer%" to apply this action on any VM across any VC Server.
9. In the VC Data Center field, leave the entry as "%DATACENTER%".
10. In the Target VM Machine field, leave the entry as "%VMNAME%".
11. Select Set CPU from the Operations drop-down list and enter Values as 10000.
The number is arbitrary and the share values are set to normal.
Note: Use higher or lower numbers to increase and decrease the share allocations accordingly.
12. If the changes require approval, enable Help Desk Approval.
A Message will appear in the Event Console after the Action is created.
CAAP4521 Policy: Action <action name> was created.

Create a Rule for CPU Metric to Increase Allocation

Creating a rule for CPU metric to increase CPU allocation ensures dynamic resource allocation when the usage exceeds the threshold.

Follow these steps:

1. Click on the Resources tab, Policy, Rules.
2. Click '+' on the upper right side bar to add a rule.
3. Enter the name of rule and click Next.
4. Select the action from the "Action Selection" list for the rule and click Next.
5. Enter the metric-based rule where CPU usage is greater than 50 percent to increase the CPU shares of the VM.

Create a Rule for CPU Metric to Decrease Allocation

Create a Rule for CPU metric to decrease allocation. This rule decreases the CPU shares when the CPU usage is less than 30 percent.

Follow these steps:

1. Click on the Resources tab, Policy, Rules.
2. Click '+' on the upper right side bar to add a rule.
3. Enter the name of rule and click Next.
4. Select the action from the "Action Selection" list for the rule and click Next.
5. Enter the rule-based on metric where CPU usage is less than 30 percent to decrease the CPU shares of the VM.

vApp Support

A *vApp* is a specific resource pool which treats a collection of VMs as a single unit. vApp uses the Open Virtualization Format (OVF). OVF is a standard to specify and encapsulate all components of a multi-tier application and the operational policies and service levels that are associated with it. CA 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.

A vApp is represented at the host and cluster level.

CA Server Automation supports the following management operations on the vApp level:

- Discover
 - Server
 - Network
 - vCenter Server
- Capture Service
- Add Resource
- Clone vApp
- Power On vApp
- Power Off vApp
- Suspend vApp
- Delete from VMware vCenter
- Unregister from VMware vCenter
- Edit Sort Order

CA Server Automation supports the following provisioning operations on vApps:

- Provision VMware VM
- Provision VMware vApp

Provision VMware vApp

You can create a vApp directly on the ESX host or cluster level, or as part of an existing resource pool or vApp.

Follow these steps:

1. From the host or clusters level in the Explore pane, right-click the ESX host or cluster.

A pop-up menu opens.

2. Select Provisioning, Provision VMware vApp.

The Create New vApp dialog appears.

3. Specify the following fields and click OK.

Name

Identifies the vApp.

CPU Shares

Specifies CPU shares for this vApp with respect to the total CPU resources of the parent host, resource pool, or vApp. Sibling vApps share resources according to their relative share values bounded by the reservation and limit. Specify a number of shares which expresses the appropriate proportional weight.

For example, assume that you have vApp1 and vApp2 on a host and each has 1000 CPU shares. The weight is equal and each vApp can allocate 50 percent CPU time of the parent host. However, if vApp1 has 2000 CPU shares and vApp2 has 1000, then the weight is not equal. The total number is 3000 shares, and 1000 shares represent 33.3 percent, and 2000 shares represent 66.6 percent. So vApp1 can allocate 66.6 percent and vApp2 can allocate 33.3 percent of CPU time.

CPU Reservation

Specifies the guaranteed CPU allocation for this vApp.

CPU Unlimited

Disables the CPU limit setting. The actual limit is now set to the available physical resource.

CPU Limit

Specifies the upper limit for the CPU allocation for this vApp. You can usually accept the default.

Memory Shares

Specifies memory shares for this vApp with respect to the total memory resources of the parent resource pool or vApp. Sibling vApps share resources according to their relative share values bounded by the reservation and limit. Specify a number of shares which expresses the appropriate proportional weight.

For example, assume that you have vApp1 and vApp2 on a host and each has 1000 memory shares. The weight is equal and each vApp can allocate 50 percent memory of the parent host. However, if vApp1 has 2000 memory shares and vApp2 has 1000, then the weight is not equal. The total number is 3000 shares, and 1000 shares represent 33.3 percent, and 2000 shares represent 66.6 percent. So vApp1 can allocate 66.6 percent and vApp2 can allocate 33.3 percent of memory.

Memory Reservation

Specifies the guaranteed memory allocation for this vApp.

Memory Unlimited

Disables the memory limit setting. The actual limit is now set to the available physical resource.

Memory Limit

Specifies the upper limit for the memory allocation for this vApp. You can usually accept the default.

The new vApp appears in the Explore pane.

Clone vApp

You can clone a vApp which is similar to clone a virtual machine.

Follow these steps:

1. From the host or clusters level in the Explore pane, select the vApp which you want to clone.
2. Right-click the vApp.
A pop-up menu opens.
3. Select Management, Clone vApp.
The Clone a vApp dialog appears.
4. Specify the following fields and click OK.

Name

Identifies the cloned vApp.

Location

Specify the appropriate location. Expand the object that is displayed on the pop-up menu and select the location.

Datastore

Specify the appropriate datastore from the drop-down menu.

The cloned vApp appears in the Explorer pane.

More vApp Operations

CA Server Automation supports the following additional operations on vApps:

- Power On
- Power Off
- Suspend
- Delete from VMware vCenter
- Unregister from VMware vCenter

Follow these steps:

1. From the host or clusters level in the Explore pane, select the appropriate vApp.
2. Right-click the vApp.
A pop-up menu opens.
3. Select Management, and click the desired operation.
A confirmation dialog appears.
4. Click OK.
CA Server Automation performs the selected operation.

Monitor vApps Through Events

You can monitor vApps through the following events:

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

The following traps are available:

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

To monitor vApps through events

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

The Filter panel opens.

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

The Events panel lists the filtered events.

vCenter Server in a Cluster

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

Virtual Standard Switches and Virtual Distributed Switches in the vNetwork Panel

The vNetwork panel is available in the user interface at the VMware datacenter level and ESX host level. At the VMware datacenter level vNetwork indicates the Virtual Distributed Switches of that datacenter. At the ESX host level vNetwork indicates the associated Virtual Distributed Switches and Virtual Standard Switches.

More information:

[vNetwork Standard Switches \(vSwitch\)](#) (see page 314)

[Distributed Virtual Switches](#) (see page 314)

[Properties](#) (see page 315)

[Actions](#) (see page 318)

[Monitor Distributed Virtual Switches Through Events](#) (see page 319)

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.

You can expand the Virtual Standard Switch objects to see the associated ports and portgroups.

- The port groups contain associated virtual machines that use the portgroup.

Distributed Virtual Switches

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

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

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

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

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

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

Note: If you use the Cisco Nexus 1000V Switch, the VSM VM does not appear as a special VM in the CA 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.

You can expand the Virtual Distributed Switch objects to see the associated portgroups and uplink groups.

- The port groups contain associated VMs that use the portgroup.
- The Uplink Groups list the physical uplink adapters.

Properties

The Properties pane displays the properties of Virtual Standard Switches or Virtual Distributed Switches.

Policies

The following list contains default policies or enabled policies for Virtual Standard Switches or Virtual Distributed Switches.

Promiscuous Mode

Indicates whether or not all traffic is seen on the port.

MAC Address Changes

Indicates whether or not the Media Access Control (MAC) address can be changed.

Forged Transmit

Indicates if the MAC address is different to the MAC address of the virtual network adapter.

Traffic Shaping

Indicates whether or not traffic shaper is enabled on the port.

Average Bandwidth

Indicates the average bandwidth in bits per second if shaping is enabled on the port.

Peak Bandwidth

Indicates the peak bandwidth during bursts in bits per second if traffic shaping is enabled on the port.

Burst Size

Indicates the maximum burst size allowed in bytes if shaping is enabled on the port.

Network Failure Detection

Indicates whether network failure detection is enabled. Valid values are:

- false (1)
- true (2)

Notify Switches

Specifies whether or not to notify the physical switch if a link fails

Fallback

Indicates if fallback is enabled.

Policy Inbound Frames

Indicates whether or not the teaming policy is applied to inbound frames.

Active Adapters

Displays a list of active network adapters used for load balancing.

Standby Adapters

Displays a list of standby network adapters used for failover.

vSwitch Properties

The following vSwitch properties indicate port number characteristics:

Number of Ports

Indicates the current number of ports of the Virtual Distributed Switch or Virtual Standard Switch.

Maximum Number of Ports

Indicates the maximum number of ports of the Virtual Distributed Switch.

Note: For Virtual Distributed Switches, this information is only available on the VMware datacenter level. On the ESX host level, it is not available.

Port Group Properties

The following port group property indicates the VLAN ID:

VLAN ID

Indicates the VLAN ID of a port group.

Port Properties

The following properties specify port characteristics:

VLAN ID

Indicates the VLAN ID of a port.

Type

Indicates the type of a port, for example, VMkernel Port or Service Port.

Network Properties

The following properties specify network characteristics of the virtual switch:

- IPv4 Address
- IPv6 Address
- MAC Address

Virtual Machine Counts

The following values provide statistical information about VMs associated with a port group.

- Powered On
- Powered Off
- Suspended
- Unknown

Actions

Use the appropriate actions to manage your Virtual Standard Switches and Virtual Distributed Switches. The following actions are available:

- Add vSwitch
- Update vSwitch
- Remove vSwitch
- Add Portgroup
- Update Portgroup
- Remove Portgroup
- Rename Portgroup

When you apply these actions, a dialog opens and prompts you to enter the required information. Possible fields are:

Switch Name

Specifies the switch name to perform the operation on.

NICs

(Optional) Specifies lists of physical NICs associated with the ESX host members.

Uplink Port Names

(Optional) Specifies a list of uplink port names to use.

Maximum Number of Ports

(Optional) Specifies the maximum number of ports for the Virtual Distributed Switch.

Bindtype

(Optional) Specifies the bind type of the port group. Valid values are:

earlyBinding

Assigns the ports when the VM binds to the portgroup. This type of binding ensures connectivity at all times, but permanently reserves the port. This binding type is the default.

lateBinding

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. LateBinding is configurable through vCenter.

ephemeral

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. Ephemeral binding is configurable through the ESX Host and vCenter.

Number of Ports

(Optional) Specifies the number of ports of the port group.

Portgroup Name

Specifies the port group name.

New Portgroup Name

Specifies the new port group name.

LAN ID *vlanID*

(Optional) Specifies an Integer value (vlan ID) used for the virtual portgroup operations.

Monitor Distributed Virtual Switches Through Events

You can monitor Distributed Virtual Switches through the following events:

- Add switch:

Distributed Virtual Switch VM-dvSwitch added to Datacenter MyDC. vSphere: vcserver.mycomp.com

- Delete switch:

Distributed Virtual Switch VM-dvSwitch removed from Datacenter MyDC. vSphere: vcserver.mycomp.com

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

To monitor Distributed Virtual Switches through events

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

VMware vCenter Provisioning and Common Use Cases

This section provides instructions how to provision virtual resources and to perform common use cases.

Add a Virtual Machine (vCenter Server)

You can use one of two methods to add a VM:

- Clone a predefined template
- Clone an existing VM and a customization specification. The customization specification defines the characteristics of the Guest OS.

VM provisioning supports Standard Switches and Distributed Virtual Switches. When provisioning a VM that is attached to a Distributed Virtual Switch, you can specify the appropriate discovered dvPort Group in the user interface. dvPort Groups define how a connection is made to the network through the Distributed Virtual Switch.

To add a VM

1. Right-click VMware vCenter Server in the Explore pane and select Provisioning, Provision VMware VM.

VMware vCenter Provisioning dialog appears.

2. Select options from the drop-down lists to specify the settings.

Note: The virtual machines listed for cloning are limited to virtual machines that are monitored by CA 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.

Clone a Virtual Machine

Cloning a virtual machine creates a copy of the virtual machine that you can place anywhere in the same virtual machine farm. You can also customize the guest operating system when creating a clone. You can only clone a virtual machine when it is in the powered off state.

To clone a virtual machine

1. Open the Explore pane.

Available groups, services, and systems appear.

2. Find and right-click the virtual machine to clone on the Explore pane and select Management, Cloning.

The Cloning pane appears.

3. Complete the following fields and click Clone:

Name

Specifies the VM clone name.

Datastore

Specifies the datastore under which to store the cloned VM. The datastore must be in the same farm as the source VM.

Custom Spec

Specifies the guest operating system specification to use. You can select the default or a customization.

Destination Resource Pool

Specifies the pool from which the cloned VM obtains resources.

A message appears confirming the request submission.

4. Click the Summary tab for the virtual machine.

Verify that an event appears to confirm the operation. The clone appears in the Explore pane after the operation completes.

Manage VM Status (VMware)

You can control the status of vCenter Server virtual machines by performing one of the following VM operations:

- Power On
- Power Off
- Suspend
- Reset
- Shut Down

You can perform any of these operations on multiple VMs simultaneously.

To control VM status

1. Select the virtual machine on which you want to perform a status operation in the Explore pane.
2. Right-click the VM, select Management. You can also click Quick Start and click the related link of power control. Select *one* of the following:

Power On

Starts the virtual machine and boots the guest operating system. You can only power on a virtual machine that is currently powered off or suspended.

Power Off

Powers off the virtual machine. You can only power off a virtual machine that is currently powered on or suspended.

Suspend

Pauses the virtual machine and saves its current state. All activity is suspended until you resume the machine.

Reset

Shuts down the guest operating system and restarts it.

Shutdown

Shuts down the guest operating system. You can only shut down a virtual machine that is currently powered on.

A confirmation dialog appears.

3. Click OK.

The status operation occurs, and a confirmation message appears. Refresh the interface to view the new VM status. An event should appear confirming the result of the operation.

The following icons indicate VM status:



Indicates that the VM is in critical state.



Indicates that the VM is in warning state.



Indicates that the VM is in normal state.



Indicates that the VM is in unknown state.

Convert a Template to a Virtual Machine

You can convert a virtual machine template to a virtual machine. When you convert a template to a VM, the VM uses the template name and settings.

To convert a template to a virtual machine

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine template on the Explore pane and select Management, Convert to Virtual Machine.
The Conversion page appears.
3. Select the ESX server and resource pool for the virtual machine and click Convert.
A message appears confirming the request submission.
4. Click the Summary tab for the virtual machine template.
Verify that an event appears to confirm the operation. After the operation completes, the template appears as a virtual machine on the Explore pane when you refresh the interface.

Convert a Virtual Machine to a Template

You can convert a powered off virtual machine to a template to use the virtual machine's configuration as a base for other virtual machines.

To convert a virtual machine to a template

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Convert to Template.
A confirmation dialog appears.
3. Click OK.
A message appears confirming the request submission.
4. Click the Summary tab for the virtual machine.
Verify that an event appears to confirm the operation. After the operation completes, the virtual machine appears as a template in the Explore pane when you refresh the interface.

Create a Snapshot

Create a snapshot to preserve the current state of a virtual machine so that you can return to the same state at a later time. A snapshot preserves the entire state of the virtual machine, including memory contents, settings, and virtual disk state. You can create snapshots for virtual machines that are powered on, powered off, or suspended.

To create a snapshot

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots.
The Snapshots pane appears.
3. Select New from the Actions drop-down menu.
The Create New Snapshot dialog appears.
4. Enter the snapshot name and description, specify whether to enable capture memory, and click OK.
A confirmation message appears.
5. Click Summary for the virtual machine.
6. Verify that an event confirms the operation.
The snapshot appears in the Snapshots pane once the operation is complete.

Delete a Snapshot

You can delete a snapshot that you no longer need.

To delete a snapshot

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots.
The Snapshots pane appears and displays all existing snapshots for the virtual machine.
3. Select a snapshot and select Delete from the Actions drop-down menu.
A confirmation dialog appears.
4. Click OK.
A message appears confirming the request submission.
5. Click the Summary tab for the virtual machine.
Verify that an event appears to confirm the operation. The snapshot disappears from the Snapshots pane after the operation completes.

Delete all Snapshots

You can delete all existing snapshots for a virtual machine in one operation.

To delete all snapshots

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots.
The Snapshots pane appears and displays all existing snapshots for the virtual machine.
3. Select Delete All from the Actions drop-down menu.
A confirmation dialog appears.
4. Click OK.
A confirmation message appears.
5. Click Summary for the virtual machine.
Verify that an event appears to confirm the operation. All snapshots disappear from the Snapshots pane after the operation completes.

Delete a Virtual Machine

When you delete a virtual machine from VMware vCenter Server, the virtual machine is deleted from the virtual disk.

To delete a virtual machine

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Delete from vCenter Server.
A confirmation dialog appears.
3. Click OK.
A message appears confirming the request submission.
4. Click the Summary tab for the virtual machine.
An event should appear confirming the result of the operation. If successful, the virtual machine is deleted from the virtual disk, and the virtual machine disappears from the Explore pane after you refresh the interface.

Deploy a Virtual Machine from a Template

You can deploy a virtual machine from a template to use the template's settings to create and deploy a new virtual machine.

To deploy a virtual machine from a template

1. Right-click VMware vCenter Server in the Explore pane and select Provisioning, Provision VMware VM.
VMware vCenter Provisioning dialog appears.
2. Specify the all required fields and select the appropriate VC Template.
Click Next
3. Perform the remaining steps to specify the virtual hardware for that virtual machine. Click Finish.
A message appears confirming the request submission.
4. Verify that an event confirms the operation.
After the operation completes, the new virtual machine appears in the Explore pane when you refresh the interface.

Manage Cluster Services

You can control the status of the following services on VMware vCenter clusters:

HA

Allows automatic migration and restarting of VMs when a host fails.

DRS

Lets you manage hosts as a collection of resources. The DRS service migrates VMs to hosts and resources to VMs as necessary.

To manage cluster services

1. Select a VMware vCenter cluster on the Explore pane.
The Overview pane appears on the right side, displaying the status of the HA and DRS services.
2. Select Enable or Disable from the drop-down menu.
The status of the service changes.

Migrate a Virtual Machine

You can migrate a virtual machine to move it to another ESX host. You can migrate a powered off machine or migrate a powered on machine with VMotion. You cannot migrate a virtual machine that is suspended.

To migrate a virtual machine

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click the virtual machine to migrate on the Explore pane and select Management, Migration.
The Migration pane appears.
3. Enter the destination ESX server and resource pool for the virtual machine, and click Migrate.
Note: VM migration between ESX hosts is only supported when the VMs datastore/disk is shared between the two ESX hosts.
A confirmation message appears.
4. Click the Summary tab for the virtual machine.
Verify that an event appears to confirm the operation. The virtual machine appears in its migrated location in the Explore pane after the operation completes.

Monitor a Virtual Machine

You can monitor the status and the properties of VMs in detail.

To monitor virtual machines

1. Click Resources.

The Resources page appears.

2. Open the Explore pane.

Available groups, services, and systems appear.

3. Expand the VMware vCenter Server folder and the ESX server object.

A VM list appears.

4. Click the Summary tab.

The right pane displays general information, FT properties, overview, CPU and memory usage, disk usage (logical volumes), and events.

In the Overview panel, the disk states indicate the virtual hardware state of the virtual disk, as calculated by SystemEDGE, and based on monitors configured in the vCenter AIM. This information is based on true performance data of the virtual disk, in terms of reads and writes per second.

In the Disk Usage panel, the disk states indicate the usage of the logical volumes as viewed through the guest operating system. The state is calculated by SystemEDGE, and is based on monitors configured in the vCenter AIM. This information is only valid when the VM and the guest operating system are running.

The General Information panel provides details about the connection state of the VM. Valid connection state values are as follows:

- Not connected
- Connected
- Orphaned

The orphaned connection state can happen during Cluster failover situations. When a virtual machine has been marked as orphaned, states reflected under the Overview section are based upon data collected prior to becoming orphaned.

Monitor an ESX Server

You can monitor the status and the properties of ESX servers in detail.

To monitor ESX servers

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Expand the VMware vCenter Server folder and select an ESX server.
4. Click the Summary tab.
The right pane displays general information, FT attributes, overview, CPU and memory usage, utilization, and events.
5. Click the vNetwork tab.
The right pane displays a list of associated virtual standard switches (vSwitches) and virtual distributed switches (vDS).
6. Select a virtual switch from the list.
The right pane displays the properties of the virtual switch.

Revert to a Snapshot

When you revert to a snapshot, you return the virtual machine to its exact state when the snapshot was taken.

To revert to a snapshot

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Snapshots.
The Snapshots pane appears and displays all existing snapshots for the virtual machine.
3. Select a snapshot and select Revert from the Actions drop-down menu.
A confirmation dialog appears.
4. Click OK.
A confirmation message appears.
5. Click the Summary tab for the virtual machine.
Verify that an event appears to confirm the operation.

Unregister a Virtual Machine

When you unregister a virtual machine from the vCenter Server, the virtual machine still exists but is removed from VMware vCenter Server inventory.

To unregister a virtual machine

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Find and right-click a virtual machine on the Explore pane and select Management, Unregister from vCenter Server.
A confirmation dialog appears.
4. Click OK.
A message appears confirming the request submission.
5. Click the Summary tab for the virtual machine.
An event should appear confirming the result of the operation. If successful, the virtual machine is removed from the vCenter inventory.

vCenter Automation and Policy Actions

The following action types are available for use with VMware vCenter Server:

- [Add Disk](#) (see page 572)
- [Add Network Interface](#) (see page 574)
- [Configure Shares](#) (see page 596)
- [Configure CPU/Memory](#) (see page 583)
- [Configure Power](#) (see page 592)
- [Convert Template to Virtual Machine](#) (see page 597)
- [Convert Virtual Machine to Template](#) (see page 599)
- [Delete Machine](#) (see page 607)
- [Manage VM Snapshots](#) (see page 615)
- [Modify CPU](#) (see page 623)
- [Modify Memory](#) (see page 624)
- [Provision Machine](#) (see page 636)
- [Remove Disk](#) (see page 639)
- [Remove Network Interface](#) (see page 640)
- [Migrate Machine](#) (see page 622)

You can use these action types to create new actions that automate vCenter power, resource allocation, and other operations when assigned rule criteria are met. You can also schedule these actions to occur at specific times.

For more information about using actions and rules to create automation policy, see the "Policy" section.

View Custom Specifications

Custom specifications are custom versions of guest operating systems that you are using on virtual machines. You can view all current custom specifications, the date of the last update, and the current version number.

To view custom specifications

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and select a VMware vCenter server.
The server page appears in the right pane.
3. Click the Configuration tab and select the Customization Specifications submenu.
The Customization Specifications section appears and displays existing customization specifications.

View General Information

CA Server Automation displays General Information in the right hand pane and provides resource properties at the following levels in the object hierarchy:

- vCenter Server
- ESX Server
- Resource Pool
- VM

Resource properties include information about the following categories:

- Names, item types, versions
- Quantitative characteristics of CPU and memory
- Number of VMs and Resource Pools
- Current mode of the resource

Additionally, CA Server Automation displays connection state, power state, and Fault Tolerance information on the VM level.

Valid Fault Tolerance status values are:

- Not Fault Tolerant
- Protected
- Not Protected (Starting)
- Not Protected (Need Secondary VM)
- Not Protected (Disabled)
- Not Protected (VM Not Running)

The Secondary Location values are:

- Not Available
- Total Secondary CPU Usage
- Total Secondary Memory

The General Information panel provides details about whether Fault Tolerance is configured, version, and various counts of supported FT VMs. The counts consider:

- Total Primary VMs
- Total Secondary VMs
- Powered On Primary VMs
- Powered On Secondary VMs

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

More Information

[Monitor an ESX Server](#) (see page 331)

[Monitor a Virtual Machine](#) (see page 330)

VMware vCloud

VMware vCloud Director lets you build secure, multitenant clouds by pooling virtual infrastructure resources into virtual data centers and exposing them to users. CA Server Automation supports VMware vCloud Director management.

vCloud Director resources depend on underlying vSphere resources such as CPU, memory, storage, or vNetwork Distributed Switches to run virtual machines. You can use these underlying vSphere resources to create virtual machines and vApps in vCloud.

A *vCloud Organization* is a unit of administration that represents a collection of users, groups, and computing resources. Associated virtual datacenters provide the required computing resources. After users authenticate at the organization level, they can create, use, and manage virtual machines or vApps.


A *virtual datacenter (vDC)* provides virtual computing resources to a vCloud organization. You can provision, run, and store virtual systems in a virtual datacenter. A vCloud organization can have multiple virtual datacenters.

Organizations provide *catalogs* to store vApp templates and media files. The members of an organization can use the vApp templates and media files in the catalog to create their own vApps.

How to Configure the vCloud Director Management Components

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 vCloud Server Management Components



Follow these steps:

[Review vCloud Requirements](#) (see page 338)
[Interactions Between vCloud Management Components](#) (see page 339)
[Add a New vCloud Director Connection to the Manager](#) (see page 340)
[Troubleshoot the vCloud Server Connection](#) (see page 341)
[vCloud Server Connection Failed](#) (see page 342)
[Add the AIM Instance for the vCloud Server](#) (see page 344)
[Troubleshoot the vCloud AIM Instance Connection](#) (see page 345)
[Verify the VMware vCloud Folder in the Resources Tree](#) (see page 349)

Review vCloud Requirements

Review the following requirements before you start configuring the vCloud Director management components of CA Server Automation:

- 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, VMware vSphere, and VMware vCloud.
- You can access a CA Server Automation manager installation that includes the VMware Platform Management Module (PMM), vCloud Application Insight Module (AIM), and Monitoring Agent (CA SystemEDGE).
- You can access the CA Server Automation user interface.
- You have valid credentials available (user name and password) to access the vCloud Director server that you want to manage.
- You have found out which protocol (HTTP or HTTPS) and port to use for accessing the vCloud Director through web services. Default: HTTPS, Port 443
- You have verified that the vSphere environment and the vCloud Director run properly.
- If the VMware PMM and vCloud AIM are installed on different systems, you have verified that the SNMP settings on these systems are consistent. Read and write community strings and SNMP port number must be identical.
- You have verified that the CA Server Automation manager has discovered any remote vCloud AIM Servers that you want to use.


More information:

[Add a New vCloud Director Connection to the Manager](#) (see page 340)
[Add the AIM Instance for the vCloud Server](#) (see page 344)
[Verify the VMware vCloud Folder in the Resources Tree](#) (see page 349)

Interactions Between vCloud Management Components

The following diagram illustrates how the components involved in vCloud Director management interact. SystemEDGE and the vCloud AIM run on a Windows server. The AIM communicates with one or more remote vCloud Director servers to manage the virtual environment. The vCloud AIM collects the data for an entire view of the virtual resources associated with the vCloud Director. The underlying vSphere environment provides the required resources to run virtual machines and vApps.

Interaction Between vCloud Director Management Components



You can configure vCloud management through the Administration tab of the user interface.

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

Add a New vCloud Director Connection to the Manager

You can add a vCloud Director 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 vCloud Server from the Provisioning section in the left pane.

The right pane refreshes and displays the managed vCloud Servers, associated vCloud AIM Servers, and the AIM Instance for the vCloud Server.

3. Click  (Add) on the vCloud Servers pane toolbar.

The New vCloud Server dialog appears.

4. Enter the required connection data (server name, username, password, protocol, port), specify the preferred AIM, enable Managed Status (checkbox), and click OK.

When specifying the username, you can use the following syntax to consider user roles and access levels:

- System Administrator (Full access): administrator@System
- Limit operation at the organization level and the role assignment (Organizational Access) *username@organization_name*

If the network connection has been established successfully, the vCloud Server is added to the top right vCloud Servers pane with a green status icon. CA Server Automation discovers the vCloud Server automatically.

If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Server Automation adds the vCloud Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added. For troubleshooting the connection, see [Troubleshoot the vCloud Server Connection](#) (see page 341).

More information:

[Add the AIM Instance for the vCloud Server](#) (see page 344)


[Verify the VMware vCloud Folder in the Resources Tree](#) (see page 349)

[Troubleshoot the vCloud Server Connection](#) (see page 341)

Troubleshoot the vCloud Server Connection

The vCloud Server connection has failed. Follow the troubleshooting information indicated in the following diagram:

How to Troubleshoot the vCloud Server Connection



Follow these steps:

[vCloud Server Connection Failed](#) (see page 342)

[Add the AIM Instance for the vCloud Server](#) (see page 344)

[Verify the VMware vCloud Folder in the Resources Tree](#) (see page 349)

vCloud Server Connection Failed

Symptom:

After I have added a new vCloud Server connection under Administration, Configuration, the validation of the connection to the vCloud Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used vCloud Server connection data (server name, user, password, protocol, port) is still valid. If necessary, update the connection data.
- Verify, if the vCloud Server system is running and accessible.

To update the vCloud Server connection data

1. Click  (Add) or  (Edit) that is associated with the failed connection.

The New or Edit vCloud Server dialog appears.

2. Add the valid server name, user, password, protocol, and port. Specify the preferred AIM. Enable Managed Status and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the vCloud Server cannot be established, continue with the next procedure.

To verify, if the vCloud Server system is running and accessible

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <vCloud Server Name>  
ping <IP Address of vCloud Server>
```

2. Verify the output of the commands to find out whether the vCloud Server has a valid DNS entry and IP address.

If the vCloud Server is not in the DNS, add the vCloud Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the vCloud Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

ipaddress <vCloud Server Name>

Enter the correct IP address and vCloud Server name. For example:

192.168.50.50 myvCloud

4. Click  (Validate) in the upper-right corner.

If the connection to the vCloud Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

Work with the vCloud administrator or VMware support to fix the vCloud Server connection problem.

Add the AIM Instance for the vCloud Server

After adding a new vCloud Server connection to the CA Server Automation manager, add a vCloud AIM instance to manage the new vCloud Server. CA Server Automation then discovers the entire vCloud environment with all its virtual components, such as Organizations, vApps, VMs, and so on.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select vCloud Server from the Provisioning section in the left pane.

The right pane refreshes and displays the managed vCloud Servers, associated vCloud AIM Servers, and the AIM Instances for managed vCloud Servers.

3. Click  (Add) on the vCloud AIM Servers pane toolbar.

The New vCloud AIM Server dialog appears.

4. Open the vCloud AIM Server drop-down list.

The list of discovered vCloud AIM Servers appears. If you have installed the vCloud AIM on the local system, the name of the local system appears in the list too.

5. Select a vCloud AIM Server from the drop-down list.

CA Server Automation populates the vCloud Server drop-down list with the vCloud Servers listed in the vCloud Servers pane. That is, you can only manage those vCloud Servers for which your CA Server Automation manager has a valid connection established.

6. Select the vCloud Server you want to manage and click OK.

A new AIM instance for the selected vCloud Server is added. If the instance is not in an error or stopped state, CA Server Automation starts to discover the associated vCloud environment. When the discovery process is complete, you can start managing the virtual resources of vCloud.

More information:

[Verify the VMware vCloud Folder in the Resources Tree](#) (see page 349)

[Troubleshoot the vCloud AIM Instance Connection](#) (see page 345)


Troubleshoot the vCloud AIM Instance Connection

The vCloud AIM Connection is in not-ready status. One of the following status icons appears:

-  Discovery in progress - Wait until the platform manager synchronizes all data.
-  Error - Unable to connect to the AIM. Check the network configuration.
-  No Polling - The CA Server Automation manager does not poll this AIM instance.
-  Disabled - This instance is not managed.

Follow the troubleshooting information indicated in the following diagram:

How to Troubleshoot the vCloud AIM Instance Connection



More information:

[vCloud AIM Instance Status Icon Shows Discovery in Progress](#) (see page 347)

[vCloud AIM Instance Status Icon Shows Error](#) (see page 347)

[vCloud AIM Instance Status Icon Shows No Polling](#) (see page 348)

[vCloud AIM Instance Status Icon Shows Disabled](#) (see page 349)

vCloud AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I have added a vCloud AIM instance for a vCloud Server under Administration, Configuration, the status icon shows  (Discovery in Progress).

Solution:

Wait until the discovery process of the vCloud environment has completed. The discovery duration depends on the amount of managed objects related to virtual resources in vCloud. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job has completed, CA Server Automation adds a vCloud Server folder to the Resources tree. Then you can start managing vCloud and its entire virtual infrastructure.

vCloud AIM Instance Status Icon Shows Error

Symptom:

After I have added a vCloud AIM instance for a vCloud Server under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the vCloud AIM:

- Verify, if the vCloud AIM Server is accessible.
- Verify, if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify, if the vCloud AIM server system is accessible

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
ping servername
```
2. Verify the output of the commands to find out whether the vCloud AIM server has a valid DNS entry and IP address.

If the vCloud AIM server is not in the DNS, add the vCloud AIM server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the vCloud Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

ipaddress servername

Enter the correct IP address and vCloud AIM server name. For example:

192.168.50.51 myvCloudAIM

4. Click  (Validate) in the upper-right corner of the vCloud AIM Server pane.
If the error status remains unchanged, continue with the next procedure.

To verify, if SystemEDGE is running

1. Log in to the vCloud AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.

The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.

2. Start or restart SystemEDGE.

Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.

3. Change to the CA Server Automation user interface, vCloud AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the vCloud AIM Server connection.

If the error status remains unchanged, verify whether the data you gathered according to the requirements for this scenario is still valid.

vCloud AIM Instance Status Icon Shows No Polling

Symptom:

After I add a vCloud AIM instance for a vCloud Director under Administration, Configuration, the status icon shows  (no polling).

Solution:

No specific actions are required for the associated instance. This icon informs you that the CA Server Automation manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular vCloud Director, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

vCloud AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered vCloud AIM instances in the network, the status icons of several instances show  (Disabled). This vCloud AIM instance is not managed.

This status appears, if CA Server Automation has discovered a vCloud AIM with the following relationships:

- The vCloud AIM is configured for a vCloud Server that has a valid connection to the CA Server Automation manager but is in unmanaged state.
- The AIM is connected to a vCloud Server that has not been configured in the vCloud Servers pane.

Solution:

To change the status of the AIM instance to ready, do one of the following:

- Add the missing vCloud Server connection to the CA Server Automation manager.
- Edit the existing vCloud Server connection and change its managed status to enabled.

Verify the VMware vCloud Folder in the Resources Tree

After a successful configuration and discovery, the new vCloud Server is listed in the Resources Explore pane under the VMware vCloud folder.

Follow these steps:

1. Click Resources, Explore.
The Resources tree appears.
2. Expand VMware vCloud.
The managed vCloud Director Servers appear.
3. Expand the new vCloud Director Server entry.
The managed vCloud infrastructure appears: Organizations, vApps, VMs, ...

CA Server Automation is now ready to manage the added vCloud environment with its virtual infrastructure.

Remote and Multi-instance vCloud Director Support


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

Each vCloud Director is uniquely associated with one Preferred vCloud AIM that you specify during configuration. Setting the preferred AIM indicates which AIM should be used for polling if multiple AIMs manage one vCloud Director.

vCloud Folder Structure

After a successful configuration of the connection to a vCloud Director server, CA Server Automation discovers the vCloud Director environment that consists of organizations, vApps, and virtual machines. When the discovery completes, the VMware vCloud folder appears in the Explore pane of the Resources tab. You can expand the folder and can manage your vCloud environment.

The following diagram shows the object hierarchy underneath the VMware vCloud folder.



The VMware vCloud folder represents the service level at the top. The vCloud service can consist of multiple VMware vCloud Director servers. Each vCloud Director usually has multiple organizations with vApps and virtual machines.

At the organization level, you can provision vApps based on templates stored in a catalog.

vApp Support in vCloud

The concept of vApps that are created in vCloud or vSphere environments is similar. Both types represent an application object that can be operated on as a single entity. Usually, a vApp contains multiple VMs each with its own purpose to the complete vApp application or service that it provides to the end user. Operations that are performed on the vApp are also performed on all VMs in the vApp. For example, both types define start and stop orders for all VMs in a vApp and define CPU and memory resource limits that all VMs in the vApp can use.

The purpose of vApps in vCloud is to be able to define an application or service once as a template, and make it accessible to multiple organizations through the organizations catalog. vCloud stores its data in the vCloud database which is different from the vCenter Server database.

Important! Do not operate on the VMs defined in vCloud directly from a vCenter Server. Those operations can cause the vCloud database to become out-of-sync with the actual defined VM. CA Server Automation provides a limited set of operations for those VMs which appear under vCloud and vCenter Server so that the databases do not become out-of-sync.

Differences between vCloud and vSphere vApps

- A vCloud vApp does not provide the ability for a nesting hierarchy. vSphere vApps can contain other vApps and Resource Pools.
- In vCloud, CPU and memory resource limits are defined through Virtual Data Centers (vDC), and the vApp is mapped to one of those Virtual Data Centers.
In vSphere, vApp resource limits are defined on the vApp itself.
- vCloud vApps can contain VMs that are defined on many different vCenter Servers and ESX Hosts.
VMs in vSphere vApps are limited to VMs in a particular data center and cluster.
- vCloud vApps have lease limits. You can define a runtime and storage limit on the vApps. When the runtime limit is reached, a vCloud vApp can no longer be used. When the storage limit is reached, the vApp is deleted from the vCloud or is moved to the Expired Items folder, depending on organization lease policy.
vSphere vApps remain in existence until a user manually deletes them.
- vCloud vApps are created from vApp templates. vApp templates are created by importing a VM from a vCenter Server or by importing an OVF package. vApps are created by deploying the template to the cloud for the organization on which the template was created. After deployment, additional VMs can be moved into the vApp.
vSphere vApps are created by defining a vApp along with the CPU and memory resource limits desired. Then VMs for the data center where the vApp is defined can be moved into the vApp.

vCloud Organizations

A vCloud organization is a unit of administration for a collection of users, groups, and computing resources. Organizations provide catalogs to store vApp templates and media files. The members of an organization can use the vApp templates and media files in the catalog to create their own vApps.

A virtual data center (vDC) provides virtual computing resources to a vCloud organization. You can provision, run, and store virtual systems in a virtual data center. A vCloud organization can have multiple virtual data centers.

Provision vApp from Template

From the vCloud Organization level, you can provision vApps from templates which are stored in catalogs.

Follow these steps:

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Expand the VMware vCloud folder.
The vCloud folder structure appears.
4. Right-click the organization object.
The Provisioning pop-up menu appears.
5. Click Provision vApp from Template.
The Provision new vApp from Template dialog appears.
6. Specify Name, vApp Template, Deployment lease, and Storage lease. Click OK.
CA Server Automation creates a vApp in the organization.

Operations on vApps in vCloud

From the vCloud Organization level, you can provision vApps from templates which are stored in catalogs.

Follow these steps:

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.

3. Expand the VMware vCloud folder.
The vCloud folder structure appears.
4. Right-click a vApp object.
The Management pop-up menu appears.
5. Select one of the available operations.

Power On vApp

Powers On the vCloud vApp.

Power Off vApp

Powers Off the vCloud vApp.

Reset vApp

Resets the vCloud vApp.

Suspend vApp

Suspends the vCloud vApp.

Resume vApp

Resumes the vCloud vApp.

Clone vApp

Creates a vCloud vApp from an existing vApp.

Move vApp

Moves a vCloud vApp to another virtual datacenter.

Delete vApp

Deletes the vCloud vApp.

Modify vApp Lease

Modifies the deployment and storage lease.

Specify the required parameter values and click OK.

6. Click Events to verify the new status of the vApp.
The list of events appears.

IBM PowerVM (LPAR)

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

The LPAR AIM enables SystemEDGE to monitor PowerVM resources.

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

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

Server level

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

- Provision LPARs
- Delete LPARs

Power operations level

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

- Activate LPARs
- Shutdown LPARs
- Restart LPARs

Resource adjustments level


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

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

How to Configure the PowerVM Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the Management Components



Follow these steps:

[Review Requirements](#) (see page 356)

[Interaction Between AIX LPAR Management Components](#) (see page 357)

[Add an HMC or an IVM Server Connection to the Manager](#) (see page 361)

[Server Connection to the Manager Failed](#) (see page 362)

[Add the LPAR AIM Instance](#) (see page 363)

[Troubleshoot the AIM Instance Connection](#) (see page 365)

[Verify the Group in the Resources Tree](#) (see page 368)

Review Requirements


Review the following requirements before you start configuring the management components of CA Server Automation:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Server Automation, SystemEDGE, and your environment.
- You can access a CA Server Automation manager installation that includes:
 - Platform Management Module (PMM)
 - Application Insight Module (AIM)
 - Monitoring Agent (SystemEDGE)
- You can access the CA Server Automation user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Server Automation manager discovered remote AIM Servers that you want to use.

Interaction Between AIX LPAR Management Components

The following diagram illustrates how the components involved in IBM LPAR management interact. The AIM Server is a Windows Server on which SystemEDGE and the LPAR AIM run. The communication between the AIM and the HMC/IVM Server is based on SSH (Secure Shell). Because CA 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



After the installation, configure your environment by adding the required connection information for each required HMC/IVM and Virtual I/O Server. Use *one* of the following methods:

- Administration tab of the user interface
- NodeCfgUtil.exe utility on the AIM Server

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

IBM PowerVM Configuration Use Cases

The following use cases describe the handling of LPAR AIM instance entries for managed PowerVM environment in the Administration tab:

- You add an HMC Server and an LPAR AIM instance.

The AIM discovers:

- Power Systems associated with the HMC.
- Virtual I/O Servers associated with the Power systems. The AIM applied the default VIOS credentials specified when adding the HMC.

Important! If you do not specify the default VIOS credentials for an HMC Server, provide the VIOS credential for *each* VIOS in the Virtual I/O Servers panel to complete the configuration of the discovered VIOS. If the default VIOS credentials do not apply to a particular VIOS, you can overwrite the credentials in the Virtual I/O Servers panel.

- Preferred AIM

Two AIMs can manage one HMC. When second AIM is added it becomes the redundant AIM, the AIM added first becomes the preferred AIM. The status of the HMC under the redundant AIM becomes Suspended. This status reflects that the HMC is managed by the preferred AIM. You can change the preferred AIM in the HMC/IVM Servers panel.

- Dual HMC feature supports configurations where a P System is associated with two HMC servers.

P-Server and associated HMC servers are one atomic management entity and as such have to be managed by one AIM. The dual HMC configuration is only supported in the scope of one AIM. For example, one Power System, P1 is connected to two HMC servers, HMC1 and HMC2. Both HMC servers are managed by one AIM, AIM1.

- Dual HMC failover

If the preferred HMC fails, the redundant HMC starts managing your System automatically. The redundant HMC becomes the current one. However, when the preferred HMC becomes available, the current HMC is not changed. To manage your System by the preferred HMC again, change the current HMC in the LPAR AIM Instances panel on the Administration, Configuration tab.

Note: If a preferred HMC fails, the redundant HMC manages your System. After the failover, you can change the current HMC for your System manually.

- You do not specify the default VIOS credentials or enter incorrect Virtual I/O Server credentials.

The user interface displays a message about the failure of the operation. When you try to apply incorrect Virtual I/O Server credentials () , the Virtual I/O Server changes to "authentication failed" state. The managed system instance changes from "pending VIOS" state to "out of date" due to connection problems.

- You add a managed system instance without Virtual I/O Servers.

The managed system instance appears in the instances table in "ready" state.

- You add a P-server to a managed system.

The new P-server and VIOS are discovered automatically. If VIOS credentials match the default VIOS credentials, no configuration is required. If the VIOS credentials do not match the default VIOS credentials, set the VIOS credentials on AIM instances () .

- You remove a Virtual I/O Server in "invalid configuration" state from a managed system.

The LPAR AIM removes the corresponding record from the instance table and the managed system changes to 'ready' state.

- You remove a Virtual I/O Server in "ready" state from a managed system.

The LPAR AIM removes the corresponding record from the instance table.

- You remove a managed system instance with one or two Virtual I/O Servers.

The managed system instance and associated Virtual I/O Server entries disappear from the instances table.

- The IBM PowerVM administration pane displays status information through icons and tooltips.

Detailed tooltips become visible when you hover the cursor over warning and error icons.

The following icons can appear:

 Discovery in progress

 No polling

 Error

 Warning

 Disabled

 Unknown

Add an HMC or an IVM Server Connection to the Manager

You can add an HMC or an IVM Server connection using the Administration tab of the 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 IBM PowerVM from the Provisioning section in the left pane.

The right pane refreshes and displays the managed HMC and IVM Servers, associated Virtual I/O Servers, and the LPAR AIM Servers.

3. Click  (Add) on the HMC/IVM Servers pane toolbar.

The New HMC/IVM Server dialog appears.

4. Enter the required connection data (server name, user, password), specify the preferred AIM, enable Managed Status (checkbox).

Note: The preferred AIM field is active only if you specify more than one AIM instance for a given HMC or IVM server.

5. (Optional) Specify the Virtual I/O Servers Default Credentials.

The default VIOS credentials apply to newly discovered VIO servers.

Important! If you do not specify the default VIOS credentials for an HMC Server, provide the VIOS credential for *each* VIOS in the Virtual I/O Servers panel to complete the configuration of the discovered VIOS. If the default VIOS credentials do not apply to a particular VIOS, you can overwrite the credentials in the Virtual I/O Servers panel.

6. Click OK.

If the network connection has been established successfully, the Server is added to the top right HMC/IVM Servers pane with a green status icon. CA Server Automation discovers the HMC/IVM Server automatically.

If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Server Automation adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Server Connection to the Manager Failed

Symptom:

After I have added a Server connection under Administration, Configuration, the validation of the connection to the Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Server connection data is still valid. If necessary, update the connection data.
- Verify, if the Server system is running and accessible.
- Verify, if the Management Service on the Server system is running properly.

To update the Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, enable Managed Status, and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the Server cannot be established, continue with the next procedure.

To verify if the Server system is running and accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <Server Name>  
ping <IP Address of Server>
```

2. Verify the output of the commands to find out whether the Server has a valid DNS entry and IP address.

If the Server is not in the DNS, add the Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

ipaddress <Server Name>

Enter the correct IP address and Server name. For example:

192.168.50.50 myServer

4. Click  (Validate) in the upper-right corner.

If the Server credentials and connection data are correct and you can ping the Server, the connection can still fail. In this case, it is possible that the Server causes the problem. If the connection to the Server cannot be established, continue with the next procedure.

To verify, if the Management Service on the Server system is running properly:

1. Contact the Administrator to access the Server system.
2. Log in to the Server system and open Administrative Tools, Services from the Start menu.

The Services window opens.

3. Select the service and start or restart the service.
4. Change to the CA Server Automation user interface, Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the Server connection.

If the connection to the Server fails, verify the validity of the data you gathered according to the requirements for this scenario.

Work with the administrator or support to fix the Server connection problem.

Add the LPAR AIM Instance

After adding an HMC or an IVM Server connection to the CA Server Automation manager, add an AIM instance to manage the new Server. CA Server Automation then discovers the PowerVM environment.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select IBM PowerVM from the Provisioning section in the left pane.

The right pane refreshes and displays the managed HMC and IVM Servers, associated Virtual I/O Servers, and the LPAR AIM Servers.

3. Click  (Add) on the LPAR AIM Servers pane toolbar.

The New LPAR AIM Server dialog appears.

4. Select the LPAR AIM Server from the drop-down list.

The list of discovered LPAR AIM Servers appears. If you have installed the LPAR AIM on the local system, the name of the local system appears in the list too.

5. Select the HMC or IVM Server from the drop-down list.

CA Server Automation populates the HMC/IVM Server drop-down list with the HMC and IVM Servers listed in the HMC/IVM Servers pane. You can only manage those HMC or IVM Servers for which your CA Server Automation manager has a valid connection established.

Note: If the AIM resides on a remote system, CA Server Automation must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

A new AIM instance for the selected Server is added. If the instance is not in an error or in a stopped state, CA Server Automation starts to discover the associated PowerVM environment:

- For each HMC server, the AIM discovers all Power Systems and the Virtual I/O servers.
- For each IVM server, the AIM discovers a Power System that the IVM manages.

When the discovery process is complete, you can start managing your PowerVM environment.

The Administration Tab shows an aggregated state for all the P Systems and VIO Servers the AIM discovered. To view their individual configuration state press the Show Managed Systems  icon.

Change the Preferred HMC for the Managed Power System

If your Power System uses a dual HMC, you can change the preferred HMC.

Important! Verify that one LPAR AIM manages both the primary and the redundant HMC servers.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.
The Configuration page appears.
2. Select IBM PowerVM from the Provisioning section in the left pane.
The right pane refreshes and displays the managed HMC and IVM Servers, associated Virtual I/O Servers, and the LPAR AIM Servers.
3. Click  (Configure Managed/VIO Servers) that is associated with the HMC Server.
The IBM PowerVM dialog with Managed/VIO Servers appears.
4. Click  (Switch the preferred HMC) under the Actions row and confirm.
The redundant HMC is set as the preferred HMC.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

-  Discovery in progress
-  No polling
-  Error
-  Warning
-  Disabled
-  Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the amount of managed objects that are related to virtual and physical resources in your environment. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Server Automation adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Server Automation manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the AIM:

- Verify if the AIM Server is accessible.
- Verify if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
ping servername
```

2. Verify the output of the commands to find out whether the AIM server has a valid DNS entry and IP address.

If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress servername
```

Enter the correct IP address and AIM server name. For example:

```
192.168.50.51 myAIM
```

4. Click  (Validate) in the upper-right corner of the AIM Server pane.

If the error status remains unchanged, continue with the next procedure.

To verify if SystemEDGE is running:

1. Log in to the AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.

The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.

2. Start or restart SystemEDGE.

Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.

3. Change to the CA Server Automation user interface, AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the AIM Server connection.

If the error status remains unchanged, verify the validity of the data you gathered according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Server Automation discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Server Automation manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Server Automation manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Group in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand IBM PowerVM group.
The managed HMC and IVM Servers appear.
3. Expand the HMC or IVM Server entry.
The managed systems appear.

CA Server Automation is now ready to manage the added PowerVM environment with its virtual infrastructure.

calpara.xml File

The main purpose of the calpara.xml file is to store LPAR AIM configuration data, such as: persistent data and default values. The settings for monitoring can be adjusted to specific environments.

This document is intended for system administrators, who are familiar with xml format. We recommended caution when changing this file. To change the calpara.xml file, stop SystemEDGE before and start SystemEDGE again after changing the file.

Important! If you require Monitoring Threshold, Lag, or Severity adjustment, change only the Default Values. Modify Poll Group and DisableOutOfDate setting only when CA Support asks you to do so.

The calpara.xml file is at the following location:

```
<SystemEDGE_InstallDir>\plugins\calpara\calpara.xml
```

Persistent Data

Persistent data is available the next time the AIM starts up. This data can change during the AIM lifetime and is user settable, by SNMP set requests.

The following list presents the examples of persistent data:

- Instances
- Systems
- Partitions
- Slots
- Poll Groups

Instances

For each configured Instance in the Instance Table (lparAimInstanceTable) a section similar to the following example is stored:

```
<ManagedInstance>
  <InstIndex>7</InstIndex>
  <SerialNr>1010101</SerialNr>
  <ServerName>vios1.company.com</ServerName>
  <ServerType>vios</ServerType>
  <RowStatus>1</RowStatus>
</ManagedInstance>
```

ServerType

Specifies one of the following server types: hmc, vios, or ivm.

RowStatus

Specifies the status as active (1) or notInService (2).

Systems

For each managed Power System in the System Table (lparAimStatSysTable) a section similar to the following example is stored inside the related ManagedInstance section:

```
<System>
  <MonitorIndices>530091,530092,530093,530094,530095</MonitorIndices>
</System>
```

MonitorIndices

Stores the indices of the SystemEDGE monitors that the AIM created to monitor the operational status, CPU, and Memory usage of the Power System.

Note: If the AIM does not manage the Power System any more, the corresponding monitors are deleted.

Partitions

For each managed Logical Partition in the Partition Table (lparAimStatLPTable) four corresponding entries similar to the following examples are stored in the Partitions section inside the related ManagedInstance section:

```
<Partitions>
...
  <LparIndex>7</LparIndex>
  <LparId>7</LparId>
  <LparName>LPAR12345</LparName>
  <MonitorIndices>530141,530142,530143,530144,530145</MonitorIndices>
...
</Partitions>
```

Note: If the AIM does not manage the Power System any more, the corresponding monitors are deleted.

Slots

For each Physical Slot in the Slot Table (lparAimStatSlotTable) four corresponding entries similar to the following examples are stored in the Slots section inside the related ManagedInstance section:

```
<Slots>
...
  <SlotIndex>3</SlotIndex>
  <DRCName>U787B.001.DNFFF77-P1-C3</DRCName>
  <DRCIndex>553713666</DRCIndex>
  <SlotName>C3</SlotName>
...
</Slots>
```

The LPAR AIM uses this data to detect any Slot-related change right after start-up, potentially causing to send a corresponding SNMP trap.

Poll Groups

A Poll Group is a set of related commands, all executed at the same poll interval. For each Poll Group in the Poll Table (lparAimPollTable) three corresponding entries in the respective Poll Group section are stored. The following example shows the Basic Poll Group:

```
<Basic>
  <PollDefault>5</PollDefault>
  <PollSpecific>30,30</PollSpecific>
  <PollInstances>4,6</PollInstances>
</Basic>
```

PollDefault

Stores the default poll interval (in minutes) to apply to all instances, except those instances listed in PollInstances (listing the indices of the instances).

PollSpecific

Stores a list of poll intervals (in minutes) to apply one-to-one to the corresponding list of instances stored in PollInstances.

Note: Initially PollSpecific and PollInstances are empty.

Default Values

The following section describes default values stored in calpara.xml file that specify lags, thresholds, and severities, the AIM uses when creating new SystemEDGE monitors.

Important! Default values cannot change during the AIM lifetime and are not user settable.

```
<LowestPollInterval>5</LowestPollInterval>
<DisableOutOfDate>0</DisableOutOfDate>
<MonitorIndexStart>530001</MonitorIndexStart>
<SysAliveSev>fatal</SysAliveSev>
<CpuLagValue>3</CpuLagValue>
<CpuThresh1Val>95</CpuThresh1Val>
<CpuThresh1Sev>warning</CpuThresh1Sev>
<CpuThresh2Val>98</CpuThresh2Val>
<CpuThresh2Sev>critical</CpuThresh2Sev>
<MemLagValue>2</MemLagValue>
```

```
<MemThresh1Val>95</MemThresh1Val>  
<MemThresh1Sev>warning</MemThresh1Sev>  
<MemThresh2Val>98</MemThresh2Val>  
<MemThresh2Sev>critical</MemThresh2Sev>
```

LowestPollInterval

Stores the lowest allowed poll interval (in minutes).

DisableOutOfDate

Specifies whether the data status (lparAimInstDataStatus) outOfDate is excluded.

Note: Set the variable to 1 to disable the data status becoming outOfDate(7) if any command execution failed.

Default: 0

MonitorIndexStart

Specifies the index of the first SystemEDGE monitor the AIM creates after startup.

Note: When creating new monitors, the AIM always searches for the next free index being equal or greater.

SysAliveSev

Specifies the severity of the SystemEDGE monitor that the AIM creates for monitoring the operational status of a Power System or a Logical Partition.

Valid values: ok, warning, minor, major, critical, fatal.

Note: Changing this value has no effect on existing monitors.

CpuThresh1Val and CpuThresh2Val

Specify the threshold values of the two SystemEDGE monitors that the AIM creates for monitoring the CPU usage of a Power System or a Logical Partition.

Limits: 0 to 100.

Note: Changing this value has no effect on existing monitors.

CpuThresh1Sev and CpuThresh2Sev

Specify the severities of the two SystemEDGE monitors that the AIM creates for monitoring the CPU usage of a Power System or a Logical Partition.

Valid values: ok, warning, minor, major, critical, fatal.

Note: Changing this value has no effect on existing monitors.

MemThresh1Val and MemThresh2Val

Specify the threshold values of the two SystemEDGE monitors the AIM creates for monitoring the Memory usage of a Power System or a Logical Partition.

Limits: 0 to 100.

Note: Changing this value has no effect on existing monitors.

MemThresh1Sev and MemThresh2Sev

Specify the severities of the two SystemEDGE monitors that the AIM creates for monitoring the Memory usage of a Power System or a Logical Partition.

Valid values: ok, warning, minor, major, critical, fatal.

Note: Changing this value has no effect on existing monitors.

CpuLagValue and MemLagValue

Specify the lag values of the SystemEDGE monitors that the AIM creates for monitoring the CPU and Memory usage of a Power System or a Logical Partition. The lag value specifies the number of consecutive poll intervals (Basic Poll Group) that the monitor reaches the threshold before the monitor changes its status.

Note: Changing this value has no effect on existing monitors.

IBM PowerVM Server Administration Overview

The CA IBM PowerVM component of CA Server Automation lets you monitor and manage IBM PowerVM resources. The monitored and managed resources consist of the following types:

- Hardware Management Console (HMC)
- Integrated Virtualization Manager (IVM)
- Virtual I/O Server (VIOS)
- Managed System (POWER Server)
- Logical Partition (LPAR)

The *Hardware Management Console (HMC)* is an external appliance that is used to perform management tasks on IBM PowerVM Systems. HMC can be used to create or change logical partitions, including dynamically assigning resources to a partition. The HMC communicates with the server firmware layers of POWER Systems, providing a single point of control in large PowerVM environments.

The *Integrated Virtualization Manager (IVM)* is an enhancement of the Virtual I/O Server (VIOS) and allows you to manage a single POWER System. IVM lets you create and manage LPARs. IVM enables management of VIOS functions and provides a web-based user interface.

A *Virtual I/O Server (VIOS)* is a special logical partition that is configured to own all physical I/O resources and provides its virtualization capabilities to other LPARs. LPARs access disk, network, and optical devices through the Virtual I/O Servers as virtual devices. Each PowerVM system with virtualized resources has a Virtual I/O Server.

A *Logical Partition (LPAR)* is a subset of hardware resources, virtualized as a separate system. A physical system can be partitioned into multiple LPARs, each providing a separate operating system and applications. The number of logical partitions depends on the hardware configuration of the system. LPARs communicate in the network as separate systems.

To manage IBM PowerVM resources, provide SSH access credentials to HMC/IVM Servers and Virtual I/O Servers.

You can configure CA Server Automation to manage PowerVM resources by using CA Server Automation Administration, Configuration, Provisioning, the IBM PowerVM Group.

The following panels are available:

- HMC/IVM Servers
- Virtual I/O Servers
- LPAR AIM Servers

The LPAR AIM Server is the system on which SystemEDGE and the LPAR AIM run. The LPAR AIM can run on the local CA Server Automation manager system or on a remote Windows server. The LPAR AIM is a multi-instance AIM that can connect to multiple HMCs or IVMs. Once the AIM starts managing an HMC or IVM server, the AIM discovers and manages all P-Servers that are connected to this HMC or IVM server.

LPAR Monitoring

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

Example

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

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

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

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

mon Index	monOID	mon ObjClass	monObjInstance	mon ObjAttribute	mon Severity	mon CurrState
530001	lparAimStatSys Status.1	System	lpar://System:SerialNumber/Total	Alive	critical	ok
530002	lparAimStatSys CPUUsage PerMil.1	CPU	lpar://System:SerialNumber/Total	PercentUsed	warning	ok
530003	lparAimStatSys CPUUsage PerMil.1	CPU	lpar://System:SerialNumber/Total	PercentUsed	minor	ok
530004	lparAimStatSys MemoryUsage PerMil.1	Memory	lpar://System:SerialNumber/Total	PercentUsed	warning	warning
530005	lparAimStatSys MemoryUsage PerMil.1	Memory	lpar://System:SerialNumber/Total	PercentUsed	minor	minor
530006	lparAimStatLP Status.1.1	System	lpar://System:SerialNumber/lpar01/Total	Alive	critical	critical

mon Index	monOID	mon ObjClass	monObjInstance	mon ObjAttribute	mon Severity	mon CurrState
530007	lparAimStatLPCPU Usage.1.1	CPU	lpar://System:SerialNumber/lpar01/Total	PercentUsed	warning	ok
530008	lparAimStatLPCPU Usage.1.1	CPU	lpar://System:SerialNumber/lpar01/Total	PercentUsed	minor	ok
530009	lparAimStatLP MemoryUsage.1.1	Memory	lpar://System:SerialNumber/lpar01/Total	PercentUsed	warning	ok
530010	lparAimStatLP MemoryUsage.1.1	Memory	lpar://System:SerialNumber/lpar01/Total	PercentUsed	minor	ok
530011	lparAimStatLP Status.1.2	System	lpar://System:SerialNumber/lpar02/Total	Alive	critical	critical
530012	lparAimStatLPCPU Usage.1.2	CPU	lpar://System:SerialNumber/lpar02/Total	PercentUsed	warning	ok
530013	lparAimStatLPCPU Usage.1.2	CPU	lpar://System:SerialNumber/lpar02/Total	PercentUsed	minor	ok
530014	lparAimStatLP Memory Usage.1.2	Memory	lpar://System:SerialNumber/lpar02/Total	PercentUsed	warning	ok

Add a Logical Partition for an IBM AIX Computer

You can use the Provisioning wizard to manage the logical partitions on 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, for an HMC server, the profile name. Specify the minimum, desired, and maximum memory for the partition. Click Next.

The Processors page appears.

4. Specify whether to allocate partial processor units or dedicated processors and the minimum, desired, and maximum processors units. Advanced settings are available for shared modes and virtual processors. Click Next.

The I/O Components page appears.

5. Select the I/O devices to associate with the partition, and click Next.

Note: For each I/O device, you can specify that the I/O device is required or optional for logical partition activation. If the I/O device is required, the partition cannot be activated if the I/O device is unavailable or in use by another logical partition. If the I/O device is optional, and if the desired I/O device is available when the partition is activated, the managed system commits the I/O device to the partition. If the optional I/O device is not available, the managed system skips the I/O device.

The I/O Pools page appears.

6. (Optional) To create a new I/O pool, click the + (Add) on the I/O Pools table, enter a numerical value, and click Save.

Note: When you add an I/O device to a partition, 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.

If an IVM server was selected, the Virtual Ethernet page appears. Proceed to Step 10.

8. (Optional) Specify the maximum virtual adapters for the partition. To create a new virtual serial adapter, click + (Add) and specify the Adapter ID, Remote Partition, and Remote Slot Number. You can require that the virtual adapter must be allocated and the managed system must have enough memory to run the required virtual adapters for the partition profile, or the logical partition does not activate.

9. Click Next.

The Virtual Ethernet page appears.

10. Specify the maximum virtual adapters for the partition. (Optional) You can add new virtual ethernet adapters by clicking + (Add) and selecting an Adapter ID, Virtual LAN ID, Access External Network, Trunk Priority, IEEE 802.1 Q Compatibility, additional Virtual LAN IDs, and whether the Ethernet adapter is Required.

11. Click Next.

The Virtual Disks page appears.

12. Specify the virtual SCSI devices or physical fibre channel port 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 select 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, it is not associated with a virtual SCSI server adapter yet. When the job to create the partition takes place, the virtual SCSI server adapter is created and assigned to the device.

Note: If a physical fibre channel port that supports NPIV is selected, a virtual fibre channel server adapter and virtual fibre channel client adapter are created for the partition.

13. Click Next.

The Provision Storage wizard appears only when adding a new backing device. If not adding a new backing device, skip to Step 21.

14. Select a provisioning method. For more information about storage provisioning, see Storage Provisioning Manager for NetApp in the *Administration Guide*.
15. Select HMC server, Managed System.
16. Select a Virtual I/O Server. To add a new virtual Virtual I/O Server, click New to create a Virtual I/O Server. Select the Virtual I/O Server partition name, user name, and password, and click Save to validate the virtual I/O server and add the configuration.
17. Select the NetApp Data Fabric Manager, desired storage services, and any other advanced options.
18. Click Finish to confirm the request.

The NetApp iSCSI storage is attached to the Virtual I/O Server, and the new provisioned storage is selected.

The New SCSI Adapter dialog appears.

19. Select this device or another device, and click OK.

20. Click Next.

The NIM page appears. You can choose whether to provision AIX using NIM by toggling Deploy Operating System.

21. Specify the necessary Build Machine, System Attributes, Software Delivery Information, and Template, and click Next.

22. Verify the Summary page, and click Add Computer

The logical partition is created, and NIM provisioning starts.

If a physical fibre channel port is selected and NIM is enabled, a dialog displays WWPNs assigned to the partition. Further configuration is required to zone and create the storage LUN for the WWPNs. There is an option to continue or to cancel the NIM provisioning. You can choose to continue after the additional configuration has been completed. You can choose to cancel and perform a separate NIM provision operation at a later time once the configuration is complete.

IBM PowerVM Management

This section describes the IBM PowerVM management operations that you can perform from the Resources page.

The Resources page lets you view events and perform management operations on LPARs. Expand the IBM PowerVM group in the Explore pane to list the following objects:

- HMC/IVM servers
- PowerVM systems
- Logical partitions (LPARs)

View Resource Summary and Events

CA Server Automation displays the Summary in the right-hand pane. The Summary page provides resource properties at the following levels in the object hierarchy:

- PowerVM Server
- LPAR

Performance Chart pane displays the utilization with available metrics and options.

General Information pane includes the following properties:

- Names, item types, versions
- Quantity characteristics of CPU and memory
- Number of LPARs and shared processor pools
- Current mode of the resource

Overview pane displays information about:

- CPU state
- Memory state
- System state
- Collection engine state
- Health state
- Propagated health state

Usage pane displays:

- Overall CPU usage
- Overall memory usage

The Summary tab lets you view information associated with that object for example, total memory, operating system, number of CPUs, IP address, overall CPU and memory usage and events associated with the resource. Click the Configuration tab on the Usage panel to configure threshold limits.

Configuring CPU and Memory

You can configure memory shares allocated to a virtual machine to adjust its allocated resources. When you add resources, the appropriate amount of unassigned memory or CPU shares must be available for the operation to succeed. If values exist for the minimum and maximum allowed memory or CPU shares, any resource allocation change must stay within these limits. You can edit VM CPU and memory allocation using the Quick Start link on the Resources tab. You can also use create and schedule policy with specific VM resource allocation actions.

Important! For Dynamic LPAR operations, like adding or removing CPU and Memory, install AIX version 5.2 or 5.3 or 6.0 or higher on each LPAR system. Alternatively, run AIX resource control daemon IBM.DRM on the LPAR system.

Configure CPU

To configure VM CPU allocation

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Configuration, Configure Processor...

The Configure Processor Resource Allocation dialog appears.

3. Select *one* of the following Adjustment Types:

Dynamic Adjustment

Updates the running VM.

Profile Update

Updates the active profile. The VM must be restarted to pick-up the changes from the profile.

Dynamic Adjustment and Update Profile

Updates both the running VM and the active profile.

4. Edit the corresponding fields and click Ok.

A confirmation message appears.

Configure Memory

To configure VM Memory allocation

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Configuration, Configure Memory...

The Configure Memory Resource Allocation dialog appears.

3. Select one of the following Adjustment Types:

Dynamic Adjustment

Updates the running VM.

Profile Update

Updates the active profile. The VM must be restarted to pick-up the changes from the profile.

Dynamic Adjustment and Update Profile

Updates both the running VM and the active profile.

4. Edit the corresponding fields and click Ok.

A confirmation message appears.

Control Power Status for Logical Partitions

You can control the status of logical partition by performing one of the following operations:

- Activate
- Restart
- Shutdown
- Delete

You can perform any of these operations on one or multiple logical partitions simultaneously.

Follow these steps:

1. Select the managed machine on which you want to perform a status operation in the Explore pane.
2. Right-click the partition, select Management. You can also click Quick Start and click the related link of power control. Select *one* of the following:

Activate

Activates the selected logical partition that is currently powered off or suspended.

Restart

Shuts down the guest operating system and restarts it.

Shutdown

Shuts down the selected logical partition. You can only shut down a logical partition that is currently powered on.

Delete

Deletes the selected logical partition permanently. You can only delete a logical partition if it is shut down.

A confirmation dialog appears.

3. Click OK.

The status operation occurs, and a confirmation message appears. Refresh the interface to view the new logical partition status. An event appears confirming the result of the operation.

Activate Logical Partition

Activate a Logical Partition to commit resources to the partition and start the installed operating system. You can only activate a partition when it is not running.

To activate a Logical Partition

1. Right-click a partition on the Explore pane and select Management, Activate.

The Activate Logical Partition dialog appears.

2. Complete the following fields and click OK.

Profile

Specifies the partition profile used to activate the partition.

Keylock

Specifies the key lock position. Key Lock establishes the power-on and power-off modes allowed for the system. CA Server Automation supports the following valid keylock modes:

Do Not Override

The LPAR uses the keylock mode specified in the selected profile.

Normal

The LPAR starts up as normal. Use this option to perform most everyday tasks.

Manual

Consider security impact, when you set the key lock position to Manual.

Boot mode

Specifies the boot mode. Select a boot mode and select the Activate check box only if you want to use a boot mode that is different from the one specified in the selected profile. The system uses this boot mode to start the operating system on the logical partition unless you specify otherwise when activating the partition profile. CA Server Automation supports the following valid boot modes:

Do Not Override

The LPAR uses the boot mode specified in the selected profile.

Normal

The LPAR starts up as normal. Use this option to perform most everyday tasks.

Open Firmware

The LPAR boots to the open firmware prompt. This option is used by service personnel to obtain additional debug information.

3. Click the Events tab for the partition.

An event should appear confirming the result of the operation.

Delete Logical Partition

You can delete a partition from the managed system that is no longer required. When you delete a logical partition, all hardware resources are returned to the primary partition. You can only delete a partition that is powered off.

To delete a Logical Partition

1. Right-click a partition in the Explore pane and select Management, Delete.

A confirmation dialog appears.

2. Click OK.

A message appears confirming that the request was submitted.

3. Click the Summary tab for the partition.

An event should appear confirming the result of the operation. The deletion is unsuccessful if the partition is not powered off. If the deletion was successful, the partition disappears from the Explore pane after you refresh the interface.

Restart Logical Partition

You can restart a partition that is already running. Restarting a partition shuts it down and starts the operating system again.

Note: A Logical Partition must be in the Running or Open Firmware state to restart.

To restart a Logical Partition

1. Right-click a partition on the Explore pane and select Management, Restart.
The Restart Logical Partition dialog appears.
2. Select one of the following restart types using the Type drop-down list and click OK:

Immediate

Shuts down the logical partition immediately. The HMC/IVM ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

Operating System Shutdown

Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

Operating System Shutdown Immediate

Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

3. Click the Summary tab for the partition.
An event should appear confirming the result of the operation.

Shut Down Logical Partition

Shutting down a partition shuts down the operating system. A partition must be in the Running or Open Firmware state to shut down.

To shut down a Logical Partition

1. Right-click a partition on the Explore pane and select Management, Shut Down.
The Shut Down Logical Partition page appears.
2. Select one of the following shutdown types using the Type drop-down list and click OK:

Immediate

Shuts down the logical partition immediately. The HMC/IVM ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

Operating System Shutdown

Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

Operating System Shutdown Immediate

Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

3. Click the Summary tab for the partition.
An event should appear confirming the result of the operation.

Citrix XenServer

Citrix XenServer is a virtualization platform that offers near bare metal virtualization performance for virtualized server and client operating systems. XenServer uses the Xen hypervisor to virtualize each server on which it is installed, enabling each server to host multiple virtual machines (VMs) simultaneously with guaranteed performance. XenServer provides its own operating system to administer the physical and virtual resources of a XenServer host, and therefore, it does not require a specific operating system. XenServer supports Linux and Windows guest operating systems.

XenServer resources can be managed at three levels:

Host Management

A *XenServer host* object represents a physical host on which XenServer and its VMs run. A XenServer host can be a stand-alone host or associated with a XenServer pool. You can monitor virtual and physical resources available on a XenServer host, manage storage repositories containing virtual disk images, manage tasks, or run the XenServer host in maintenance mode.

Resource Pool Management

A *resource pool* is a connected group of up to 16 XenServer hosts. Combined with shared storage and dynamically controlled memory, CPU, and networking resources, the XenServer hosts in a resource pool provide an operating environment on which VMs run. You can manage the membership or role of the XenServer hosts in a pool, and can let XenServer monitor the health of the pool members for high availability. If necessary, VMs are live migrated between pool hosts to avoid downtime.

Virtual Machine Management


On the VM management level, you can perform the following tasks:

- Control VMs (Discover, Start, Suspend, Shutdown, Delete From Disk)
- Manage VMs (clone)

How to Configure XenServer Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the Management Components



Follow these steps:

1. [Review Requirements](#) (see page 356).
2. Review [Interactions Between Citrix XenServer Management Components](#) (see page 392).
3. [Add a Citrix XenServer Connection to the Manager](#) (see page 393).
4. Troubleshoot if [Server Connection to the Manager Failed](#) (see page 393).
5. [Add the Discovered Citrix XenServer AIM Instance](#) (see page 395).
6. [Troubleshoot the AIM Instance Connection](#) (see page 365).
7. [Verify the Citrix XenServer Group in the Resources Tree](#) (see page 399).

Review Requirements

Review the following requirements before you start configuring the management components of CA Server Automation:


- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Server Automation, SystemEDGE, and your environment.
- You can access a CA Server Automation manager installation that includes:
 - Platform Management Module (PMM)
 - Application Insight Module (AIM)
 - Monitoring Agent (SystemEDGE)
- You can access the CA Server Automation user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Server Automation manager discovered remote AIM Servers that you want to use.

Interactions Between Citrix XenServer Management Components

The Citrix XenServer AIM is implemented as a multi-instance, remote AIM. CA Citrix XenServer AIM is able to monitor multiple standalone Citrix XenServers and/or Citrix XenServer resource pools remotely. The Citrix XenServer AIM is implemented as x86 and x64 module.

The management API for Citrix XenServer is based on XML RPC. For a Citrix XenServer resource pool, all XML RPC communication is taking place between the AIM, PMM, and the pool master only.

Interaction Between XenServer Management Components



Add a Citrix XenServer Connection to the Manager

You can add a Citrix XenServer connection using the Administration tab of the CA 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 Citrix XenServer from the Provisioning section in the left pane.
3. Click  (Add) on the Registered Citrix XenServers pane toolbar.

The Add Citrix XenServer dialog appears.

4. Enter the required connection data (server name, username, password, resource pool UUID), specify the preferred AIM, and enable Managed Status (checkbox).

Important! Verify that you add the pool master to the Registered Citrix XenServers.

5. Click OK.

If the network connection has been established successfully, the Server is added to the top right pane with a green status icon. CA Server Automation discovers the Citrix XenServer system automatically.

If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Server Automation adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Server Connection to the Manager Failed (Citrix XenServer)

Symptom:

After I have added a Server connection under Administration, Configuration, the validation of the connection to the Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Server connection data is still valid. If necessary, update the connection data.
- Verify, if the Server system is running and accessible.
- Verify, if the Management Service on the Server system is running properly.

To update the Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, enable Managed Status, and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the Server cannot be established, continue with the next procedure.

To verify if the Server system is running and accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <Server Name>
ping <IP Address of Server>
```

2. Verify the output of the commands to find out whether the Server has a valid DNS entry and IP address.

If the Server is not in the DNS, add the Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress <Server Name>
```

Enter the correct IP address and Server name. For example:

```
192.168.50.50 myServer
```

4. Click  (Validate) in the upper-right corner.

If the Server credentials and connection data are correct and you can ping the Server, the connection can still fail. In this case, it is possible that the Server causes the problem. If the connection to the Server cannot be established, continue with the next procedure.

To verify, if the Management Service on the Server system is running properly:

1. Contact the Administrator to access the Server system.
2. Log in to the Server system and execute `xsconsole` command
The Service control console launches.
3. Verify the status of the service and resolve any reported issues
4. Change to the CA Server Automation user interface, Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the Server connection.

If the connection to the Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

Work with the administrator or support to fix the Server connection problem.

Add the Discovered Citrix XenServer AIM Instance

After adding a Citrix XenServer connection to the CA Server Automation manager, add an AIM instance to manage the Citrix XenServer.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.
The Configuration page appears.
2. Select Citrix XenServer from the Provisioning section in the left pane.

3. Click  (Add) on the Discovered Citrix XenServer AIM Instances pane toolbar.

The Add Citrix XenServer AIM dialog appears.

4. Select the Citrix XenServer AIM Server from the drop-down list.

The list of discovered XenServer AIM Servers appears. If you have installed the XenServer AIM on the local system, the name of the local system appears in the list too.

5. Select the Citrix XenServer from the drop-down list.

CA Server Automation populates the XenServer drop-down list with the XenServers listed in the Registered Citrix XenServers pane. You can only manage those XenServers for which your CA Server Automation manager has a valid connection established.

Note: If the AIM resides on a remote system, CA Server Automation must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

A new AIM instance for the selected Server is added. If the instance is not in an error or in a stopped state, CA Server Automation starts to discover the associated environment. When the discovery process is complete, you can start managing your Citrix XenServer environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:



Discovery in progress



No polling



Error



Warning



Disabled



Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the amount of managed objects that are related to virtual and physical resources in your environment. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Server Automation adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Server Automation manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the AIM:

- Verify if the AIM Server is accessible.
- Verify if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
ping servername
```

2. Verify the output of the commands to find out whether the AIM server has a valid DNS entry and IP address.

If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress servername
```

Enter the correct IP address and AIM server name. For example:

```
192.168.50.51 myAIM
```

4. Click  (Validate) in the upper-right corner of the AIM Server pane.

If the error status remains unchanged, continue with the next procedure.

To verify if SystemEDGE is running:

1. Log in to the AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.

The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.

2. Start or restart SystemEDGE.

Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.

3. Change to the CA Server Automation user interface, AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the AIM Server connection.

If the error status remains unchanged, verify the validity of the data you gathered according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Server Automation discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Server Automation manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Server Automation manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Citrix XenServer Group in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:


1. Click Resources, and open the Explore pane.
2. Expand Citrix XenServer group.
The managed Citrix Resource Pools appear.
3. Expand the Resource Pool entry.
The managed Citrix XenServers appear.

CA Server Automation is now ready to manage the Citrix XenServer environment that was configured.

How to Prepare Windows Templates for XenServer Provisioning

CA Server Automation supports customized provisioning of new virtual machines (VM) running Windows 2003 R2 Server (32 bit and 64 bit), Windows 2008 (32 bit and 64 bit) or Windows 2008 R2 Server (64 bit). Customization options include a number of settings. For example, changing the built-in Administrator account password, computer name, and the network configuration.

The following diagram illustrates how a system administrator prepares Windows templates for XenServer provisioning.



Follow these steps:

1. [Prepare an image.](#) (see page 401)
2. [Install CA provisioning helper.](#) (see page 401)
3. (Valid on Windows 2003 R2) [Install the Sysprep tool.](#) (see page 402)
4. Depending on your operating system select *one* of the following actions:
 - [Run Sysprep tool on Windows 2003 R2.](#) (see page 402)
 - [Run Sysprep tool on Windows 2008 or on Windows 2008 R2.](#) (see page 402)
5. Convert the VM to a template.

Prepare an Image

When creating a template containing the Windows operating system, prepare the image by following this procedure. Follow the steps to enable <CAVMP> provisioning operations to customize the template. The specific steps differ based on the Windows version.

Follow these steps:

1. Install the Windows operating system on a new virtual machine from scratch.
2. Install the system tools inside the virtual machine.
3. Apply any customizations like user accounts, policy, applications, hotfixes, and so on, that you would like to apply on the new virtual machines.
4. (Valid on Windows 2003) Blank out the built-in Administrator account password.

Note: If the Administrator password is not empty, SysPrep is unable to set a new password during provisioning and the existing password remains.

Prerequisites for XenServer Environments

Ensure that the following prerequisites are met:

- Each XenServer in the resource pool must have SSH or SFTP access enabled.

Install CA provisioning helper

CA provisioning helper enables <CAVMP> to change the virtual machine settings externally.

Follow these steps:

1. Find this utility at <InstallationRoot>\Utilities\Sysprep\CAProvisioningHelper.exe
2. Transfer this executable file to any location on a hard drive of the VM being prepared.
3. Execute CA provisioning helper once from the command line.

The Sysprep Tool

The Microsoft provided Sysprep tools to generalize, freeze and shut down the readily configured Windows installation. The following sections describe how to use the Sysprep tool for Windows 2003 R2 and Windows 2008 R2 in detail.

Install and Run the Sysprep Tool on Windows 2003 R2

On Windows 2003 the Sysprep tools are not installed by default, but can be found on the Windows installation CD-ROM.

Install the Sysprep Tool

Install the Sysprep tool from the Windows installation CD-ROM.

Run the Sysprep Tool on Windows 2003 R2

After you configure the Sysprep tool installation, run the Sysprep tool.

Follow these steps:

1. Locate and open the following CAB file:
`\SUPPORT\TOOLS\DEPLOY.CAB`
2. Select all files contained in the CAB file and copy them to the following location:
`%SystemDrive%\Sysprep` (normally `C:\Sysprep`).

Note: Do not change the directory name.

3. Change to the Sysprep directory and run:
`sysprep -quiet -reseal -mini -shutdown`

Run the Sysprep Tool on Windows 2008 R2

All files to perform the SysPrep process are installed as part of the regular Windows installation process. After you configure the Windows installation, perform the following steps:

1. Generate a valid XML response file using the Windows Automated Installation Kit (WAIK) for Windows Server 2008 R2. WAIK is available from the Microsoft Web site.

Note: The way provisioning requires a dummy unattended response file, or it cannot shutdown. The content of the response file is irrelevant, since the provisioning process replaces it, but the file must follow the SysPrep-specific XML schema.

2. Name the generated XML file “`sysprep.xml`” and place it into the Sysprep directory:
`%SystemRoot%\system32\sysprep`
3. Run the following command:
`sysprep /generalize /oobe /shutdown /unattend:sysprep.xml`

Convert the VM to a Template (XenServer)

The template allows you to create any number of customized Windows virtual machines.

Follow these steps:

1. Shut down the Windows virtual machine.
2. Use XenCenter to convert the prepared image to a template.

The template appears in CA Server Automation and can be used for customized provisioning.

Manage VM Status (XenServer)

You can control the status of virtual machines by performing one of the following operations:

- Discover
 - Server
 - Network
- Start
- Suspend
- Shutdown
- Delete From Disk

To control VM status:

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Right-click a VM, select Management and one of the following options:

Discover

Discovers a server or network.

Start

Starts a VM on the specified XenServer host.

Suspend

Suspends a running VM on the specified XenServer host and saves its current state. All activity is suspended until you resume the VM.

Shutdown

Shuts down a running VM on the specified XenServer host.

Delete From Disk

Deletes a VM from the Disk.

A corresponding wizard appears.

3. Fill in the required information and proceed to the next step.
4. Submit.

The status operation occurs, and a confirmation message appears. Refresh the interface to view the new VM status. An event appears confirming the result of the operation.

Provision a Citrix XenServer Virtual Machine

You can provision virtual machines by performing the following procedure. Ensure that you prepare a Windows template for VM provisioning.

Follow these steps:

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Right-click the Citrix XenServer group, select Provisioning, Provision Citrix XenServer Virtual Machine.
A provisioning wizard appears.
3. Fill in the required information:

VM Name

Defines the new VM name.

Template

Specifies the Windows provisioning template.

Administrator Password

Defines the administrator password for the new VM.

Product Activation Key

Defines the Windows 2003 Activation Key.

Full Name

Defines the full VM name.

4. (Optional) Fill in the additional information (Workgroup, Memory, CPUs, VM Host, Organization). If you want to use a static IP address, disable the DHCP and provide the IP address, mask, and default gateway.

Note: The Memory and CPUs settings depend on the Windows provisioning template used.

5. Submit.
The confirmation message appears.
6. Refresh the Jobs panel to view the progress.
An event appears confirming the result of the operation.

Solaris Zones

A Solaris Zone defines a virtualized operating system that provides an isolated, secure environment in which to run applications. This environment allows allocation of resources among applications and services, and ensures that processes do not affect other zones. Solaris manages each zone as one entity. A *container* is a zone that also uses the resource management of the operating system. The Solaris Zones PMM provides health monitoring, management, and provisioning of Solaris Zones environments.

Solaris Zones Container resources can be managed at three levels:

Solaris Zones Zone Management

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

Solaris Zones Project Management

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


Solaris Zones Resource Pool Management

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

How to Configure the Solaris Zones Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the Management Components



The Solaris Zone PMM provides health monitoring, management, and provisioning of Solaris Zone environments.

Follow these steps:

1. [Review Requirements](#) (see page 356) and [Requirements for Solaris Zones Management](#) (see page 409).
2. Review [Interaction Between Solaris Zones Management Components](#) (see page 410).
3. [Add a Solaris Zone Connection to the Manager](#) (see page 411).
4. Troubleshoot if [Server Connection to the Manager Failed](#) (see page 362).
5. [Add the Zones AIM Servers](#) (see page 414).
6. [Troubleshoot the AIM Instance Connection](#) (see page 365).
7. [Verify the Solaris Zones Group in the Resources Tree](#) (see page 418).

Review Requirements

Review the following requirements before you start configuring the management components of CA Server Automation:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Server Automation, SystemEDGE, and your environment.
- You can access a CA Server Automation manager installation that includes:
 - Platform Management Module (PMM)
 - Application Insight Module (AIM)
 - Monitoring Agent (SystemEDGE)
- You can access the CA Server Automation user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Server Automation manager discovered remote AIM Servers that you want to use.

More information:

[Requirements for Solaris Zones Management](#) (see page 409)

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 privileges to execute the following commands:
 - zlogin
 - zoneadm
 - zonecfg
- From the global zone, the user must have the permission to log in to individual Solaris Zones with zlogin and run the following commands:
 - uname -a
 - sar
 - prstat
 - netstat


Using the user interface or the NodeCfgUtil.exe utility on the Managed Node where the Solaris Zones AIM resides, add this user name and the corresponding password to CA Server Automation during configuration .

Create a resource pool from Explore, Management, Create Resource Pool to allocate resources to zones, projects, and applications on a Solaris Zones server. Assign it to a zone during zone creation.

Interaction Between Solaris Zones Management Components

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

Interaction Between Solaris Zones Management Components



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

Add a Solaris Zones Connection to the Manager

You can add a Solaris Zones connection using the Administration tab of the CA 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 Solaris Zone from the Provisioning section in the left pane.
3. Click  (Add) on the Solaris Zone Servers pane toolbar.

The Add Solaris Zone Server dialog appears.

4. Enter the required connection data (server name, user, password, port), specify the preferred AIM, enable Managed Status (checkbox).
5. Click OK.

If the network connection is established successfully, the Server is added to the top right pane with a green status icon.

Note: If the connection fails, the Validation Failed dialog appears. Click Yes, CA Server Automation adds the Server to the list with a red status icon. If you click No, nothing is added.

Server Connection to the Manager Failed

Symptom:

After I have added a Server connection under Administration, Configuration, the validation of the connection to the Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Server connection data is still valid. If necessary, update the connection data.
- Verify, if the Server system is running and accessible.
- Verify, if the Management Service on the Server system is running properly.

To update the Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, enable Managed Status, and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the Server cannot be established, continue with the next procedure.

To verify if the Server system is running and accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <Server Name>
ping <IP Address of Server>
```

2. Verify the output of the commands to find out whether the Server has a valid DNS entry and IP address.

If the Server is not in the DNS, add the Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress <Server Name>
```

Enter the correct IP address and Server name. For example:

```
192.168.50.50 myServer
```

4. Click  (Validate) in the upper-right corner.

If the Server credentials and connection data are correct and you can ping the Server, the connection can still fail. In this case, it is possible that the Server causes the problem. If the connection to the Server cannot be established, continue with the next procedure.

To verify, if the Management Service on the Server system is running properly:

1. Contact the Administrator to access the Server system.
2. Log in to the Server system and open Administrative Tools, Services from the Start menu.

The Services window opens.

3. Select the service and start or restart the service.
4. Change to the CA Server Automation user interface, Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the Server connection.

If the connection to the Server fails, verify the validity of the data you gathered according to the requirements for this scenario.

Work with the administrator or support to fix the Server connection problem.

Add the Zones AIM Servers

After adding a Solaris Zone connection to the CA Server Automation manager, add an AIM instance to manage the Solaris Zone environment.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select Solaris Zone from the Provisioning section in the left pane.
3. Click  (Add) on the Solaris Zones AIM Instances pane toolbar.

The Add Solaris Zones AIM Instance appears.

4. Select the Zones AIM Server from the drop-down list.

The list of Zones AIM Servers appears.

5. Select the ZONE Server from the drop-down list.

CA Server Automation populates the Zone Server drop-down list with the Zone Servers listed in the Registered Solaris Zones pane. You can only manage those Zone Servers for which your CA Server Automation manager has a valid connection established.

Note: If the AIM resides on a remote system, CA Server Automation must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

A new AIM instance for the selected Server is added.

Note: If the instance is not in an error or in a stopped state, CA Server Automation starts to discover the associated environment. When the Discovery process is complete, you can start managing the Solaris Zone environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

-  Discovery in progress
-  No polling
-  Error
-  Warning
-  Disabled
-  Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the amount of managed objects that are related to virtual and physical resources in your environment. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Server Automation adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows

 (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Server Automation manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the AIM:

- Verify if the AIM Server is accessible.
- Verify if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
ping servername
```

2. Verify the output of the commands to find out whether the AIM server has a valid DNS entry and IP address.

If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

ipaddress servername

Enter the correct IP address and AIM server name. For example:

192.168.50.51 myAIM

4. Click  (Validate) in the upper-right corner of the AIM Server pane.
If the error status remains unchanged, continue with the next procedure.

To verify if SystemEDGE is running:

1. Log in to the AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.
The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.
2. Start or restart SystemEDGE.
Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.
3. Change to the CA Server Automation user interface, AIM Server pane on the manager system and click  (Validate) in the upper-right corner.
CA Server Automation validates the AIM Server connection.
If the error status remains unchanged, verify the validity of the data you gathered according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Server Automation discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Server Automation manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Server Automation manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Solaris Zones Group in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand Solaris Zone group.

The managed Solaris Zone resources appear.

CA Server Automation is now ready to manage the configured Solaris Zone environment.

Solaris Zones Management

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

This section describes the management operations that you can perform on Solaris Zones resources from the Resources page. The Resources page lets you view basic information and details about the following objects:

- Solaris Zones servers
- Solaris Zones

Click Resources, open the Explore pane, and select one of the resources; then click Summary for the resource. The Summary page lets you view information associated with the resource (for example, zones, resource pools, and disks on a Solaris Zones server or network interfaces and projects on a zone) and events associated with the resource.

Note: If you select alert as Normal using the Configuration button in the Usage panel, the zone is shown in the normal state (in green) even if the CPU or memory is in the critical or warning state. Similarly, if you select alert as warning, the zone is always shows in the warning state.

The components tree displays only those resource pools that are used by a zone. Inactive resource pools are not listed in this panel.

The Details page lets you view other detailed resource information, such as system properties, software, hardware, performance, and so on.

Other pages may be available to perform resource management tasks. Right-click menus on the Explore pane also let you perform management and policy tasks.

More Information

[Create Resource Pool](#) (see page 421)

[Control Zone Status](#) (see page 422)

[Clone a Zone](#) (see page 423)

[Delete a Zone](#) (see page 424)

[Available Solaris Zones Actions](#) (see page 424)

Add a Solaris Zone

Solaris Zones servers use Zones to run applications in isolated environments to make it appear as if they are running on physically separate systems. Each Zone on a server takes its resources from a resource pool and includes virtual network interfaces, file systems, memory, and other dedicated units. When you create a Zone, you must supply all of this information. The Zone installs automatically after creation.

To add a Solaris Zone

1. Select the Resources tab, right-click the Zone Host in the Explore pane, and select Provisioning, Provision Zone.

The Solaris Zone Provisioning wizard appears.

2. Complete the following fields on the Zone Identity and Type page and click Next:

Host

Defines the host on which to create the Zone.

Name

Defines the name of the Zone.

Description

(Optional) Defines a description of the Zone.

Type

Defines whether the Zone is Native, Whole Root, or Branded. A Branded Zone is based on an existing Zone template.

Template Name

(Optional) Defines the template from which to create the Zone when you set Type to Branded.

Install Archive Path

Defines the directory path of the installation archive on the Zone. This field is only required if you set Type to Branded.

The CPU, Memory, and Additional page appears.

3. Complete the following fields and click Finish:

Type

Defines the scheduler type. Select FSS to use the Fair Share Scheduling class to control CPU allocation based on the number of CPU shares assigned to tasks.

Capacity

Defines the amount of physical memory capacity to allocate to the Zone, in megabytes.

Swap Memory

Defines the amount of swap memory to allocate to the Zone, in megabytes. The swap memory must be at least 50 MB.

Lock Memory

Defines the amount of lock memory to allocate to the Zone, in megabytes. The lock memory must be less than the physical memory.

Zone Path

Defines the root directory path of the Zone.

NIC Type

(Optional) Defines the NIC type. Select a type from the drop-down list. If you do not select a NIC, the Zone is not assigned a NIC card or IP address.

IP Address

(Optional) Defines the IP address of the Zone.

Resource Pool

Defines the resource pool to use with the Zone. Select a pool from the drop-down list. If you want to use a new resource pool with the Zone, create the pool first. If you do not select a pool, the default is used.

Auto Reboot

Defines whether to reboot the Zone automatically when the global Zone is rebooted.

Create Resource Pool

You can create a resource pool for use in allocating resources to zones, projects, and applications on a Solaris Zones server. After you create a resource pool, you can assign it to a zone during zone creation.

To create a resource pool

1. Right-click a Solaris Zones server on the Explore pane and select Management, Create Resource Pool.

The Create Resource Pool dialog appears.

2. Complete the following fields and click OK:

Name

Identifies the resource pool name.

Min CPU Shares

Identifies the minimum number of CPU shares that the pool must have at any time.

Max CPU Shares

Identifies the maximum number of CPU shares that the pool can have.

Processor Set Name

Identifies the pool's processor set name.

Scheduler Type

Identifies the type of scheduling to use when allocating resources. Select FSS to use Fair Share Scheduling to allocate resources based on workload importance (the number of CPU shares specified for a project or task).

The pool is created, and a confirmation message appears.

3. Click the Summary tab for the Zones server on which you created the pool, and select Resource Pools in the Show drop-down list to verify that the pool was created.

Control Zone Status

You can perform stop, reboot, start, and uninstall operations to control the status of a zone. You cannot perform these operations for global zones or when a zone is in the installed state.

To control zone status

1. Right-click a zone on the Explore pane, and select Management and one of the following options:

Start

Starts the zone, putting it into a running state. You can only start a zone that is in the installed state.

Stop

Halts the zone by resetting it to the installed state. Halting the zone stops all processes, removes network interfaces, and performs other operations to remove the zone's existing application environment and virtual platform. After halting a zone, you must start the zone to re-initiate the environment. You can only halt a zone that is currently running.

Reboot

Halts the zone and boots it again. You can only reboot a zone that is currently running. When you reboot a zone, the server gives it a new zone ID.

Delete

Deletes a zone. For more information, see the section [Delete a Zone](#).

Install

Installs a native or branded zone which enters the configured state when the installation completes. Installing a zone opens a dialog that asks you for the archive path of a branded zone. If you install a native zone, leave this field empty. In case of a branded zone, provide the archive path.

Note: You get an error message when you try to install a branded zone with no archive path parameter entered or a native zone with archive path parameter entered.

Uninstall

Uninstalls all the files under the zone's root file system. You can only uninstall a zone that is not currently running (installed state). You should uninstall a zone before you delete it.

Clone

Clones a zone. For more information, see the section [Clone a Zone](#)

A confirmation dialog appears.

2. Click OK.

A message appears confirming that the request was submitted.

3. Click the Summary tab for the zone host.

An event should appear confirming the result of the operation.

Note: The zone status shows incomplete, if the current operation is in progress and has not yet completed.

More Information

[Delete a Zone](#) (see page 424)

[Clone a Zone](#) (see page 423)

Clone a Zone

Cloning a zone lets you configure and install a new zone by copying the data from an existing zone. The zone that you are cloning must be halted for the Clone operation to be available. You cannot perform this operation for global zones or when a zone is in the configured or running state.

To clone a zone

1. Right-click a zone on the Explore pane and select Management, Clone.

The Cloning pane appears.

2. Complete the following fields in the Target pane and click Clone.

Name

Specifies the name of the zone to which you want to copy the cloned information.

Path

Specifies the file path of the zone to which you want to copy the cloned information.

A confirmation message appears.

3. Click the Summary tab for the zone host.

An event confirming the result of the operation appears in the dashboard. The cloned zone appears in the Explore pane under its containing host when the operation completes.

Delete a Zone

You can delete a non-global zone from a Solaris Zones server. The zone must be shut down before you delete it.

If the zone is in the installed state, this operation will uninstall and then delete the zone. If the zone is in any other state such as running, an error message displays.

To delete a zone

1. Right-click a zone on the Explore pane and select Management, Delete.

A confirmation dialog appears.

2. Click OK.

A message appears confirming the deletion.

3. Click the Summary tab for the zone host.

An event should appear confirming the result of the operation. The deleted zone should disappear from the Explore pane when the operation completes.

Available Solaris Zones Actions

The following action types are available for use with Solaris Zones:

- [Clone Zone Machine](#) (see page 578)
- [Delete Zone Machine](#) (see page 606)
- [Provision Zone Machine](#) (see page 633)

You can use these action types to create new actions that automate zone operations when assigned rule criteria are met. You can also schedule these actions to occur at specific times.

For more information about using actions and rules to create automation policy, see the chapter "Policy."

Microsoft Hyper-V Server

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

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

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

Hyper-V Server

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


Virtual Machine

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

How to Configure Hyper-V Management

The following diagram provides an overview of the required actions. The diagram includes troubleshooting strategies for connection problems.

How to Configure the Hyper-V Server Management Components



Follow these steps:

[Review Hyper-V Requirements](#) (see page 427)
[Apply Required Settings for Using Microsoft Hyper-V](#) (see page 428)
[Add a New Hyper-V Server Connection to the Manager](#) (see page 430)
[\(Optional\) Add the SCVMM Management Instance to the CA Server Automation Manager](#) (see page 433)
[Discover the Servers](#) (see page 437)
[Verify the Hyper-V Server Folder in the Resources Tree](#) (see page 438)

Review Hyper-V Requirements

Review the following requirements before you start configuring the Hyper-V management components of CA Server Automation:

- 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 Hyper-V Servers.
- You can access a CA Server Automation manager installation that includes the Hyper-V Platform Management Module (PMM), Hyper-V Application Insight Module (AIM), and Monitoring Agent (CA SystemEDGE).
- You can access the CA Server Automation user interface.
- Verify that the Hyper-V AIM is installed at the Hyper-V Server.
- You have valid credentials available (user name and password) to access the Hyper-V Server that you want to manage.
- You have verified that the Hyper-V Server runs properly.
- Verify that the SNMP settings on the CA Server Automation manager and the Hyper-V Server are consistent. Read and write community strings and SNMP port number must be identical.
- You have verified that the CA Server Automation manager has discovered the Hyper-V Server that you want to use.

More information:

[Apply Required Settings for Using Microsoft Hyper-V](#) (see page 428)
[Interactions Between Hyper-V Server Management Components](#) (see page 429)
[Add a New Hyper-V Server Connection to the Manager](#) (see page 430)
[\(Optional\) Add the SCVMM Management Instance to the CA Server Automation Manager](#) (see page 433)
[Discover the Servers](#) (see page 437)
[Verify the Hyper-V Server Folder in the Resources Tree](#) (see page 438)

Apply Required Settings for Using Microsoft Hyper-V

Verify prerequisites and apply the following settings for Microsoft Hyper-V management:

To apply required settings for using Microsoft Hyper-V


1. Verify that SystemEDGE, Advanced Encryption, and the Hyper-V AIM are installed on the Hyper-V Server. You can only assign one AIM for each managed Hyper-V Server.
2. Disable the local User Access Control (UAC) on the Hyper-V Server.
3. Disable the network UAC by setting the following registry value:
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\LocalAccountTokenFilterPolicy,1 (REG_DWORD)
4. Enable remote WMI firewall exception by running the following command from the command prompt:


```
netsh advfirewall firewall set rule group="Windows Management Instrumentation (WMI)" new enable=yes
```
5. Verify that the user configured in the management components for server access is a member of the "Distributed COM Users" group.

Interactions Between Hyper-V Server Management Components

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

Interaction Between Hyper-V Server Management Components



You can configure Hyper-V management by adding connection information. Use *one* of the following methods:

- Administration tab of the user interface
- NodeCfgUtil.exe utility on the AIM Server

More information:

[Add a New Hyper-V Server Connection to the Manager](#) (see page 430)

[\(Optional\) Add the SCVMM Management Instance to the CA Server Automation Manager](#) (see page 433)

[Discover the Servers](#) (see page 437)

[Verify the Hyper-V Server Folder in the Resources Tree](#) (see page 438)

Add a New Hyper-V Server Connection to the Manager

You can add a Hyper-V 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 Hyper-V Server from the Provisioning section in the left pane.

The right pane refreshes and displays the managed Hyper-V Servers.

3. Click  (Add) on the Hyper-V Servers pane toolbar.

The New Hyper-V Server dialog appears.

4. Enter the required connection data (server name, user, password) and click OK.

If the network connection has been established successfully, the Hyper-V Server is added to the top right Hyper-V Servers pane with a green status icon. CA Server Automation discovers the Hyper-V Server automatically.

If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Server Automation adds the Hyper-V Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added. For troubleshooting the connection, see [Troubleshoot the Hyper-V Server Connection](#) (see page 431).

More information:

[\(Optional\) Add the SCVMM Management Instance to the CA Server Automation Manager](#) (see page 433)

[Discover the Servers](#) (see page 437)

[Verify the Hyper-V Server Folder in the Resources Tree](#) (see page 438)

[Hyper-V Server Connection Failed](#) (see page 431)

Hyper-V Server Connection Failed

Symptom:

After I have added a new Hyper-V Server connection under Administration, Configuration, the validation of the connection to the Hyper-V Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Hyper-V Server connection data (server name, user, password) is still valid. If necessary, update the connection data.
- Verify, if the Hyper-V Server system is running and accessible.

To update the Hyper-V Server connection data

1. Click  (Add) or  (Edit) that is associated with the failed connection.

The New or Edit Hyper-V Server dialog appears.

2. Add the valid server name, user, and password. Enable Managed Status and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the Hyper-V Server cannot be established, continue with the next procedure.

To verify, if the Hyper-V Server system is running and accessible

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <Hyper-V Server Name>
ping <IP Address of Hyper-V Server>
```

2. Verify the output of the commands to find out whether the Hyper-V Server has a valid DNS entry and IP address.

If the Hyper-V Server is not in the DNS, add the Hyper-V Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Hyper-V Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress <Hyper-V Server Name>
```

Enter the correct IP address and Hyper-V Server name. For example:

```
192.168.50.50 myHyper-V
```

4. Click  (Validate) in the upper-right corner.


If the connection to the Hyper-V Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

Work with the Hyper-V administrator or VMware support to fix the Hyper-V Server connection problem.

(Optional) Add the SCVMM Management Instance to the CA Server Automation Manager

The following diagram provides an overview about the required actions. The diagram includes corresponding troubleshooting strategies in case of connection problems.

Add a new SCVMM Server Connection to the Manager



Follow these steps:

[Apply Required Settings for Using Microsoft SCVMM](#) (see page 434)

[Add a New SCVMM Server Connection to the Manager](#) (see page 435)

[SCVMM Server Connection Failed](#) (see page 436)

Apply Required Settings for Using Microsoft SCVMM

CA Server Automation optionally integrates with Microsoft System Center Virtual Machine Manager (SCVMM) for Hyper-V provisioning. SCVMM is not required for Hyper-V monitoring and management. Instead of SCVMM, you can also use local templates (Hyper-V Server bound) for VM provisioning.

When using the optional SCVMM integration, verify the following prerequisites and apply the required settings:

To apply required settings for using Microsoft SCVMM

1. Verify that the SCVMM Server, all potential Hyper-V target hosts for VM provisioning, and the CA Server Automation manager running the Hyper-V PMM are members of an Active Directory domain.
2. Verify that Hyper-V target hosts for VM provisioning are configured in CA Server Automation and SCVMM. CA Server Automation does not perform a SCVMM discovery.
3. Verify that SCVMM has Windows Remote Management (WinRM) configured.
4. Run the following command on the command line of the SCVMM server to configure WinRM:

```
winrm quickconfig
```

5. Allow unencrypted HTTP or enable HTTPS in addition to the basic WinRM configuration on the SCVMM server.

Run the following command to allow unencrypted HTTP traffic:

```
winrm set winrm/config/service @{AllowUnencrypted="true"}
```

To enable HTTPS, obtain an SSL certificate for your SCVMM server, install it, and run the following command:

```
winrm quickconfig -transport:https
```

Depending on the expected utilization of the SCVMM server in your environment, VM provisioning can fail due to inadequate parameter settings for the Quota Management for Remote Shells on the SCVMM server. Affected parameters are:

MaxShellsPerUser

Specifies the maximum number of shells per user.

Default: 5

MaxConcurrentUsers

Specifies the maximum number of concurrent users who can open shells.

Default: 5

If you expect more than one provisioning job at a time, you can increase the parameter settings on the SCVMM server as follows:

```
winrm set winrm/config/winrs @{MaxShellsPerUser="number"}  
winrm set winrm/config/winrs @{MaxConcurrentUsers="number"}
```

Example

```
winrm set winrm/config/winrs @{MaxShellsPerUser="30"}  
winrm set winrm/config/winrs @{MaxConcurrentUsers="10"}
```

See also the [Quota Management for Remote Shells](#) article from Microsoft.

Add a New SCVMM Server Connection to the Manager

You can add a SCVMM connection using the Administration tab of the CA 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 SCVMM Server from the Provisioning section in the left pane.

The right pane refreshes and displays the managed SCVMM Servers.

3. Click  (Add) on the SCVMM Servers pane toolbar.

The New SCVMM Server dialog appears.

4. Enter the required connection data (server name, user, password) and click OK.

If the network connection has been established successfully, the SCVMM Server is added to the top right vCloud Servers pane with a green status icon. CA Server Automation discovers the SCVMM Server automatically.

If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Server Automation adds the SCVMM 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 SCVMM Server Connection](#) (see page 436).

SCVMM Server Connection Failed

Symptom:

After I have added a new SCVMM Server connection under Administration, Configuration, the validation of the connection to the SCVMM Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used SCVMM Server connection data (server name, user, password) is still valid. If necessary, update the connection data.
- Verify, if the SCVMM Server system is running and accessible.

To update the SCVMM Server connection data

1. Click  (Add) or  (Edit) that is associated with the failed connection.
The New or Edit SCVMM Server dialog appears.
2. Add the valid server name, user, and password. Enable Managed Status and click OK.
The connection data is updated.
3. Click  (Validate) in the upper-right corner to validate the new settings.
If the connection to the SCVMM Server cannot be established, continue with the next procedure.

To verify, if the SCVMM Server system is running and accessible

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <SCVMM Server Name>
ping <IP Address of SCVMM Server>
```
2. Verify the output of the commands to find out whether the SCVMM Server has a valid DNS entry and IP address.
If the SCVMM Server is not in the DNS, add the SCVMM Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.
If the SCVMM Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

ipaddress <SCVMM Server Name>

Enter the correct IP address and SCVMM Server name. For example:

192.168.50.50 mySCVMM

4. Click  (Validate) in the upper-right corner.

If the connection to the SCVMM Server fails, verify whether the data you gathered according to the requirements for this scenario is still valid.

Work with the SCVMM administrator or Microsoft support to fix the SCVMM Server connection problem.

Discover the Servers

After adding a new Hyper-V Server connection and an optional SCVMM connection to the CA Server Automation manager, discover the Hyper-V Server and the SCVMM Server. CA Server Automation then discovers the entire Hyper-V environment with all its virtual components.

Verify, that the SNMP credentials on the CA Server Automation manager and on the Hyper-V and SCVMM Servers are consistent. If necessary, update the SNMP configurations accordingly.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Resources, Data Center.

The data Center page appears.

2. Click Discover System from Quick start in the right pane.

The right pane refreshes and displays the Discover System wizard.

3. Enter the required data and click Finish.

CA Server Automation discovers the system.

More information:

[Verify the Hyper-V Server Folder in the Resources Tree](#) (see page 438)

Verify the Hyper-V Server Folder in the Resources Tree

After a successful configuration and discovery, the new Hyper-v Server is listed in the Resources Explore pane under the Hyper-V Server folder.

Follow these steps:

1. Click Resources, Explore.
The resources tree appears.
2. Expand Hyper-V Server.
The managed Hyper-V Servers appear.
3. Expand the new Hyper-V Server entry.
The managed Hyper-V infrastructure appears.

CA Server Automation is now ready to manage the added Hyper-V environment with its virtual infrastructure.

Hyper-V Management

Hyper-V Server lets you manage your Hyper-V servers and virtual machines. The Hyper-V Server is the central location from which you can view all virtual resources and perform management operations such as start, stop, delete, and so on.

This section describes the management operations that you can perform on Hyper-V Server resources from the Resources page. The Resources page lets you view basic information and details about the following objects:

- Hyper-V Servers
- Virtual Machines

Click Resources, open the Explore pane, and select one of the Hyper-V resources; then click Summary for the resource.

The Summary page lets you view information associated with that object (for example, Hyper-V Server, or virtual machines on a Hyper-V server) and events associated with the resource.

The Details page lets you view other detailed resource information, such as system properties, software, hardware, performance, and so on.

Right-click menus on the Explore pane let you perform management and policy tasks.

More Information

[Manage VM Status \(Hyper-V\)](#) (see page 441)

[Delete a Virtual Machine](#) (see page 442)

[Rename a Virtual Machine](#) (see page 442)

[Create Action and Rules](#) (see page 443)

[Edit Startup and Shutdown Actions](#) (see page 444)

[Edit VM CPU and Memory Allocation](#) (see page 445)

[Hyper-V Management Actions](#) (see page 446)

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..

More Information

[Provision Machine: Microsoft Hyper-V](#) (see page 630)

Manage VM Status (Hyper-V)

You can control the status of Hyper-V Server virtual machines by performing one of the following VM operations:

- Start
- Stop
- Pause
- Restart
- Shut Down
- Save

You can perform any of these operations on multiple VMs simultaneously.

To control VM status

1. Select the virtual machine on which you want to perform a status operation in the Explore pane.
2. Right-click the VM, select Management. You can also click Quick Start and click the related link of power control. Select *one* of the following:

Start

Starts the virtual machine and boots the guest operating system. You can only power on a virtual machine that is currently powered off or suspended.

Stop

Powers off the virtual machine. You can only power off a virtual machine that is currently powered on or suspended.

Pause

Pauses the virtual machine and saves its current state. All activity is suspended until you resume the machine.

Restart

Shuts down the guest operating system and restarts it.

Shutdown

Shuts down the guest operating system. You can only shut down a virtual machine that is currently powered on.

Save

Saves the current status of the virtual machine. In other platforms, this option is similar to Suspend.

A confirmation dialog appears.

3. Click OK.

The status operation occurs, and a confirmation message appears. Refresh the interface to view the new VM status. An event should appear confirming the result of the operation.

Delete a Virtual Machine

When you delete a virtual machine from Hyper-V Server, the virtual machine is deleted from the virtual disk.

To delete a virtual machine

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Delete.
The Delete Hyper-V VM dialog appears with options to delete additional components.
3. Click OK.
A message appears confirming the request submission.
4. Click the Summary tab for the virtual machine.
An event should appear confirming the result of the operation. If successful, the virtual machine is deleted from the virtual disk, and the virtual machine disappears from the Explore pane after you refresh the interface.

Note: You can only delete a VM that is in the power off state, otherwise the delete link is disabled.

Rename a Virtual Machine

You can rename an existing virtual machine from Hyper-V Server.

To rename a virtual machine

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Management, Rename.
The Rename VM dialog appears.
3. Enter the New VM Name and Click OK.
The message appears confirming the request submission.
4. Click the Summary tab for the virtual machine.

Create Action and Rules

You can create actions and rules based on predetermined policies for different types of resources such as a virtual machine.

To create Actions and Rules for a virtual machine

1. Select Resources, Explore, Data Center.
2. Open the Policy tab and the Actions subtab.
3. Click + (Add) to create a new Action.
4. Select the appropriate items from the drop-down menus and follow the instructions in the user interface to complete the Action.
5. Select the Rules subtab and click + (Add) to create a new Rule.

The Rule/Template Identification and Evaluation page dialog appears.

A wizard guides you through the creation process. Assign an Action from the available Actions list to this rule.

For more information about Actions and Rules, see [Create an Action](#) (see page 567) and Create a Rule.

Edit Startup and Shutdown Actions

You can edit the actions to start up and shut down a virtual machine.

To edit Startup and Shutdown Actions

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Configuration, Startup and Shutdown Actions.
The Startup and Shutdown Actions dialog appears.
3. The Startup and Shutdown Actions dialog contains the following fields:

Start Action

Specifies the action to perform when the Hyper-V Server starts. Select one of the following from the drop-down list:

- Always
Always starts the VM when the Hyper-V Server starts.
- Auto
Automatically starts the VM when the Hyper-V Server starts, if the VM is shutdown in running mode.
- None
Does not start the VM when the Hyper-V Server starts.

Start Delay

Adjust the delay (in seconds) to start a VM after the Hyper-V Server starts.

Shutdown Action

Specifies the action to perform when the virtual machine shuts down. Select one of the following from the drop-down list:

- Off
Turns off the VM before Hyper-V Server shuts down.
- Save
Saves (Suspends) the VM before Hyper-V Server shuts down.
- Shutdown
Shuts down the VM before Hyper-V server shuts down.

Recovery Action

Specifies the action to regain the previous details of a virtual machine when the Hyper-V Server fails. Select one of the following from the drop-down list:

- None
Does not take a specific action when the Hyper-V Server starts after the server fails.
- Restart
Restarts the VM when Hyper-V Server starts after the VM server fails.
- Revert
Reverts the VM with the latest snapshots when the Hyper-V Server starts after the server fails.

4. Click Ok after you edit all the details. A confirmation message appears.

Edit VM CPU and Memory Allocation

You can edit the number of CPU and memory shares allocated to a virtual machine to adjust its allocated resources. When you add resources, the appropriate amount of unassigned memory or CPU shares must be available for the operation to succeed. If values exist for the minimum and maximum allowed memory or CPU shares, any resource allocation change must stay within these limits.

You can also create and schedule policy with specific VM resource allocation actions.

To edit VM CPU and memory allocation

1. Open the Explore pane.
Available groups, services, and systems appear.
2. Find and right-click a virtual machine on the Explore pane and select Configuration, Resource Allocation, CPU and Memory.
The CPU and Memory Resource Allocation dialog appears.
3. Adjust the number of CPUs, reserved CPU percentage, and CPU Limit percentage.
4. Adjust the memory shares allocated to the virtual machine and click Ok after you edit all the details.
A confirmation message appears.

Hyper-V Management Actions

The following action types are available for use with Hyper-V Server:

- [Delete Machine](#) (see page 605)
- [Change Machine State](#) (see page 577)
- [Configure Power](#) (see page 590)
- [Configure CPU/Memory](#) (see page 581)

You can use these action types to create new actions that automate the configuration of startup and shutdown options for Hyper-V Server, and other operations when assigned rule criteria are met. You can also schedule these actions to occur at specific times.

For more information about using actions and rules to create automation policy, see the "Policy" chapter.

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 447)

[Add a Cisco UCS Manager Server](#) (see page 448)

[Register a UCS AIM Server](#) (see page 449)

[How to Configure the Cisco UCS Management Components](#) (see page 450)


[Cisco UCS Management](#) (see page 461)

[Interaction Between Cisco UCS Management Components](#) (see page 447)

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 manager is HTTP or HTTPS.

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.
The Cisco UCS page appears.
4. Click + (Add) on the Cisco UCS 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.
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.

How to Configure the Cisco UCS Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the Management Components



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.

Follow these steps:

1. [Review Requirements](#) (see page 356) and Requirements for [Cisco UCS Server](#) (see page 452).
2. Review [Interaction between Cisco UCS Management Components](#) (see page 453).
3. [Add a Cisco UCS to the Manager](#) (see page 454).
4. Troubleshoot if [Server Connection to the Manager Failed](#) (see page 362).
5. [Add the UCS AIM Server](#) (see page 456).
6. [Troubleshoot the AIM Instance Connection](#) (see page 365).
7. [Verify the Cisco UCS in the Resources Tree](#) (see page 460).

Review Requirements

Review the following requirements before you start configuring the management components of CA Server Automation:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Server Automation, SystemEDGE, and your environment.
- You can access a CA Server Automation manager installation that includes:
 - Platform Management Module (PMM)
 - Application Insight Module (AIM)
 - Monitoring Agent (SystemEDGE)
- You can access the CA Server Automation user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.
- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Server Automation manager discovered remote AIM Servers that you want to use.

More information:

[Cisco UCS Server](#) (see page 452)

Cisco UCS Server

Verify the following conditions for Cisco UCS management:


- Launch the Cisco Java user interface to verify that the Cisco UCS Manager is running. The link to launch the Cisco Java user interface is `http://<UCS_Manager_name>` or `https://<UCS_Manager_name>`.

Interaction Between Cisco UCS Management Components

Cisco UCS integration requires the UCS AIM for SystemEDGE to provide SNMP get/set requests for retrieving UCS devices and statistic data and configuring devices. The UCS Platform Management Module (PMM) also queries UCS devices and statistic information and stores the data in the Management DB. Cisco provides an XML API for interaction with the Cisco UCS Manager.

The API allows CA Server Automation to gain access to the hardware, statistics, pools (UUID, MAC, WWPN, WWNN), and the UCS Manager service profiles information.

Interaction Between Cisco UCS Management Components



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

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

Add a Cisco UCS to the Manager

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

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select Cisco UCS from the Provisioning section in the left pane.
3. Click  (Add) on the Cisco UCS pane toolbar.

The Add Cisco UCS dialog appears.

4. Enter the required connection data (server name, user, password, port), specify the preferred AIM, enable Managed Status (checkbox).
5. Enter the required server identification information, and click OK.

If the network connection is established successfully, the Server is added to the top right pane with a green status icon.

Note:If the connection fails, the Validation Failed dialog appears. Click Yes, CA Server Automation adds the Server to the list with a red status icon. If you click No, nothing is added.

Server Connection to the Manager Failed

Symptom:

After I have added a Server connection under Administration, Configuration, the validation of the connection to the Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Server connection data is still valid. If necessary, update the connection data.
- Verify, if the Server system is running and accessible.
- Verify, if the Management Service on the Server system is running properly.

To update the Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, enable Managed Status, and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the Server cannot be established, continue with the next procedure.

To verify if the Server system is running and accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <Server Name>
ping <IP Address of Server>
```

2. Verify the output of the commands to find out whether the Server has a valid DNS entry and IP address.

If the Server is not in the DNS, add the Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

ipaddress <Server Name>

Enter the correct IP address and Server name. For example:

192.168.50.50 myServer

4. Click  (Validate) in the upper-right corner.

If the Server credentials and connection data are correct and you can ping the Server, the connection can still fail. In this case, it is possible that the Server causes the problem. If the connection to the Server cannot be established, continue with the next procedure.

To verify, if the Management Service on the Server system is running properly:

1. Contact the Administrator to access the Server system.
2. Log in to the Server system and open Administrative Tools, Services from the Start menu.

The Services window opens.

3. Select the service and start or restart the service.
4. Change to the CA Server Automation user interface, Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the Server connection.

If the connection to the Server fails, verify the validity of the data you gathered according to the requirements for this scenario.

Work with the administrator or support to fix the Server connection problem.

Register a UCS AIM Server

After adding a Cisco UCS component to the CA Server Automation manager, add the AIM instance using the Administration page of the user interface to manage the Cisco UCS environment.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select Cisco UCS from the Provisioning section in the left pane.

3. Click  (Add) on the Cisco UCSs AIM Instances pane toolbar.

The Add Cisco UCS AIM Instance appears.

4. Select the UCS AIM Server from the drop-down list.

The list of UCS AIM Servers appears.

5. Select the UCS Server from the drop-down list and enter SNMP access identification information.

CA Server Automation populates the UCS Server drop-down list with the UCS AIM Servers listed in the Cisco UCS pane. You can only manage those UCS Servers for which your CA Server Automation manager has a valid connection established.

Note: If the AIM resides on a remote system, CA Server Automation must discover the system first. After the Discovery, the AIM server appears in the drop-down list.

6. Click OK.

A new AIM instance for the selected Server is registered.

Note: If the instance is not in an error or in a stopped state, CA Server Automation starts to discover the associated environment. When the Discovery process is complete, you can start managing the Cisco UCS environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:

-  Discovery in progress
-  No polling
-  Error
-  Warning
-  Disabled
-  Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the amount of managed objects that are related to virtual and physical resources in your environment. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Server Automation adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Server Automation manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the AIM:

- Verify if the AIM Server is accessible.
- Verify if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
ping servername
```

2. Verify the output of the commands to find out whether the AIM server has a valid DNS entry and IP address.

If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress servername
```

Enter the correct IP address and AIM server name. For example:

```
192.168.50.51 myAIM
```

4. Click  (Validate) in the upper-right corner of the AIM Server pane.

If the error status remains unchanged, continue with the next procedure.

To verify if SystemEDGE is running:

1. Log in to the AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.

The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.

2. Start or restart SystemEDGE.

Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.

3. Change to the CA Server Automation user interface, AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the AIM Server connection.

If the error status remains unchanged, verify the validity of the data you gathered according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Server Automation discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Server Automation manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Server Automation manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Cisco UCS in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand Cisco UCS group.

The managed Cisco UCS resources appear.

CA Server Automation is now ready to manage the configured Cisco UCS environment.

Cisco UCS Management

The CA Server Automation integration with Cisco UCS lets you manage your UCS switches, chassis, and blades from a centralized user interface. The UCS Manager running on a switch is the location from which you can view UCS resources and perform management operations such as cloning, snapshots, power operations, and so on.

This section describes the resource management operations that you can perform on Cisco UCS resources from the Resources page. The Resources page displays basic information about the following UCS objects:

- Cisco UCS servers
- UCS Manager servers
- Chassis
- Blade servers
- Fabric interconnects
- Organizations

The Summary page lets you view information associated with that object (for example, a chassis summary can display its blades, or a blade summary can display its storage) and events associated with the resource.

If available, a Details page lets you view other resource information, such as system properties, software, hardware, performance, and so on. After you assign an automation source here, default access and management profiles are automatically created for the system and discovery is run.

Other pages can be available to perform resource management tasks. The right-click menus on Explore pane objects also let you perform UCS management tasks.

More information:

[How To Use Centralized Service Profiles](#) (see page 462)

[View Cisco UCS Resources](#) (see page 464)

[Associate Service Profiles with Blades](#) (see page 465)

[Back Up a UCS Manager Configuration](#) (see page 466)

[vNIC Templates](#) (see page 467)

[UCS Organizations](#) (see page 467)

[UCS Pools](#) (see page 468)

[How to Manage Port Profiles](#) (see page 472)

[UCS Action Types](#) (see page 474)

How To Use Centralized Service Profiles

Central service profiles that reside in the CA Server Automation Management Database 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 Database from UCS Managers, or create a central service profile in the Management Database.

From the Management Database, you can export central service profiles to any UCS Manager.

Manage Central Service Profiles

You can manage central service profiles from the Resources page. For access, select Cisco UCS Server in the Explore pane, and click Central Service Profiles in the right-hand pane.

To import service profiles from UCS Managers

1. Click the white triangle (Import) icon.
2. Use the Import Service Profiles dialog to select UCS Managers. Click Refresh to populate the Service Profiles list and select Import All Service Profiles, or select one or more from the list. To remove imported profiles from their UCS managers after import, select Delete Source Service Profiles.
3. Click OK.

The selected service profiles are imported into the Management Database.

To create or update a central service profile in the Management Database

1. Click the + (Create) icon, or select a central service profile and click the tool (Edit) icon.
2. Use the wizard pages to create or update the central service profile.

Note: You cannot specify pools and policies when you create a service profile in the Management Database; this information is for reference only. You can specify this information after you export the central service profile to the UCS Manager.

3. When the service profile is created or updated, click Finish.

To export service profiles to a UCS Manager

1. Select one or more central service profiles
2. Click the blue triangle (Export) icon.

The Export Service Profiles dialog appears.

3. On the Available UCS Managers list, select one UCS Manager and click a right arrow to transfer to the Selected UCS Managers list. Click the double right arrows to transfer all.
4. Click OK.

Note: Pools and policies are not exported; they must already reside on the target UCS Manager.

To delete service profiles from the Management Database

1. Select the service profiles to delete.
2. Click the - (Delete) icon.

View Cisco UCS Resources

The Resources page lets you view UCS resources at any level on the UCS object tree. For example, you can inspect the following objects to determine:

- Cisco UCS Servers - UCS resources by category, blade allocation, and imported service profiles
 - Centralized service profiles - UCS Manager assignments, import, and export
 - UCS Managers - Fabric interconnects, chassis, and organizations
 - Chassis - Number of blades mounted, number of fans (and their status), and input/output modules
 - Blade servers - Number of blades, whether powered on or off, whether associated with service profiles, and the OS host name
- Note:** Expand the blade in the Explore tree to see the OS host. The OS host name for a blade will be available after provisioning and discovery are complete.
- Individual blade - High-level inventory, including motherboard, CPU, memory, and storage
 - Fabric interconnects - High-level hardware and fans
 - Organizations - Pools, service profiles, and service profile templates

To view a Cisco UCS resource

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Find and select a Cisco UCS resource.
The resource page appears in the right pane.

Associate Service Profiles with Blades

Services profiles can be associated with blades, unassociated, or set to apply at failover.

To adjust a UCS service profile

1. Right-click and select Policy.
The Policy submenu appears.
2. Right-click and select Policy, Actions & Rules.
The Actions & Rules page appears.
3. Click Actions.
The Actions page appears.
4. Click + (Add new action).
The Action Definition: New page appears.
5. Click the action type Configure Service Profile on the Type drop-down list.
The Configure Service Profile form displays.
6. Specify the UCS resource details to which you want the service profile to apply.
Select the profile operation.
Note: If help desk approval is required, enter information as needed.
7. Click Save on the Actions drop-down.
The service profile relationship is modified.

More information:

[Configure Service Profile: Cisco UCS](#) (see page 594)

Back Up a UCS Manager Configuration

CA Server Automation supports backup of the following types of UCS Manager configuration information:

- Full state
- All configurations
- System configuration
- Logical configuration

The export/import capability emulates the Cisco UCS capability and allows you to create and rerun backup jobs.

To export a UCS Manager configuration

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Select a Cisco UCS Manager.
The UCS Manager page appears in the right pane.
4. Click Export/Import.
The Export/Import page appears.
5. In the Export jobs section, click + (Create).
The Create Backup Operation dialog appears.
6. Enter export information and click OK.
The export job starts and displays in the Export list.

To import a UCS Manager configuration

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Select a Cisco UCS Manager.
The UCS Manager page appears in the right pane.
4. Click Export/Import.
The Export/Import page appears.
5. In the Import jobs section, click + (Create).

The Create Backup Operation dialog appears.

6. Enter import information and click OK.

The import job starts and displays in the Import list.

vNIC Templates

CA Server Automation supports the creation and management of vNIC templates. You can specify the template target as either the service profile or a VM.

To create a vNIC template, select Use vNIC Template in the Service Profile wizard, and launch the Create vNIC Template dialog. You also can right-click on a UCS organization in the Explore pane.

UCS Organizations

You can use organizations to group related UCS resources to create a nested hierarchy of pools, service profiles, and service profile templates for UCS resource management. Organizations and sub-organizations can be created and deleted.

More information:

[Create a Sub Organization](#) (see page 467)

Create a Sub Organization

You can create organizations or sub-organizations on the UCS root tree.

To add an organization

1. Click Resources.

The Resources page appears.

2. Open the Explore pane.

Available groups, services, and systems appear.

Navigate to root, or open root and click a sub-organization

3. Right-click root or a sub-organization to select Management, Create Sub Organization.

4. Use the Create Sub Organization dialog to create the new sub-organization.

The sub-organization is created.

UCS Pools

CA Server Automation lets you create pools to manage UCS resources more efficiently.

Note: When pool range conflicts occur, a warning displays.

The following types of pools are available:

- UUID Pools
- MAC Pools
- World Wide Node Name (WWNN) Pools
- World Wide Port Name (WWPN) Pools
- Server Pools

WWNN and WWPN pools can be used for configuring blade to use remote storage (SAN).

More information:

[View a UCS Pool](#) (see page 468)

[Create a UCS Pool](#) (see page 469)

[Rename a UCS Pool](#) (see page 470)

[Delete a UCS Pool](#) (see page 471)

View a UCS Pool

The Resources page lets you view a UCS pool.

To view a UCS pool

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Click root at the top of the UCS organization tree, and navigate to the desired organization.
4. Click Summary.
The Summary page appears.
5. In the Components section, click the desired Pools type in the drop-down menu.
6. Select a pool to view, and click the tool (View) icon.
7. Click Cancel to return to the Pools list.

Create a UCS Pool

The Resources page lets you create UCS pools to manage UCS resources more efficiently.

To create a resource pool

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Click root at the top of the UCS organization tree, and navigate to the desired organization.
4. Click Summary.
The Summary page appears.
5. In the Components section, click the desired Pools type in the drop-down menu.
6. Click + (Add new pool).
The Create Pool dialog appears.
7. Use the dialog to complete the definition.
The pool is added to the pool list.

Note: You can customize the Quick Start menu to provide the function.

Rename a UCS Pool

The Resources page lets you rename a UCS pool.

To rename a UCS pool

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Click root at the top of the UCS organization tree, and navigate to the desired organization.
4. Click Summary.
The Summary page appears.
5. In the Components section, click the desired Pools type in the drop-down menu.
6. Select the pool to rename.
7. Click the double-arrow icon (>>).
The Rename Pool dialog appears.
8. Use the dialog to rename the pool.
The pool is renamed in the list.

Delete a UCS Pool

The Resources page lets you delete a UCS pool.

To delete a UCS pool

1. Click Resources.
The Resources page appears.
2. Open the Explore pane.
Available groups, services, and systems appear.
3. Click root at the top of the UCS organization tree, and navigate to the desired organization.
4. Click Summary.
The Summary page appears.
5. In the Components section, click the desired Pools type in the drop-down menu.
6. Select the pool to delete.
7. Click the - (Delete) icon.
The Delete confirmation appears.
8. Click Yes.
The pool is deleted.

How to Manage Port Profiles

Use the following process to manage port profiles using CA Server Automation.

1. Export plug-ins.

To establish the communication between Cisco UCS Manager and vCenter, generate and install one or more extension XML files in the target vCenter as follows:

- For vCenter 4.0, multiple extension files are required.
- For vCenter 4.0 update 1 version and above, export a single extension file from Cisco UCS Manager.

After the required files are exported, import them into vCenter as new plug-ins using vSphere Client. This is a one-time requirement for each UCS Manager and vCenter combination; a Cisco UCS Manager cannot use files exported from a different UCS Manager.

2. Export .vib file to ESX server.

Based on the ESX version, configure the ESX server by installing the appropriate .vib file component from the Cisco Nexus 1000V Virtual Ethernet Module Software. This package (jointly designed by Cisco and VMWare) enables a distributed virtual switch solution that is fully integrated with the VMware Virtual Infrastructure.

3. [Create Port Profile Network Topology](#) (see page 473)
4. [Create Port Profiles and Port Profile Clients](#) (see page 473)

Create Port Profile Network Topology

To push a port profile to VMware, the Cisco UCS Manager must have defined vCenter, datacenter, DVS folder, DVS, and profile client objects. The topology of these objects must match the topology in VMware. CA Server Automation lets you create the required topology.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Right-click a UCS Manager in the Explore tree, and click VMware to launch the vCenter Layout dialog
3. Expand vCenter and highlight a vCenter, datacenter, DVS folder, DVS, or profile client object.
4. Select Create New in the Actions drop-down menu.
5. Enter the required information, and click Finish. Selecting Enable on the DVS panel automatically pushes the associated port profiles to vCenter.

The vCenter, datacenter, DVS folder, DVS, or profile client object is added.

Note: To use this dialog to delete topology objects, select Delete on the Actions drop-down menu.

Create Port Profiles and Port Profile Clients

You can use the CA Server Automation vCenter Layout dialog to create port profiles and port profile clients.

Follow these steps:

1. Right-click a UCS Manager in the Explore tree and click VMware to launch the vCenter Layout dialog.
2. Highlight Port Profiles to create a port profile or an existing port profile to create a port profile client.
3. Select Create New in the Actions drop-down menu.
4. Enter the required information, and click OK.

The port profile or port profile client is created.

Note: You can also use this procedure to edit or delete port profiles and port profile clients. Highlight an existing port profile or port profile client, and select Edit or Delete on the Actions drop-down menu.

UCS Action Types

Cisco UCS resources can use CA Server Automation action types to create new actions that automate UCS power, resource allocation, and other operations when the assigned rule criteria are met. You can also schedule these actions to occur at specific times.

Chapter 7: 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:

[Add Proxy Servers](#) (see page 475)

[Cisco UCS Server](#) (see page 476)

[AppLogic Management](#) (see page 478)

[How to Configure AIX NIM Imaging](#) (see page 484)

[Solaris JumpStart Provisioning](#) (see page 489)

[Configure CA Network Automation](#) (see page 507)

[Storage Provisioning Features](#) (see page 508)

[How to Configure Software Delivery](#) (see page 511)

[Automating Processes with CA Process Automation](#) (see page 512)

[Event Forwarding](#) (see page 518)

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 AMLs (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 479)

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 477)

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\AIPCommon` 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.

AppLogic Management

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*.

AppLogic Prerequisites

CA Server Automation supports CA AppLogic 3.0 and 3.1. Review the following requirements.

CA AppLogic 3.0 and 3.1

- The AppLogic grid controller requires WS_API_R12 or later.

Note: If your AppLogic grid controller does not have the required version of the web services application (WS_API), contact CA Support.

- Verify that a WS_API application is provisioned and configured using the [Application Catalog Reference Guide](#).
- Set up ACL access to the grid for the default WS_API user.

For example, for a user *saapi*:

1. Add the default user to the admin group.
Group modify admin +local:user:saapi
2. Enable grid administrator access for the default user.
Grid modify_acl local:user:saapi=grid_administrator

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, including the default WS_API username with grid administration rights.
4. Select a proxy server to use, and click OK.

Note: If the AppLogic grid resides within your network, you may not require a proxy.

The server connection validates and the resources that are associated with the AppLogic server are added to Explore pane in the Resources tab.

More information:

[Add Proxy Servers](#) (see page 475)

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.

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.

Manage AppLogic Resources

CA Server Automation lets you manage your AppLogic grids, grid servers, applications, templates, and appliances.

This section describes the management operations that you can perform on AppLogic resources from the Resources page. The Resources page lets you view basic information and details about the following objects:

- AppLogic Grids
- AppLogic Applications
- AppLogic Grid Servers
- AppLogic Appliances

Click Resources, open the Explore pane, and select an AppLogic resource.

Information and operations for AppLogic resource display, including the following tabs:

Quick Start

Perform operations on a grid or application, such as provisioning a new application or starting an existing application.

Summary

Displays resource information (for example, number of CPUs in a grid, memory allocation for an application) and events associated with the resource.

Model

Displays a graphical representation of the components of AppLogic applications.

Right-click an AppLogic resource in the Explore pane to open the context menu for the resource. The context menu provides quick access to features such as Discovery, Refresh Grid, Provisioning, and application and appliance management.

Manage AppLogic Applications

You can perform management operations on an AppLogic application from the Explore pane.

Follow these steps:

1. Right-click an application in the Explore pane, select Management and one of the following options:

Start Application

Starts the application.

Stop Application

Stops the application.

Restart Application

Stops the application and then starts it again.

Delete Application

Deletes an application.

Copy Application

Creates a copy of an existing application with an input for the new name.

Migrate Application

Copies an application to another grid.

Rename Application

Renames the application with an input for the new name.

Configure Application

Set new values for the resources and configuration boundary properties for the application.

A dialog opens for the management operation.

2. Complete the dialog for the management operation.

A message appears confirming that the request was submitted.

Manage AppLogic Appliances

You can perform management operations on an AppLogic appliance from the Explore pane. These operations enable you to control the individual components of an application.

Follow these steps:

1. Right-click an appliance in the Explore pane, select Management and one of the following options:

Start Appliance

Starts the appliance.

Stop Appliance

Stops the appliance.

Restart Appliance

Stops the appliance and then starts it again.

2. Click OK.

A message appears confirming that the request was submitted.

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. 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 specify a root password, 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.

Note: The `ca_post_install.sh` script file has a default hash password that translates to the plaintext value of **admin**.

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 or Stop the NIM Adapter Daemon

The NIM adapter daemon starts automatically after installation or when the system is booted up.

To start the NIM adapter daemon manually, run the following command:

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

To stop the NIM adapter daemon manually, run the following command:

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

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.

Solaris JumpStart Provisioning

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/productname`) 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/productname/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/productname/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/productname/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 496)

[Extract an Installable Image from the Media](#) (see page 497)

[Add Packages to the Image](#) (see page 498)

[Add Patches to the Image](#) (see page 499)

[How To Modify Configuration Files](#) (see page 500)

[Configure SSH for JumpStart](#) (see page 505)

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 SMCssl /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 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 500)

[Edit the Package Table of Contents File](#) (see page 501)

[Edit the Profile File](#) (see page 502)

[Edit the Rules File](#) (see page 502)

[Edit the Configuration File](#) (see page 503)

[Edit the Finish File](#) (see page 504)

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.

SMCssl

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

SMCzlib

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

SMCssh501

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/SMCssl/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 fileys entries:

```
package SMClgcc add  
package SMCoss1 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.

Configure CA Network Automation

CA Network Automation lets you discover and manage Cisco routers and switches running the Cisco Internetwork Operating System (Cisco IOS) in 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 that it manages.

CA Server Automation provides CA Network Automation scripts in the folder *install-directory\nma_scripts*.

To use this integration, CA Network Automation must be preconfigured and scripts imported into the NetMRI system. 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. In Management mode, 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.

Storage Provisioning Features

As a System Administrator, you can provision storage in CA Server Automation for virtual and physical servers. To save time, you can create enhanced storage policies and associate them with managed servers. When servers need more space, you can select a different storage policy.

Storage provisioning is supported in Reservation Manager. You also can automate storage provisioning using CA Process Automation connectors.

The following table lists the supported storage vendors, protocols, servers, and features.

Storage Servers	Supported Protocols and Servers	Supported Provisioning Methods in User Interface
NetApp OnCommand	■ SAN-based iSCSI – VMware vCenter ESX, AIX LPAR Virtual I/O, physical and virtual servers (Windows, AIX, Linux, HP-UX, Solaris) ■ SAN-based FCP – VMware vCenter ESX, Cisco UCS, physical and virtual servers (Linux) ■ NAS-based CIFS – physical and virtual servers ■ NAS-based NFS – physical, virtual, and VMware vCenter ESX	■ Enhanced storage policies ■ Storage services (NetApp Provisioning Manager) ■ Provisioning policy (NetApp Provisioning Manager)
EMC SMI-S	■ SAN-based iSCSI – VMware vCenter ESX, AIX LPAR Virtual I/O, physical and virtual servers (Windows, AIX, Linux, HP-UX, Solaris) ■ SAN-based FCP – VMware vCenter ESX, Cisco UCS, physical, and virtual (Linux)	■ EMC SMI-S ■ Enhanced storage policies

For other storage tasks (discover, resize, move, deprovision), use the cadpmspm CLI.

Note: EMC SMI-S does not support the move option.

Storage Prerequisites

Review the following prerequisites before configuring storage:

- Review Release Notes for supported storage arrays.
- Register the Storage Provisioning Manager component in the user interface (Administration tab, Configuration tab, left panel, CA Server Automation Services, Storage Provisioning Manager).
- You must have a basic understanding of the CA Server Automation user interface and how to provision resources.

How to Configure Storage

Follow these steps to configure storage for CA Server Automation:

1. Configure your storage servers.

NetApp OnCommand

- a. Verify that your data center uses NetApp Storage Systems that include NetApp Operations, Provisioning, and Protection Managers.
- b. Create a storage service catalog using the NetApp Provisioning Manager.

EMC SMI-S

Verify that EMC storage systems have been configured on the SMI-S server.

2. If applicable, configure your IBM AIX LPAR iSCSI devices.
 - a. Access the Virtual I/O servers.

Example: `chdev -dev iSCSIAdapterDeviceName -attr initiator_name=UNIQUE-ID`
 - b. Verify that the ID is unique (for RFC 3270).

Example:
`iqn.2011-06.com.domain:storage:myMachine.0xa1afc0f.800010EBCD900004.`
 - c. Verify that Virtual I/O Servers have log in access using ssh.
3. If applicable, configure your AIX (physical and virtual) iSCSI devices.

Install the specific iSCSI ODM filesets (download from the EMC site) when connecting to EMC storage arrays.
4. If applicable, configure your VMware vCenter ESX iSCSI devices.
 - a. Configure the iSCSI Adapter.

On the vSphere Client Configuration tab, click Storage Adapters in the Hardware panel.
 - b. For Fiber Channel HBA, use the defaults.
 - c. Configure Discovery Addresses for iSCSI Initiators so storage resources on the network are available.
5. If applicable, configure your Windows (physical) iSCSI devices.
 - a. Install Microsoft iSCSI initiator on the remote Windows host for iSCSI storage attachment.

For Windows 2003, you can download the iSCSI initiator from the Microsoft website under downloads.
 - b. 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.
6. If applicable, configure your NFS devices.

Windows (physical or virtual)

- a. Install the NFS Client on the remote Windows host for NFS storage attachment.
 - b. 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.
7. If applicable, configure your CIFS devices.

Windows (physical or virtual)

For CIFS attachment, if the storage is already mapped and you attach to the same storage with the same CIFS credentials, an error occurs.

AIX, Linux, HP-UX, Solaris (physical or virtual)

For CIFS attachment on a UNIX platform, enable (or install and configure) the CIFS client package.

8. Add your storage servers to the user interface.
- a. In the user interface, click the Administration tab, Configuration tab.
 - b. In the left panel, go to Storage, Storage Providers, and click + (Add).
 - c. Enter the server information.

Note: The credentials on this dialog are used only to verify whether the storage server exists. This user must have privileges to schedule and provision storage on the storage system (NetApp Power Users and Administrators have these privileges). (The CA Server Automation login credentials are used to connect to storage server for storage provisioning.)
 - d. Click Validate to verify the connection.
9. Define credentials for CA Server Automation that match existing credentials for your storage server.
10. (Optional) If you are using CA Process Automation connectors for storage management, review the CA Process Automation documentation for prerequisites.

How to Configure Software Delivery

The Software Delivery service lets you deploy and patch software to managed systems in CA Server Automation.

1. Verify that the SDApapter service is running (Administration, Configuration page, Services List).

If the SDApapter service is not running, configure a Software Delivery server; the SDApapter service is started automatically.

2. Verify that each CA ITCM instance is configured for management by a single CA Server Automation instance, or provisioning can fail.
3. If you are using Active Directory, configure the Software Delivery server to use it as an external directory before authentication.

Note: When using a Windows Domain account, the user can be a local user that is a member of the Administrators group (full control access class permissions to all available security object classes); or the user can be a domain account that is also a member of the Local Administrators group.

Automating Processes with CA Process Automation

CA Process Automation automates routine administration tasks, improves operations efficiency and incident response handling, and ensures 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 513).
- 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

Specify a CA Server Automation server for single sign-on for a CA Process Automation server. To use single sign-on, CA Process Automation must be configured with CA EEM.

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

1. [Log in to the CA EEM UI](#) (see page 31) on the CA Process Automation server.
2. Verify that the required groups and users are available in Manage Identities, Groups in the CA EEM UI.
3. Reset the passwords of the users to the desired passwords.

4. 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: C:\Program

Files\CA\PAM\server\c2o\.c2orepository\public\certification\PAM.p12

EEM Certificate Password

Identifies the password for the certificate file.

Example: itpamcertpass

The CA EEM security settings are configured.

Access the CA Process Automation User Interface

All CA Process Automation functions (design, deploy, monitor, control, and audit) are available through a single browser-based UI. 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.

Follow these steps:

1. Select Start, Programs, CA, CA Process Automation Domain, Start CA Process Automation 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. 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 567) 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 SDM 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 569)

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.

Important! Before you edit the sysedge.cf file on the monitored server, verify that SystemEDGE runs in unmanaged mode, that is, the server is not registered in Policy Configuration. If SystemEDGE runs in managed mode, Policy Configuration in CA Server Automation can overwrite your changes.

To configure SNMPv1 traps by editing sysedge.cf

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

1. Configure the event manager to receive CA Server Automation traps or events.
2. Configure CA Server Automation to forward the events.

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

Follow these steps:

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

A confirmation message appears.

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

Your settings are updated and the configuration information appears. CA 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 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 Calndication.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 IM

For information about how to configure CA Spectrum to receive SNMP traps, see the *CA Spectrum SNMPv3 User Guide*.

Chapter 8: Monitoring Clusters

This section contains the following topics:

[Microsoft Cluster Service](#) (see page 523)

[IBM High Availability Cluster Multiprocessing](#) (see page 537)

Microsoft Cluster Service


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

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

How to Configure Microsoft Cluster Service Management Components

The following diagram provides an overview of the required actions to configure the management components. The diagram includes corresponding troubleshooting strategies in case of connection problems.

How to Configure the Management Components



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

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

Follow these steps:

1. [Review Requirements](#) (see page 356).
2. Review [Interactions Between MSCS Management Components](#) (see page 527).
3. Add a Microsoft Cluster Service Connection to the Manager.
4. Troubleshoot if [Server Connection to the Manager Failed](#) (see page 362).
5. Add the Discovered Microsoft Cluster Service AIM Instance.
6. [Troubleshoot the AIM Instance Connection](#) (see page 365).
7. Verify the Microsoft Cluster Service Group in the Resources Tree.

Review Requirements

Review the following requirements before you start configuring the management components of CA Server Automation:

- You are familiar with TCP/IP, SNMP, web services, and Windows Server operating systems.
- You are familiar with CA Server Automation, SystemEDGE, and your environment.
- You can access a CA Server Automation manager installation that includes:
 - Platform Management Module (PMM)
 - Application Insight Module (AIM)
 - Monitoring Agent (SystemEDGE)
- You can access the CA Server Automation user interface.
- You have valid credentials (user name and password) to access the servers in the environment that you want to manage.
- You know which protocol (HTTP or HTTPS) and port to use to access the server in your environment through web services. Default: HTTPS, Port: 443.
- You verified that the servers in your environment are running properly.

- If the PMM and AIM are installed on different systems, verify that the SNMP settings on the PMM and AIM systems are consistent. Read and write community strings and SNMP port number must be identical.
- You verified that the CA Server Automation manager discovered remote AIM Servers that you want to use.


Interactions Between MSCS Management Components

The following diagram illustrates how the components involved in MSCS monitoring interact. SystemEDGE and the MSCS AIM run on a Windows Server.

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

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

Interaction Between MSCS Management Components



Add a Microsoft Cluster Service to the Manager

You can add a Microsoft cluster using the Administration tab of the CA 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 Microsoft Cluster Services from the Provisioning section in the left pane.
3. Click  (Add) on the Microsoft Cluster Service pane toolbar.

The Add Microsoft Cluster Services dialog appears.

4. Enter the required connection data (server name, user, password, port), specify the preferred AIM, enable Managed Status.
5. Click OK.

The Microsoft cluster is registered.

When the network connection has been established successfully, the Server is added to the top right pane with a green status icon.

Note: If the connection fails, the Validation Failed dialog appears. If you click Yes, CA Server Automation adds the Server to the list with a red status icon indicating a connection failure. If you click No, nothing is added.

Server Connection to the Manager Failed

Symptom:

After I have added a Server connection under Administration, Configuration, the validation of the connection to the Server failed.

Solution:

The following procedures address the most common issues which can cause a connection failure:

- Verify, if the used Server connection data is still valid. If necessary, update the connection data.
- Verify, if the Server system is running and accessible.
- Verify, if the Management Service on the Server system is running properly.

To update the Server connection data:

1. Click  (Add) or  (Edit) that is associated with the failed connection.
2. Add the connection details, enable Managed Status, and click OK.

The connection data is updated.

3. Click  (Validate) in the upper-right corner to validate the new settings.

If the connection to the Server cannot be established, continue with the next procedure.

To verify if the Server system is running and accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
nslookup <Server Name>
ping <IP Address of Server>
```

2. Verify the output of the commands to find out whether the Server has a valid DNS entry and IP address.

If the Server is not in the DNS, add the Server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress <Server Name>
```

Enter the correct IP address and Server name. For example:

```
192.168.50.50 myServer
```

4. Click  (Validate) in the upper-right corner.

If the Server credentials and connection data are correct and you can ping the Server, the connection can still fail. In this case, it is possible that the Server causes the problem. If the connection to the Server cannot be established, continue with the next procedure.

To verify, if the Management Service on the Server system is running properly:

1. Contact the Administrator to access the Server system.
2. Log in to the Server system and open Administrative Tools, Services from the Start menu.

The Services window opens.

3. Select the service and start or restart the service.
4. Change to the CA Server Automation user interface, Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the Server connection.

If the connection to the Server fails, verify the validity of the data you gathered according to the requirements for this scenario.

Work with the administrator or support to fix the Server connection problem.

Add the Discovered MSCS AIM Instance

After adding a Microsoft Cluster Service connection to the CA Server Automation manager, add the AIM instance to manage the Microsoft Cluster Service environment.

Follow these steps:

1. Open the CA Server Automation user interface from the Start menu. Click Administration, Configuration.

The Configuration page appears.

2. Select Microsoft Cluster Service from the Provisioning section in the left pane.

3. Click  (Add) on the Discovered Microsoft Cluster Service AIM Instances pane toolbar.

The Add Microsoft Cluster Service AIM Instance appears.

4. Select the MSCS AIM Server from the drop-down list.

The list of discovered MSCS AIM Servers appears.

5. Select the MSCS Server from the drop-down list.

CA Server Automation populates the MSCS Server drop-down list with the MSCS Servers listed in the Registered Microsoft Cluster Service pane. You can only manage those MSCS Servers for which your CA Server Automation manager has a valid connection established.

Note: If the AIM resides on a remote system, CA Server Automation must discover the system first. After discovery, the AIM server appears in the drop-down list.

6. Click OK.

A new AIM instance for the selected Server is added. If the instance is not in an error or in a stopped state, CA Server Automation starts to discover the associated environment. When the discovery process is complete, you can start managing your Microsoft Cluster Service environment.

Troubleshoot the AIM Instance Connection

If the AIM Connection is in not-ready status, one of the following status icons appears:



Discovery in progress



No polling



Error



Warning



Disabled



Unknown

See the tooltips for more information about the AIM Instance status. The following troubleshooting sections provide detailed information and procedures to solve the issue.

The AIM Instance Status Icon Shows Discovery in Progress

Symptom:

After I add an AIM instance for a Server under Administration, Configuration, the status icon shows  (Discovery in progress).

Solution:

Wait until the Discovery process of the environment has completed. The discovery duration depends on the amount of managed objects that are related to virtual and physical resources in your environment. You can hover the cursor over the icon to display a tooltip that indicates the number of outstanding discovery requests. When the discovery job finishes, CA Server Automation adds a Server folder to the resources tree. Then you can start managing your environment.

The AIM Instance Status Icon Shows No Polling

Symptom:

After I add an AIM instance under Administration, Configuration, the status icon shows  (No polling).

Solution:

No specific actions are required for the associated instance. This icon indicates that the CA Server Automation manager does not poll this AIM. The AIM is not the preferred one.

If more than one AIM is configured to manage a particular server, PMM selects one of the AIMs as the current AIM. If you like to use another AIM, you can set the preferred AIM under Administration, Configuration, Provisioning. Click Edit of the server entry and select the preferred AIM. The preferred AIM becomes the current AIM.

The AIM Instance Status Icon Shows Error

Symptom:

After I have added an AIM instance under Administration, Configuration, the status icon shows  (Error). Unable to connect to the AIM.

Solution:

The following procedures address the most common issues which can cause a connection failure to the AIM:

- Verify if the AIM Server is accessible.
- Verify if SystemEDGE is running. Start or restart SystemEDGE if necessary.

To verify if the AIM server system is accessible:

1. Open a command prompt on the CA Server Automation manager system and run the following commands:

```
ping servername
```

2. Verify the output of the commands to find out whether the AIM server has a valid DNS entry and IP address.

If the AIM server is not in the DNS, add the AIM server to the Windows hosts file on the CA Server Automation manager system. Continue with Step 3.

If the Server is in the DNS, continue with Step 4.

3. Open the hosts file in the %windir%\system32\drivers\etc directory with an ASCII editor and add the following line:

```
ipaddress servername
```

Enter the correct IP address and AIM server name. For example:

```
192.168.50.51 myAIM
```

4. Click  (Validate) in the upper-right corner of the AIM Server pane.

If the error status remains unchanged, continue with the next procedure.

To verify if SystemEDGE is running:

1. Log in to the AIM server and run sysedge.cpl from the %windir%\Program Files\CA\SystemEdge\bin directory.

The SystemEDGE Control Panel appears, showing the running state of SystemEDGE.

2. Start or restart SystemEDGE.

Wait until the SystemEDGE Control Panel indicates that SystemEDGE is running.

3. Change to the CA Server Automation user interface, AIM Server pane on the manager system and click  (Validate) in the upper-right corner.

CA Server Automation validates the AIM Server connection.

If the error status remains unchanged, verify the validity of the data you gathered according to the requirements for this scenario.

The AIM Instance Status Icon Shows Disabled

Symptom:

After CA Server Automation has discovered AIM instances in the network, the status icons of several instances show  (Disabled). This AIM instance is not managed.

This status appears, if CA Server Automation discovers an AIM with the following relationships:

- The AIM is configured for a Server that has a valid connection to the CA Server Automation manager but is in unmanaged state.
- The AIM is connected to a Server that has not been configured.

Solution:

To change the status of the AIM instance to ready, do *one* of the following:

- Add the missing Server connection to the CA Server Automation manager.
- Edit the existing Server connection and change its managed status to enabled.

Verify the Microsoft Cluster Service in the Resources Tree

After successful configuration and discovery, newly discovered resources are listed in the Resources, Explore pane under the corresponding group.

Follow these steps:

1. Click Resources, and open the Explore pane.
2. Expand MSCS group.

The MSCS resources appear.

CA Server Automation is now ready to manage the configured MSCS environment. You can monitor the status and the properties of MSCS resources.

Register a Cluster

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

To register a Microsoft cluster from the user interface

1. Click Administration.
The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Microsoft Cluster Services.
The Microsoft Cluster Services section appears on the right.
3. Click + (Add) on the Registered Microsoft Clusters toolbar.
The Register New Cluster dialog appears.
4. Enter the required cluster name and access identification information, and click OK.
The Microsoft cluster is registered.

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

Remove a Cluster

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

Follow these steps:

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

Modify Cluster Properties

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

To modify cluster properties

1. Click Administration.
The Administration page appears.
2. In the Provisioning section of the Configuration pane, click Microsoft Cluster Services.
The Microsoft Cluster Services section appears on the right.
3. Select the cluster that you want to edit.
4. Click the Edit icon on the Registered Microsoft Clusters toolbar.
The Modify Cluster Properties dialog appears.
5. Edit the required properties and click OK.
The cluster properties are modified.

Microsoft Cluster Service Management

Microsoft Cluster Service Management lets you manage your Microsoft clusters, services and applications, and nodes. The Microsoft Cluster Service is the central location from which you can view all clusters and perform management operations.

This section describes the management operations that you can perform on Microsoft Cluster resources from the Resources page. The Resources page lets you view basic information and details about the following objects:

- Microsoft Clusters
- Services and Applications
- Nodes

Click Resources, open the Explore pane, and select one of the cluster resources; then click Summary for the resource.

The Summary page lets you view information associated with that object and events associated with the resource.

Monitor MS Cluster Services

You can monitor the status and the properties of MS cluster resources in detail.

To monitor cluster resources

1. Click Resources.

The Resources page appears.

2. Open the Explore pane.

Available groups, services, and systems appear.

3. Expand the MS Cluster Service folder and click the cluster object.

A list of cluster nodes and services object appears.

4. Click the Services and Applications object.

A list of service appears.

5. Click the service object.

The right pane displays general information, resources, and events.

The General Information panel displays the service name, status, and the name of the cluster it belongs to.

The Overview tab in the Resources panel displays resource details such as the resource name, type, and status. The Private Properties tab in the Resources panel displays private properties of each resource.

The Events panel displays the current events.


IBM High Availability Cluster Multiprocessing

CA Server Automation monitors High Availability Cluster Multiprocessing (HACMP) clusters remotely, detect any failure, and provides details on alerts in the cluster and any other environmental problems.

Interaction Between IBM HACMP Management Components

The following diagram illustrates how the components involved in IBM High Availability Cluster Multiprocessing (HACMP) management interact. The AIM Server is a Windows Server on which SystemEDGE and the HACMP AIM run. The communication between the AIM and the HACMP cluster uses Secure Shell (SSH). Because CA Server Automation can connect to multiple clusters, CA Server Automation gains an overall view of your HACMP environment.

Interaction Between HACMP Management Components



To add the required connection information for each required HACMP cluster, use *one* of the following methods:

- NodeCfgUtil.exe utility on the AIM Server

The connection information is written to the configuration file on the managed node. The HACMP AIM polls the configuration file and starts monitoring your HACMP environment through the master node.

Configure SSH

To monitor cluster nodes, configure SSH for remote access.

Follow these steps:

1. Install and run the SSH daemon on the cluster (nodes).
2. Configure local firewall to allow SSH connections.

Configure HACMP AIM with NodeCfgUtil in Dialog Mode

The *NodeCfgUtil.exe* is a utility that lets you modify the AIM configurations. Use the utility in dialog mode to configure which nodes the appropriate AIM manages.

Follow these steps:

1. Log in as Administrator and open Windows Explorer on the computer on which the AIM is installed.
2. Change to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory, and start NodeCfgUtil.exe.

NodeCfgUtil discovers and lists the installed AIMs in subsequent dialogs.

3. Enter 1 to add a new managed node.
4. Follow the on-screen instructions to complete the configuration. Each node requires a valid user name and password for authentication.
5. After the configuration, enter 0 to return to previous menus, or to exit the utility.

NodeCfgUtil writes a configuration file for HACMP (hacmp.cfg) to the *SystemEDGE_InstallPath*\plugins\AIPCommon directory. You can also use the NodeCfgUtil utility to edit or remove existing entries.

Example

The following example shows the Install Managed Node dialog for *mycluster* that has been successfully added to the configuration of the HACMP AIM. The HACMP AIM is a multi-instance AIM. You can repeat this procedure and can add more entities that you want to manage with this AIM.

```
**** Choose Managed Node ****
1. Microsoft Cluster
2. IBM HACMP
0. Go Back to Previous Menu
*****
Enter choice: 2
Enter following information for the IBM HACMP Node...
(At any point to go back to the previous menu, Enter 'CTRL Q').
```

```
1. Cluster Name: mycluster
2. User Name: administrator
3. Password: *****
4. Port [default=22]:
CAAC1016 Authenticating, please wait...
CAAC1019 Authentication SUCCESSFUL.
CAAC1023 Added Node Successfully.
Press any key to continue...
```

Configure HACMP AIM with NodeCfgUtil in Command Mode

The *NodeCfgUtil.exe* is a utility that lets you modify the AIM configurations. When you use the utility in command mode, you can only add managed nodes to an AIM configuration.

Note: Run NodeCfgUtil.exe as Windows Administrator.

This command has the following format:

```
(1) nodecfgutil -help
(2) nodecfgutil hacmp -u user -p password -h cluster_name [-t port]
```

-help

Displays usage information about the console.

hacmp

Specifies the virtual or physical environment.

-u *user*

Specifies the name of an administrative user.

-p *password*

Specifies the password of that user.

-h *cluster_name*

Specifies the name of the cluster.

-t *port*

(Optional) Specifies the port number.

Default: 22

Follow these steps:

1. Open a command prompt on the system on which the AIM is installed.

The command prompt appears.

2. Enter *one* of the following commands:

(1) `nodecfgutil -help`

(2) `nodecfgutil hacmp -u user -p password -h hostname [-t port]`

(1) Displays the usage information about the console.

(2) Authenticates and stores the passed credentials for IBM HACMP

The utility writes a configuration file for IBM HACMP (`hacmp.cfg`) to the `SystemEDGE_InstallPath\plugins\AIPCommon` directory.

CA IBM HACMP SystemEDGE AIM Traps

CA HACMP SystemEDGE AIM Trap Types

The following list provides the trap types for CA HACMP AIM. Refer to the MIB file for complete varbind descriptions.

hacmpAimInstanceAddedTrap

Sends a trap when a new Instance or Server is added.

Trap ID: 165800

hacmpAimInstanceRemovedTrap

Sends the trap when an Instance or Server is removed.

Trap ID: 165801

hacmpAimInstanceDataStatusChanged

Sends a trap when Instance or Server Data Status is changed.

Trap ID: 165802

hacmpAimNodeAddedTrap

Sends a trap when a node is added.

Trap ID: 165803

hacmpAimNodeRemovedTrap

Sends a trap when a node is removed.

Trap ID: 165804

hacmpAimResourceGroupAddedTrap

Sends a trap when a resource group is added.

Trap ID: 165805

hacmpAimResourceGroupRemovedTrap

Sends a trap when a resource group is removed.

Trap ID: 165806

hacmpAimResourceGroupMigration

Sends a trap when a resource group is migrated.

Trap ID: 165807

hacmpAimResourceAddedTrap

Sends a trap when Instance or Server a resource is added.

Trap ID: 165808

hacmpAimResourceRemovedTrap

Sends a trap when a resource is removed.

Trap ID: 165809

Chapter 9: Agent-less Monitoring

CA Server Automation provides agent-less monitoring of the supported virtual environments (except Hyper-V) and Windows systems (Remote Monitoring).

This section contains the following topics:

[Remote Monitoring](#) (see page 543)

Remote Monitoring

Remote Monitoring (RM) lets you monitor the health state of agent-less systems. RM provides the flexibility of monitoring systems without the need to install the monitoring agents (such as SystemEDGE) on the remote systems.

RM employs a mid-level manager named RM AIM to monitor the remote systems. RM AIM collects the metrics information using WMI queries on the remote Windows systems.

More information:

[Interaction Between Remote Monitoring Components](#) (see page 544)

[Advantages of Remote Monitoring](#) (see page 545)

[Features and Benefits](#) (see page 545)

[Architecture](#) (see page 547)

[Use Case Scenario](#) (see page 549)

[Configuration Prerequisites](#) (see page 550)

[Configuring Remote Monitor Systems](#) (see page 551)

[Managing Systems Using Remote Monitoring](#) (see page 554)


Interaction Between Remote Monitoring Components

The Remote Monitoring AIM accesses an RM system through a WMI connection to the root\CIMV2 namespace utilizing DCOM. DCOM requires local system administrator user and password credentials. If you want to monitor a Windows computer, you must provide these credentials which the RM AIM stores in a file. The password is encrypted.

Remote Monitoring collects and provides Windows system information performing WMI queries (port 135) on the monitored RM systems. WMI uses port 135 (default).

The following diagram illustrates these relationships.

Interaction Between Remote Monitoring Components



Advantages of Remote Monitoring

Remote Monitoring involves agent-less rather than agent-based technology and there are advantages to both strategies. Use this information when deciding whether to use RM or deployed agents.

RM offers the following benefits:

- Costs less to set up, configure, and deploy
- Simplifies software upgrades and maintenance
- Deploys quickly and is less intrusive on the monitored environment
- Utilizes fewer resources on the managed server

A deployed agent offers the following benefits:

- Provides more detailed data and higher levels of functionality for the monitored servers and applications
- Requires less network bandwidth to operate
- Provides a higher degree of scalability, scaling to thousands of servers
- Continues to monitor server health and conduct data gathering when network connections are unavailable (as agent can work autonomously)
- Provides stronger command and control functions over the managed servers

Features and Benefits

Remote Monitoring provides *seamless* integration of monitoring from an end-user perspective (that is, equal look-and-feel of management interfaces for the monitored systems whether by agent or RM).

RM includes features that let you manage systems by monitoring health states and key performance indicator (KPI) metrics. RM provides reports on system status and utilization metrics. RM includes benefits such as resilience, scalability, integration, and automation. The primary features and benefits are described in the sections that follow.

Agent-less Monitored Systems

Remote Monitoring enables seamless health monitoring for systems managed with agent-based and agent-less technologies.

The RM manager component (RM PMM) creates CIM system objects representing the RM systems and their health state.

This information is presented in the Dashboard and the Resources Panel.

Follow these steps:

1. Open Resources, Explore and expand the Remote Monitoring folder.
The discovered systems appear in the components tree.
2. Select a system.
The page of that system appears in the right pane.
3. Open the Remote Monitoring tab.
The agent-less gathered data appears.

Key Performance Indicator Metrics

Remote Monitoring collects and provide Windows metric information by performing WMI queries on the monitored RM systems. A rich set of information is available in various Win32 CIM classes, made available through the RM AIM.

Visualization

The RM UI lets you configure the following information:

- What systems are remotely monitored
- What metrics are collected for those systems
- If and how those metrics are monitored (including severity and threshold)

Configuration

Remote Monitoring monitors KPIs out-of-the-box on a remote system when it is selected for monitoring without requiring configuration of the monitoring being performed. You can adjust the out-of-the-box monitoring thresholds to suit your needs.

You can also define and store configuration settings in a configuration set, which can then be assigned to one or more RM systems.

Access Control

When a user logs in to the UI as admin or as a nonadmin user, security mechanisms provide authentication and authorization functionality. Remote Monitoring allows or disallows certain actions (such as configuring an RM system) based on whether the user is an admin or nonadmin user.

The RM AIM accesses the RM system through a WMI connection to the root\CIMV2 namespace (using DCOM). The local RM system administrator user and password credentials are required for access. These credentials (provided by the user when an RM system is to be monitored) are stored in a file using password encryption.

Resilience

The RM AIM is a separate process from SystemEDGE; an error in the RM AIM does not cause SystemEDGE to crash. If the RM AIM crashes or no longer responds to SystemEDGE requests, the *RM AIM alive* check in SystemEDGE restarts it.

Scalability

There is one RM AIM per SystemEDGE and each RM AIM can monitor approximately 200 RM systems. There is a single RM PMM per manager and each RM PMM can manage approximately 20 RM AIMs. The default configuration set contains ten monitored metrics with two monitors for each metric.

In terms of SystemEDGE scalability, this results in the following:

- $10 * 2 * 200 = 4000$ monitorTable entries
- $10 * 200 = 2000$ aggregateTable entries

Integration

RM monitor information is exposed in an SNMP MIB to enable easy access for 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
- Svc_StoppedAuto

Managing Systems Using Remote Monitoring

Access the RM information and settings necessary to manage your systems by highlighting a managed resource in the Resource pane and clicking Remote Monitoring. The Remote Monitoring pages let you perform the following actions:

- Add remote systems for monitoring
- Manage queries
- Manage credential settings
- Create configuration sets
- Manage configuration entries

For the Dashboard, the following RM modules are available:

- CA SystemEDGE Machines Status
- CA SystemEDGE Objects Status

Add Remote Systems for Monitoring

Use the Systems page to enter system information for systems you would like to remotely monitor.

To add a system

1. Click + (Create New).

The Create New pane appears.

2. Enter the name of the system you want to monitor remotely in the RM System name field and edit the following settings (if necessary):

RM System Name

Specifies the name of the RM System. Using user interface, you must enter RM System name in FQDN notation only, for example "vm1234.ca.com". Using "rmonwatch" utility, you can also specify RM system name by Short name or IP Address.

Status

Specifies whether the system is active or in maintenance.

Protocol

Specifies whether the protocol is DCOM or SOAP.

Max instances

Specifies the maximum number of instances created in the instance table by any query to this system.

Credentials

Specifies the user credentials for the remote system.

Config sets

Specifies the config set (or group of metrics) that will be collected for the remote system.

3. Click Save.

The system is added to the list of systems you are remotely monitoring.

Viewing Query Results

You can use the Queries page to view the query results associated with RM systems.

The Query page lets you perform the following actions:

- View detailed query results and settings (highlight a query in the Queries table and select Results or Settings).

- Filter query results based on system, status, configuration set, or specific query (use the binoculars to show or hide the query filters).
- Manage the information that displays in the query table (click a column header to sort columns in ascending or descending order and to add or remove columns).

Managing Credential Settings

You can use the Credentials page to manage the individual credential settings associated with RM systems.

The Credential page lets you perform the following actions:

- Add credentials (use the Create New (+) icon, enter settings in the Details of Individual Credentials pane, and click Save).
- Delete credentials (highlight existing credentials, and click (-) icon).
- Edit credentials (highlight existing credentials, update the settings in the Details of Individual Credentials pane, and click Save).

Create Configuration Sets

Remote Monitoring provides several out-of-the-box configuration sets that should not be changed. Use the Configuration Sets page to create custom configuration sets to suit your needs.

To create a configuration set

1. Click + (Create New).

The Details of Individual Configuration Set pane appears.

2. Enter a name for the new configuration set, enter a description, and highlight the configuration sets to include in the new set (press the Ctrl key to highlight more than one entry).
3. Click Save.

The new configuration set is added to the Config set list.

Note: You can also use the Actions drop-down list to clone and delete your custom configuration sets.

Collection Engine and Report Support for Remote Monitoring Metrics

The Collection Engine and Reports support a fixed set of RM metrics in the default config set.

As a result, assign the default config set (or a config set or group of sets containing those metrics) to all systems for which you want to use reports.

The supported default config set metrics are as follows:

- CPU_PercentIdle
- Disk_PercentIdle
- Event_SystemErrors
- Mem_PercentUsed
- FSys_FreeMB
- Fsys_FreeMBDecrease
- Net_QueueLength
- Proc_PercentCPU
- Proc_PercentMemory
- Svc_StoppedAuto

Managing Configuration Entries

Use the Configuration Entries page to view and manage the configuration settings associated with queries.

To view or manage configuration settings

1. Highlight a query in the Configuration Entries table. You can apply filters to the entries for configuration set, severity, query class, and escalation severity using the show and hide filter icon (binoculars).

The Details of Individual Configuration Entry pane appears.

2. View or update the following values and click Save.

Index

Specifies a unique index for this configuration entry within the configuration set.

Query name

Specifies the name of the query (cannot include the '.' symbol).

You can use the same query name in a different config set; however, when applying more than one config set to a system, ensure the uniqueness of all query names. If the Qualifier is set to fixed entry, you cannot rename the query.

Description

Specifies the description for the configuration entry.

Interval

Specifies the interval (in seconds) between successive executions of the query and evaluations of the monitor (the value must be a multiple of 30 seconds).

Query class

Specifies the query class for the configuration entry.

Query scope

Specifies the scope to apply to the query.

Query property

Specifies the property of the Query class.

Config set

Specifies the name of the configuration set (do not use ',').

Query depends

Specifies that a query (Q2) depends on another query (Q1); such that Q2 is only created based on the result of Q1.

Qualifiers

Specifies additional information related to the configured queries and monitors.

Possible values are as follows:

- Entry cannot be deleted and query name cannot be changed (fixed entry)
- Query is executed only once (at least one time successful)
- Query is no longer executed if it was unsuccessful
- Query is not shown in the query table
- Results are shown per instance
- Results show previous values instead of current ones
- Results show increasing delta values
- Results show decreasing delta values
- Aggregate Monitors with the same object data and severity as AND relation

Query total

Specifies the property of the query class to apply as a total reference.

Query scale

Specifies the scale to apply for the property value (for example, *100, /1024 or /1024*100) This scale is used as default for the query scale in the query table. The value of the query property is multiplied or divided by the scale before storing the value in the result attributes.

Query instance

Specifies the properties of the query class to use for instance naming in the instance table.

Condition

Specifies the condition for comparing the result attribute value to the threshold and escalation threshold.

Threshold

Specifies the threshold that the result attribute value is compared against.

Severity

Specifies the severity to use for the SysEDGE object state model if the threshold condition is met.

Obj. Class

Specifies the class name to use for the SysEDGE object state model (do not use '*').

Obj. Instance

Specifies the instance name to use for the SysEDGE object state model (do not use '*').

Obj. Attribute

Specifies the attribute name to use for the SysEDGE object state model (do not use '*').

Lag

Specifies the number of times the threshold (escalation) condition must be met to cause a state change in the SysEDGE object state model.

Escal. delta

Specifies the difference to the threshold required to indicate an escalation condition.

Escal. severity

Specifies the severity to use for the SysEDGE object state model if the escalation condition is met.

Result

Specifies the result attribute of the query in the query table or instance table to monitor with SysEDGE monitors.

The configuration settings are updated to reflect any changes.

Chapter 10: Using Rules and Actions

This section contains the following topics:

[Rules and Actions](#) (see page 561)

[Use Cases for Policies](#) (see page 653)

[Configuring Data Collection](#) (see page 655)

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 SDM, which is a customer support application that manages calls, tracks problem resolution, shares corporate knowledge, and manages IT assets. If you want to open tickets automatically when your action is triggered, you can set your actions to interact with CA SDM. This arrangement is useful for actions requiring third-party approval. After the third party approves your ticket in CA SDM, the action will automatically run.

You can also schedule your actions to run at specified times using the initiation component. The current parameters for the action are saved when you create the job. If you change the action details after the job has been submitted, it will not have an impact on jobs that you have already scheduled to run. If you must change the action details of a job that has already been scheduled, open the job that uses the action and save it again to update it with the new action details.

Configure CA SDM

For CA SDM releases before Version 12.5, configure CA SDM properly with the appropriate ticket status codes, so that you can set your actions to open issues automatically when necessary.

Note: The release number of CA Server Automation and CA SDM need not be the same, as long as the two products do not share a database.

To configure CA SDM

1. Log in to your CA SDM server by typing the following information in your web browser:
`http://servicedesk_servername:8080`
The CA SDM splash screen appears.
2. Enter your user name and password, and click Log In.
The CA SDM main page appears.
3. Click Administration and expand the Service Desk tree node in the left pane.
4. Select Requests\Incidents\Problems and then Status.
The Request\Incident>Status List appears.
5. Click Create New.
A Create New Request Status window opens.

6. Type **Approved** in the Symbol text box, select Active from the Record Status drop-down list, type **APP** in the Code text box, and click Save.

The new request status appears in the list.

7. Type **Rejected** in the Symbol text box, select Active from the Record Status drop-down list, type **REJ** in the Code text box, and click Save.

The new request status appears in the list.

CA SDM setup is complete and you can now automatically open requests when an action is triggered.

Configure the CA SDM Ticket Status Setting

CA SDM versions before 12.5 used default status code settings of APP (Approved) and REJ (Rejected) for help desk tickets. CA 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 SDM Version 12.5, new ticket status codes are supported. PRBAPP (Approved) and PRBREJ (Rejected) must be associated to the existing approval codes in CA 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 SDM 12.5 status codes.

4. Save and close the file to enable the configuration change.

Rule Planning

Consider the following points when setting up rules and actions:

- Which VMs, servers, and services do you want to analyze?
- What actions do you want to take when CA 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 SDM must be configured to use this option.

The Ticket Types and Templates fields become enabled.

Note: Actions that specify a help desk approval requirement cannot be used for action scheduling. If you need the same action for a scheduled action, create a second action that does not include the help desk approval requirement.

9. Select Auto close ticket on approval if you want to close the ticket automatically after it is approved.
10. Select a ticket type from the Ticket Types drop-down list. The following types are valid options, but depend on your configuration:

- Default
- Incident
- Problem
- Request

The Templates drop-down list is updated with the templates associated with the ticket type you selected.

11. Select a template from the Templates drop-down list.

The fields are populated with predetermined values depending on the ticket model you are using.

12. Click Save.

A confirmation message notifies you that your save was successful.

For testing purposes, you can run the action from the Actions page by selecting the action and clicking the Run action icon.

Action Types

Several categories of action types are available.

Note: When using special or reserved characters in any operation, consider operating system and shell behavior. Behaviors include, but are not limited to, the invocation of custom scripts run by the operating system shell. For more information about shell behavior and how to escape special characters, see the Microsoft TechNet website at <http://technet.microsoft.com/en-us/library/cc723564.aspx>.

Predefined Action Types

Predefined action types are commonly used actions that are available for you to use when creating actions for your rules. Action types are calling command-line utilities. All action types are listed in a drop-down list in the Policy, Actions & Rules pages of the user interface.

Note: For detailed descriptions of action types, see the *Online Help*.

Custom Action Types

You can create custom action types using substitution strings rather than typing the full command line. The custom action types are added to the drop-down list of predefined action types. You can control user access to custom actions, in general, or you can control access to individual custom actions through the Administration page in the user interface.

The Run Command Script action type provides string substitutions that let you perform multiple actions on servers. String substitutions provide more flexible rules and reduce the need for custom scripts. The following string substitutions are available:

- %ACTIONNAME%
- %EVENTMESSAGE%
- %EVENTSOURCE%
- %RULENAME%
- %SERVER%
- %SERVICE%

The following string substitutions are only valid for actions running in an action sequence:

- %STDOUT% - standard output
- %STDERR% - standard error
- %EXITCODE% - action exit code

Action Sequences

Action sequencing is treated as an action type and is listed in the drop-down with the other action types in the Policy page. Action sequencing lets you define multiple actions for a rule in a specified sequence and run them as a single action. You can save the sequence of actions you specified with a name and that sequence is saved to the Management DB for repeat usage. You can schedule your action sequences as a job using the Policy, Actions & Rules pages in the user interface. CA SDM support for action sequencing is handled differently from other action types. You can set help desk approval for individual actions running in a sequence, but you cannot set help desk approval for the overall action sequence.

Consider these key points when using action sequences:

- Do not configure sequences that create infinite loops. The action sequence is performed synchronously, but some actions are performed asynchronously. Therefore, if you are expecting certain actions to have completed their tasks when they return, use care. Some actions that are typically long running and asynchronous have a -wait parameter that causes them to wait until their task is complete before returning or after a specified timeout.
- If you attempt to delete an action that is associated with an action sequence, the product prevents you from deleting that action.

- If your action sequence terminates abnormally, it restarts at the last known sequence when the Policy Manager restarts. You can manually cancel an action sequence in progress through the user interface or from a web service.
- When you specify the %STDOUT% (Standard Output), %STDERR% (Standard Error), or %EXITCODE% (Action Return Code) substitution string actions in a custom action running in an action sequence, the standard output/standard error/exit code of the previous action can be piped into the current action. Piping uses the output of the first action as input for the next action. If you redirect the output in your action, then it cannot be piped to the next action. For example, if the custom action *ipconfig* is redirected to a text file named *ipconfig_output.txt*, then that output is not available for piping to the next action.

List of Predefined Action Types

This section describes the following predefined action types that are available to create actions for policy rules.

Add Disk: VMware vCenter

The Add Disk action type lets you add a disk to a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to add a disk. Select one from the drop-down list.

Datastore

Specifies the name of the datastore associated with the ESX server for the selected VM. Select one from the drop-down list.

Drive Size

Specifies the size of the additional disk. Enter a value and select MB or GB from the drop-down list.

SCSI Controller

Specifies the SCSI controller to use to create the additional disk. Select one from the drop-down list.

Thin Provisioning check box

Specifies whether to enable thin provisioning.

Disk Mode

Specifies the disk mode. Select one of the following from the drop-down list:

- Persistent
- Independent Persistent
- Independent Non-persistent

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Add Network Interface: VMware vCenter

The Add Network Interface action type lets you add a virtual NIC to a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to add a virtual NIC. Select one from the drop-down list.

Device Type

Specifies the device type. Select one from the drop-down list.

Network

Specifies the network associated with the ESX server for the selected VM. Select one from the drop-down list.

You can distinguish the names of Standard Switches and Distributed Virtual Switches based on the following naming convention:

- For Standard Switches, the name is the network name.
- For Distributed Virtual Switches, the name is a concatenation of the dvPort group name followed by the Distributed Virtual Switch name enclosed in parentheses: dvPortGroupName (dvSwitchName)

MAC Address

(Optional) Specifies a MAC address. Leave the field blank if you want the MAC address to be autogenerated.

Wake on LAN check box

Specifies whether to set the virtual NIC to wake on LAN.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Add Server to Service

The Add Server to Service action lets you add servers to an existing service.

The Details section of the action definition contains the following fields:

Service Name

Specifies the name of the service.

Server List (comma delimited)

Specifies the list of servers to add to the service.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Change Machine State: Microsoft Hyper-V

The Change Machine State action type controls the state changes of the virtual machines in your Hyper-V environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the server on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine for which to change the state. Select one from the drop-down list.

State

Specifies the desired state of the virtual machine. Select one of the following from the drop-down list:

- Turn off
- Shutdown
- Save
- Pause
- Start

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Clone Machine: Solaris Zones

The Clone Solaris Zone Machine action type configures and installs a new zone by copying the data from an existing zone. You cannot perform this operation for global zones or when a zone is in the installed state.

The Details section of the action definition contains the following fields:

Zone Host

Defines the Solaris Zones host that contains the zone to clone.

Zone

Defines the zone to clone. You can use text extracted from event messages.

Name

Defines the name of the new zone. You can use automatically generated text or text extracted from event messages.

Path

Defines the installation path of the new zone.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure CPU/Memory: IBM LPAR

The Configure CPU/Memory action type lets you set limits on CPU and memory resources allocated to a virtual machine in IBM LPAR environment.

The Details section of the action definition contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Displays the unique name for the partition.

Profile Name

Specifies the name of an existing profile for the selected LPAR.

Operations

Specifies the operation to perform. Select one of the following from the drop-down list:

- Add Memory Units
- Subtract Memory Units
- Add Processors
- Subtract Processors

Processors

Specifies the number of processors for addition or removal.

Adjustment Type

Specifies the Adjustment Type. Select one option:

- Dynamic Adjustment Only
- Dynamic Adjustment and Update Profile

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure CPU/Memory: Microsoft Hyper-V

The Configure CPU/Memory action type controls the number of CPU and memory shares allocated to a virtual machine in your Hyper-V environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the server on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine for which to change the state. Select one from the drop-down list.

CPU Allocation

Specifies the CPU Allocation of the virtual machine. Adjust one of the following from the drop-down list:

- Number of CPUs
- CPU Reserved %
- CPU Weight
- CPU Limit %
- Present CPUID

Memory Allocation

Specifies the memory share allocated to the virtual machine in megabytes.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure CPU/Memory: VMware vCenter

The Configure CPU/Memory action type lets you set limits on CPU and memory resources.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Target VM Machine

Specifies the name of the virtual machine for which to adjust resources. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Operations

Specifies the operation to perform. Select one of the following from the drop-down list:

- Set CPU Limit
- Set Memory Limit
- Set CPU Reservation
- Set Memory Reservation

MHz, MB

Enter a value appropriate to the operation you select.

Unlimited check box

Allows the operation you select unlimited use of the resource.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: Cisco UCS

This action type lets you configure a power management action for a UCS blade server.

The Details section of the action definition contains the following fields:

UCS Manager

Specifies the name of the UCS Manager

UCS Chassis

Specifies the name of the UCS Chassis

UCS Blade

Specifies the name of the UCS Blade

Power Operations

Select an operation from the drop-down list:

Cycle Immediate

Immediately power cycle the blade

Cycle wait

Power cycle the blade which notifies all applications about its shutdown

Hard Reset Immediate

Plug back to power on the blade similar to unplug the power of the blade

Hard Reset wait

Unplugs the power of the blade. Prior to unplugging, the blade notifies all applications about its shutdown

Soft Shut Down

Shuts down the blade. Prior to shutdown, the blade notifies all applications about its shutdown

Shut Down

Shuts down the blade immediately

Boot up

Boots up the blade

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: IBM LPAR

The LPAR Configure Power action type controls power settings on LPARs.

The Details section of the action definition contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Specifies the partition name to be controlled.

Operation

Specifies the power operation to perform. If you select Activate, complete the following fields in the Operation Options section:

Partition profile

Specifies the partition profile used to activate the power settings.

Key Lock

Specifies the key lock mode in the partition profile. Key Lock establishes the power-on and power-off modes allowed for the system and is either manual or normal. For security reasons, it is not recommended that you set the key lock position to manual.

Boot mode

Specifies the boot mode in the partition profile. The system uses this boot mode to start the operating system on the logical partition unless you specify otherwise when activating the partition profile. CA Server Automation supports the following valid boot modes:

normal

Starts the logical partition as normal. (Use this option to perform most everyday tasks).

open_firmware

Boots the logical partition to the open firmware prompt. Service personnel use this option to obtain additional debug information.

If you select Shutdown, complete the following fields in the Operation Options section:

Delayed

Shuts down the logical partition using the delayed shutdown sequence. This sequence allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it ends abnormally. The next restart may be longer than typical.

Immediate

Shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

OS Shutdown

Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

OS Shutdown Immediate

Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

If you select Restart, select *one* option from the Operation Options section:

Partition profile

Specifies the partition profile used to restart the partition.

Immediate

Shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option causes undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

OS Shutdown

Shuts down the logical partition typically by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

OS Shutdown Immediate

Shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: Microsoft Hyper-V

The Configure Power action type controls the startup and shutdown of a virtual machine in your Hyper-V environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the server on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine for which to change the state. Select one from the drop-down list.

Start Action

Specifies the action to perform when the Hyper-V Server starts. Select one of the following from the drop-down list:

- Always
Always starts the VM when the Hyper-V Server starts.
- Auto
Automatically starts the VM when the Hyper-V Server starts.
- None
Does not start the VM when the Hyper-V Server starts.

Start Delay

Adjust the delay (in seconds) to start a VM after the Hyper-V Server starts. Select one from the drop-down list.

Shutdown Action

Specifies the action to perform when the virtual machine shuts down. Select one of the following from the drop-down list:

- Off
Turns off the VM before Hyper-V Server shuts down.
- Save
Saves (Suspends) the VM before Hyper-V Server shuts down.
- Shutdown
Shuts down the VM before Hyper-V server shuts down.

Recovery Action

Specifies the action to regain the previous details of a virtual machine when the Hyper-V Server fails. Select one of the following from the drop-down list:

- None
Does not take a specific action when the Hyper-V Server starts after the server fails.
- Restart
Restarts the VM when Hyper-V Server starts after the server fails.
- Revert
Reverts the VM with the latest snapshots when the Hyper-V Server starts after the server fails.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Power: VMware vCenter/Adjust vApp Power

The Configure Power action type controls power settings on the virtual machines and vApps in your VMware vCenter environment.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

VM/vAPP

Radio button to specify the target type, VM or vApp.

Target

Specifies the name of the virtual machine or vApp for which to adjust power. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Power Operation

Specifies the power operation to perform. Select one of the following from the drop-down list:

- VC Power On
- VC Power Off
- VC Power Reset
- VC Power Suspend
- VC Power Shutdown
- Power on vApp
- Power off vApp
- Suspend vApp

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Service Profile: Cisco UCS

The Configure Service Profile action type lets you associate, disassociate or failover a service profile to a UCS blade server.

The Details section of the action definition contains the following fields:

UCS Manager

Specifies the name of the UCS Manager

UCS Chassis

Specifies the name of the Cisco UCS chassis

UCS Blade

Specifies the name of the Cisco UCS blade

Service Profile

Specifies the name of the service profile

Profile Operations

Select a profile from the drop-down list:

Associate

Associates a service profile to a blade

Unassociate

Unassociates a service profile from a blade

Failover

Use this option, a check box appears against a service profile that is used to automatically failover the service profile to next available blade. By default, the check box is selected, and both chassis and blade drop-down are disabled.

Clear the check box to select the desired chassis and blade to failover a specific service profile.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Configure Shares: VMware vCenter

The Configure Shares action type controls CPU and memory shares for virtual machines in your VMware vCenter environment.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Target VM Machine

Specifies the name of the virtual machine for which to adjust shares. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Operations

Specifies the operation to perform. Select one of the following from the drop-down list:

- Set CPU
- Add CPU
- Subtract CPU
- Set Memory
- Add Memory
- Subtract Memory

Values

Enter a value appropriate to the operation you select.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Convert Template to VM: VMware vCenter

The Convert Template to VM action type lets you convert a template to a virtual machine.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which the VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the data center where the VM is located. Select one from the drop-down list.

VC Compute Resource

Specifies the cluster or VMware ESX host where the VM is to be created. Select one from the drop-down list.

VC ESX Servers

Specifies the VMware ESX server where the VM will reside. Select one from the drop-down list.

VC Resource Pool

Specifies the name of the resource pool from which you want to select the VM for cloning. Select one from the drop-down list.

VC Template

Specifies the name of the template you want to convert. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Convert VM to Template: VMware vCenter

The Convert VM to Template action type lets you convert a powered off virtual machine to a template.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

VC Virtual Machine

Specifies the name of the virtual machine you want to convert. Select one from the drop-down list or use text extracted from event messages.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Create Event

The Create Event action type lets you create events like system discovery, system deletion, multiple system discovery, and system management status changes.

The Details section of the action definition contains the following fields:

Event Status

Specifies the status of the event.

Event Component

Specifies the component name involved in the event.

Event Message

Specifies the message that the event generated.

Event Source

Specifies the source of the event.

Event Target

Specifies the target of the event.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Create Report

The Create Report action type helps you automatically generate reports. You can schedule this action so that it creates reports regularly. This action type can also be created from the Reporting tab.

The Details section of the action definition contains the following fields:

Report Type

Specifies the type of the report that is created. For explanation of the available report types and the related creation options, see Reporting.

The generated report can be viewed on the Reporting tab in the Schedule Reports folder.

Create Service

The Create Service action type lets you organize the servers you monitor into a logical service that reflects the resources required by your business needs.

The Details section of the action definition contains the following fields:

Service Name

Specifies the name of the service.

Server List (comma delimited)

Specifies the list of available servers.

Lower Threshold

Specifies the lower threshold of the service as a whole.

Upper Threshold

Specifies the upper threshold of the service as a whole.

Lag

Defines how often the rule must evaluate as true before the action triggers. Some actions should trigger after a single occurrence, while others should trigger only after a number of occurrences signal a persistent problem.

Priority

Specifies the order in which to execute actions in a single polling cycle.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: IBM LPAR

The LPAR Delete Machine action type lets you delete a specified LPAR.

The Details section of the action definition contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Defines the partition name to be deleted.

Note: The partition being deleted must be powered off for this action. This action erases the logical partition and the logical partition configuration data stored in the partition profiles.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: Microsoft Hyper-V

The Delete Hyper-V VM action type deletes a virtual machine from your Hyper-V Server environment.

The Details section of the action definition contains the following fields:

Hyper-V Host

Specifies the name of the host on which Hyper-V Server resides. Select one from the drop-down list.

Hyper-V VM Name

Specifies the name of the virtual machine that you want to delete. Select one from the drop-down list.

Attached Resources

Specifies the resources attached to the virtual machine that you want to delete. Select the resources that you want to delete:

- Hard Drive
- Floppy Drive
- DVD/ISO Image

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: Solaris Zones

The Delete Solaris Zone Machine action type deletes a zone from a Solaris Zones host.

The Details section of the action definition contains the following fields:

Zone Host

Defines the Solaris Zones host that contains the zone to delete.

Zone

Defines the zone to delete. You can use automatically generated text or text extracted from event messages.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Delete Machine: VMware vCenter

The Delete vCenter VM action type deletes a virtual machine from your VMware vCenter Server environment.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center on which the virtual machine resides. Select one from the drop-down list. Your selection populates the Target VM Machine drop-down list with the names of VMs associated with the data center.

Target VM Machine

Specifies the name of the virtual machine that you want to delete. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Discover Host by Name

The Discover Host by Name action type lets you discover a host using a specified host name.

The Details section of the action definition contains the following fields:

Host Name

Specifies the host name.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Discover Network

The Discover Network action type lets you discover networks that are available in your domain.

The Details section of the action definition contains the following fields:

Network ID

Specifies the network ID to be discovered.

Network Name

Specifies the network name to be discovered.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Distributed Switch: VMware vCenter

Use this action type to manage distributed virtual switches.

The Details section of the action definition contains the following fields:

Operation

Select one of the following operations:

- Add Port Group
- Remove Port Group
- Update Port Group

Virtual Center

Specifies the vCenter Server. Select one from the drop-down list.

Virtual Switch

Specifies the virtual switch you want to manage. Select one from the drop-down list.

Port Group

Specifies the port group name. Select one from the drop-down list.

Bind Type (Optional)

Select one of the following bind types:

earlyBinding

Assigns the ports when the VM binds to the portgroup. This type of binding ensures connectivity at all times, but permanently reserves the port. This binding type is the default.

lateBinding

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. LateBinding is configurable through vCenter.

ephemeral

Assigns a port to a VM if the VM is powered on and its NIC is in connected state. This binding type reassigns the port when the VM is powered off or its NIC is disconnected. Ephemeral binding is configurable through the ESX Host and vCenter.

VLAN ID (Optional)

Specifies an Integer value used for the virtual port group operations.

Number of Ports (Optional)

Specifies the number of ports of the port group.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Fault Tolerance: VMware vCenter

Use this action type to manage fault tolerance.

The Details section of the action definition contains the following fields:

Operation

Select one of the following operations for the specified VM:

- Turn On
- Turn Off
- Enable
- Disable
- Migrate Secondary VM

Virtual Center

Specifies the vCenter Server host name. Select one from the drop-down list.

Datacenter

Specifies the datacenter to which the VM belongs. Select one from the drop-down list.

Virtual Machine

Specifies the fault-tolerant VM. Select one from the drop-down list.

Secondary Host

Specifies the ESX server where the secondary VM resides. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage VM Snapshots: VMware vCenter

The Manage VM Snapshots action type lets you create, revert, or delete a virtual machine snapshot on a specified target system.

Note: If a Manage VM Snapshots action fails, because the ESXi host was removed from vCenter and then added back, select the corresponding snapshot again and save the action.

The Details section of the action definition contains the following fields:

Operation

Specifies one of the following actions:

- Create Snapshot
- Revert Snapshot
- Delete Snapshot

If you select Create Snapshot, complete the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select a server from the drop-down list.

Data Center

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select a data center from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine on which to create the snapshot. Select a virtual machine from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

Name

Defines a name for the virtual machine snapshot to create. You can use automatically generated text or text extracted from event messages.

Description

(Optional) Describes the virtual machine snapshot.

Capture Memory check box

(Optional) Specifies whether to create the snapshot with system running memory as part of the snapshot.

If you select Revert Snapshot, complete the following fields:

VC Server

Specifies the name of the server on which vCenter resides. Select a server from the drop-down list.

Data Center

Specifies the name of the data center in vCenter where the virtual machine resides. Select a data center from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine on which to revert the snapshot. Select a virtual machine from the drop-down list.

Name

Defines the name of the virtual machine snapshot to revert.

Enter the name or click the binocular icon and select the snapshot you want to revert to from the dialog.

ID

Defines the ID of the virtual machine snapshot to revert.

Note: You can use Name or ID to revert a snapshot, you do not need both. ID is required if you have multiple snapshots with the same name for a virtual machine.

If you select Delete Snapshot, complete the following:

VC Server

Specifies the name of the server on which vCenter resides. Select a server from the drop-down list.

Data Center

Specifies the name of the data center in vCenter where the virtual machine resides. Select a data center from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine on which to delete the snapshot. Select a virtual machine from the drop-down list.

Name

Defines the virtual machine snapshot name to delete.

Type the name or click the binocular icon and select the snapshot you want to delete from the dialog that opens.

ID

Defines the ID of the virtual machine snapshot to delete.

Note: You can use Name or ID to delete a snapshot, you do not need both. ID is required if you have multiple snapshots with the same name for a virtual machine.

Delete Children check box

(Optional) Specifies whether to delete all the children of the snapshot.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Virtual Switch: VMware vCenter

Use this action type to manage virtual switches.

The Details section of the action definition contains the following fields:

Operation

Select one of the following operations:

- Add Port Group
- Remove Port Group
- Update Port Group

Virtual Center

Specifies the vCenter Server. Select one from the drop-down list.

Datacenter

Specifies the Datacenter. Select one from the drop-down list.

ESX Server

Specifies the ESX Server to which the virtual switch belongs. Select one from the drop-down list.

Virtual Switch

Specifies the virtual switch you want to manage. Select one from the drop-down list.

Port Group

Specifies the port group name. Select one from the drop-down list.

VLAN ID (Optional)

Specifies an Integer value used for the virtual port group operations.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Manage Windows Service

The Manage Windows Service action type controls Windows services using the AutoShell command line and scripting environment.

The Details section of the action definition contains the following fields:

Operation Options

Specifies the operation performed with the service.

Note: The Windows service status that the query service operation returns can be obtained only from the %STDOUT% parameter; it is not available from the Events table. This parameter is only valid for actions running in an action sequence.

Note: The following behavior differs from service management performed directly in Windows:

- Even if a service is in the "Stopped" status, it is possible to perform the Restart Service operation. The service status will change to "Started".
- If the service is in the "Started" status, and the Disable Service operation is performed, the service will be disabled and its status will change to "Stopped".

Host Name

Defines the name of the computer on which the service is running.

User Name

Defines the user name.

Password

Defines the password. Reenter the password for confirmation.

Service Name

Defines the name of the service for which the operation is performed. Type the name or use a text extracted from event messages.

Note: Check the Service name in the Properties dialog of the service in Windows. Do not confuse it with its Display name visible in the Computer Management window.

If you select Change Service Startup Type, complete the following field:

Startup Type

Specifies the startup type that is set for the service. The options include Automatic, Manual, Disabled. The Boot option means that a device driver is loaded by the boot loader. The System option means that a device driver is started during kernel initialization.

If you select Change Service Dependencies, complete the following field:

Dependencies

Defines the dependencies (other services, system drivers, or load order groups) that must be running before the service can be started. If you define multiple dependencies, separate them by a forward slash.

If you select Change Service Account, complete the following field:

Local System Account / This Account

Specifies the account under which the service logs in. You can use the LocalSystem account or define an account here.

Migrate Machine: VMware vCenter

The vCenter VMotion Migration action type uses VMware VMotion to migrate a virtual machine. The VMware ESX servers must be configured correctly for this operation and the VMotion license must be present on the target computer.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Source Data Center

Specifies the name of the data center where the source virtual machine resides. Select one from the drop-down list.

Source Virtual Machine

Specifies the name of the server to use as the source VM. Select one from the drop-down list.

Destination ESX Server

Specifies the name of the ESX server that is to be the target of the migration. Select one from the drop-down list.

Note: VM migration between ESX hosts is only supported when the VMs datastore/disk is shared between the two ESX hosts.

Destination Resource Pool

Specifies the name of the resource pool to use. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Modify CPU: VMware vCenter

The Modify CPU action type lets you modify the number of CPUs allocated to a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to modify memory. Select one from the drop-down list.

CPU

Specifies the number of CPUs to allocate to the VM.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Modify Memory: VMware vCenter

The Modify Memory action type lets you modify memory allocation for a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to modify memory. Select one from the drop-down list.

Memory

Specifies the amount of memory to allocate to the VM.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Open HelpDesk Ticket

The Open HelpDesk Ticket action type lets you define the properties with which to open help desk tickets.

The Details section of the action definition contains the following fields:

Summary

Summarizes the ticket details.

Description

Describes the ticket.

Entity

(Optional) Defines the name of the server or service that is used to match the ticket with a known configuration item in the help desk system. If the configuration item host name is the same as the entity name, the ticket is associated with that configuration item.

Type

Specifies the type of the ticket.

Template

Specifies the template for the ticket.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Provision Machine: IBM LPAR

This action type provisions an LPAR.

The Build Partition section contains the following fields:

HMC/IVM Name

Specifies the HMC/IVM that is associated with the managed server where the selected partition resides.

System Name

Specifies the name of the data center in IBM LPAR where the virtual machine resides. Select one from the drop-down list.

Partition Name

Defines the name of the partition for image creation.

Profile Name

Defines the name of an existing profile for the selected LPAR.

The Memory Settings section contains the following fields:

Installed Memory

Identifies installed memory.

Available Memory

Identifies installed memory.

Minimum

Specifies the minimum memory.

Desired

Specifies the desired memory.

Maximum

Specifies the maximum memory.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

The Processor section contains the following fields:

Processing Mode

Specifies the processing mode.

Select between the following options:

- Partial Processor Units (Shared)
- Entire Processors (Dedicated)

Units Available

Identifies available processor units.

Minimum

Specifies the minimum processor units.

Desired

Specifies the desired processor units.

Maximum

Specifies the maximum processor units.

I/O Components

Specifies the I/O components that you want to associate with the LPAR.

I/O Pools

Lets you add, delete and modify the I/O pools.

Maximum Virtual Adapters

Defines the maximum number of virtual adapters.

Virtual Adapters

Identifies the number of virtual adapters

Virtual Serial Adapters

Lets you add, delete and modify the virtual serial adapters.

Virtual Ethernet Adapters

Lets you add, delete and modify the virtual ethernet adapters.

Virtual SCSI Adapters

Lets you add, delete and modify the virtual SCSI adapters.

The provisioning process starts on the client computer and a confirmation message notifies you when the job has completed successfully.

Provision Machine: Microsoft Hyper-V

The Provision Hyper-V VM action type creates and installs a VM. Specify the following parameters.

The Details section of the action definition contains the following fields on the first page:

SCVMM Server

Specifies the Microsoft System Center Virtual Machine Manager (SCVMM) library server. Select one from the drop-down list.

Hyper-V Server

Specifies the Hyper-V Server. Select one from the drop-down list.

Template

Specifies the template. Select one from the drop-down list.

Destination Path

Specifies the destination path of the VM that you want to create (template is stored). Select one from the drop-down list.

VM Name

Specifies the name of the VM.

Start VM

Starts the VM automatically after it is created. By default, the new VM is in powered-off state.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

The Details section contains the following fields on the second page after you click Next:

Hardware Profile

Specifies the name of the hardware profile defined by the Microsoft System Center Virtual Machine Manager (SCVMM) library server.

Virtual Processors

Specifies the number of virtual processors that you want to assign to the VM.

Default: 1

Memory

Specifies the RAM memory in megabytes (MB) for the VM that you want to create.

Default: 1024

The Details section contains the following fields on the third page after you click Next:

Guest OS Profile

(Optional) Specifies the name of the guest operating system profile defined by the Microsoft System Center Virtual Machine Manager (SCVMM) library server. This parameter overwrites the operating system configuration settings stored in the SCVMM library server. This parameter is valid when you use SCVMM integration to provision VMs.

Product Key

(Optional) Specifies the Windows product activation key for the VM. Support for this parameter requires a Windows image created using Sysprep tool. This option is invalid for asynchronous execution of command.

Full Name

Specifies the user name of the Windows image (created using sysprep tool) which is installed on the new VM.

Organization

(Optional) Specifies the organization name of the Windows image (created using sysprep tool) which is installed on the new VM. Support for this parameter requires a Windows image created using Sysprep tool. This option is invalid for asynchronous execution of command.

Admin Password

(Optional) This option is used to set the default administrator account password for the VM. Support for this parameter requires a Windows image created using Sysprep tool. This parameter is ignored in asynchronous execution.

Note: To set this option successfully, set the Windows Server administrator password which is created using Sysprep tool as empty.

Join Workgroup

Specifies the workgroup that you want to create for the VM. Domain and workgroup specifications are mutually exclusive.

Join Domain

Specifies the domain name for the VM. Domain and workgroup specifications are mutually exclusive.

Domain User

Specifies the domain user name that you want to create as a part of the default Administrators group.

Domain User Password

Specifies the password of the domain user account that you want to create as a part of the default Administrators group.

The Details section contains the following fields on the fourth page after you click Next:

Use DHCP

Specifies an option to enable DHCP for a network interface of the VM. If the template image has more than one network adapter, DHCP is turned on for the first interface. If enabled, the other network parameters are not accessible.

IP Address

Specifies the static IPv4 address that you want to assign to the VM.

Network Mask

Specifies the subnet mask that you want to assign for the VM.

Default Gateway

Specifies the default gateway for VM.

DNS Server

Specifies the DNS server that you want to set for the VM.

IP Metric

(Optional) Specifies the interface metric that you want to set for the VM. This option is used with -ip4addr option. If an interface name is specified in the -ip4addr option, same interface name must be used in this option. Support for this parameter requires a Windows image created using Sysprep tool. This option is invalid for asynchronous execution of command.

Default: 1

Provision Machine: Solaris Zones

The Provision Solaris Zone Machine action type creates and installs a zone. Specify a Solaris Zones host, a zone name, the zone type, and other zone properties. The zone installs automatically after creation.

The Details section of the action definition contains the following fields on the first page:

Host

Defines the Solaris Zones host on which to create the zone.

Name

Defines the zone name. You can use automatically generated text or text extracted from event messages.

Description

(Optional) Defines a description of the zone.

Type

Defines whether the Zone is Native, Whole Root, or Branded. A Branded Zone is based on an existing Zone template.

Template

(Optional) Defines the template from which to create the zone when you set Type to Branded.

Install Archive Path

Defines the directory path of the installation archive on the zone. This field is only required if you set Type to Branded.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

The Details section contains the following fields on the second page after you click Next:

Type

Defines the scheduler type. Select FSS to use the Fair Share Scheduling class to control CPU allocation based on the number of CPU shares assigned to tasks.

Capacity

Defines the amount of physical memory capacity to allocate to the zone, in megabytes.

Swap Memory

Defines the amount of swap memory to allocate to the zone, in megabytes. The swap memory must be at least 50 MB.

Lock Memory

Defines the amount of lock memory to allocate to the zone, in megabytes. The lock memory must be less than the physical memory.

Zone Path

Defines the root directory path of the zone.

NIC Type

Defines the NIC type. Select a type from the drop-down list. If you do not select a NIC, the zone is not assigned a NIC card or IP address.

IP Address

Defines the IP address of the zone.

Resource Pool

Defines the resource pool to use with the zone. Select a pool from the drop-down list. If you want to use a new resource pool with the zone, create the pool first. If you do not select a pool, the default is used.

Auto Reboot

Defines whether to reboot the zone automatically when the global zone is rebooted.

Provision Machine: VMware vCenter

The Provision vCenter machine action type provisions a Virtual Machine (VM). A template and a target vCenter specification that works with the template is required. If a service rule exists for provisioning a VM, that new VM is placed in the service for which the rule was created.

The Details section of the action definition contains the following fields:

VC Server

Specifies the name of the server on which vCenter resides. Select one from the drop-down list.

VC Data Center

Specifies the name of the data center on which to provision the machine. Select one from the drop-down list.

VC Compute Resource

Specifies the name of the server on which the compute resource resides. Select one from the drop-down list.

VC ESX Server

Specifies the name of the VMware ESX server that is to be the target of the provisioned VM. Select one from the drop-down list.

VC Datastore

Specifies the name of the data store to use. Select one from the drop-down list.

VC Target Location

Specifies the VC target location. Select one from the drop-down list.

Hostname/VM Name

Specifies the name or VC name to use from the specification. Select one from the drop-down list. Alternatively, you can use automatically generated text or text extracted from event messages.

User Name

Specifies the user name credentials to access the specification.

Password

Specifies the password to access the specification.

VC Virtual Machine

Specifies which available VC virtual machine to use. If selected, click one from the drop-down list.

VC Template

Specifies which available VC template to use. If clicked, select one of the software package groups that has previously been created from the drop-down list.

NICs (VC Template)

Specifies the number of network interface cards used by the VC template.

VC Specification

Specifies the name of the VC specification to use. Select one from the drop-down list.

NICs (VC Specification)

Specifies the number of network interface cards used by the VC specification.

OS System Type

Displays the type of operating system for the provisioned VM.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Memory

Specifies the amount of memory to allocate to the VM, in megabytes.

Virtual Processors

Specifies the number of virtual processors to allocate to the VM.

Datastore

(Optional) Specifies the storage datastore under which to create additional hard disks.

Drive Size

(Optional) Specifies the size of an additional hard drive.

SCSI Controller

(Optional) Specifies the SCSI controller to use to create the additional hard drive.

Network Management

Lets you change network connection settings.

Global NIC Settings

Lets you add DNS search suffixes.

Remove Disk: VMware vCenter

The Remove Disk action type lets you remove a disk from a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to add a disk. Select one from the drop-down list.

Hard Drive

Specifies the disk to be removed. Select one from the drop-down list.

Delete Disk File(s) check box

Specifies whether to delete the disk data or not.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Remove Network Interface: VMware vCenter

The Remove Network Interface action type lets you remove a virtual NIC from a virtual machine.

The Details section of the action definition contains the following fields:

Virtual Center

Specifies the name of the server on which VMware vCenter resides. Select one from the drop-down list.

Datacenter

Specifies the name of the data center in VMware vCenter where the virtual machine resides. Select one from the drop-down list.

Virtual Machine

Specifies the name of the virtual machine for which to remove the virtual NIC. Select one from the drop-down list.

Network Interface

Specifies the virtual NIC to be removed. Select one from the drop-down list.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Remove Server From Service

The Remove Server From Service action type lets you remove a server from an existing service.

The Details section of the action definition contains the following fields:

Service

Specifies the name of the service.

Server List (comma delimited)

Specifies the list of servers to remove from the service.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Run Action

The Run Action action type lets you run actions.

The Details section of the action definition contains the following fields:

Action Name

Specifies the action.

Event Source

Specifies the source of the action.

Event Message

Specifies the event message.

Rule Name

Specifies the rule for the action.

Server Name

Specifies the server for the action.

Service Name

Specifies the service for the action.

Propagate

Specifies that the action runs against all servers in the service specified in the -service_name option.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Run Action Sequence

The Run Action Sequence action type lets you select multiple actions for a rule and run them in a defined sequence.

The Details section of the action definition contains the following fields:

Restart if interrupted

Restarts the action sequence if it is interrupted. The action does not resume from the point that it was interrupted, but it starts the sequence over from the beginning.

Action Sequence

Provides single or multiple row selections of actions and conditions. When you select Action Sequence, a drop-down list is enabled. If it is not selected, text is displayed in the table cell.

Sequence

Specifies the sequence of actions.

Note: If an action sequence exits without a condition being met, the default return code is -1.

Action

Specifies the action name. You can select an action from the available actions in the drop-down list.

Condition Name

Specifies the condition that determines the next action to run. You can create your own custom conditions or use one of the following predefined conditions:

- On Failure
- On Success

Note: Conditions are evaluated in the order that they are created.

Next Step

Specifies the next action to run based on the results of the condition.

Continue

Continues to the next action when the condition evaluates to true.

Exit (RC=0)

Exits the sequence and returns a code of 0 to the log when the condition evaluates to true.

Exit w/RC (RC=action RC)

Exits the action sequence and the action's return code when the condition evaluates to true.

Abort (RC=-1)

Stops the action sequence and returns a code of -1 to the log when the condition evaluates to true.

Go to #

Continues to the action sequence number specified when the condition evaluates to true.

Add Action

Adds a new action to the table and automatically generates a new sequence number.

Add Condition

Adds a new condition to the action.

Delete

Deletes the selected row and updates the sequence numbers. A row can contain an action or a condition. This function permits you to delete a condition of an action without deleting the entire action.

Save

Saves the action sequence.

Note: If the Next Step in the last action in the sequence is set to *Continue*, the setting is automatically changed to “Exit w/RC (RC=action RC)”, and an informational message notifies you that the Next Step for the last action in the sequence has been changed.

Run Command Script

The Run Command Script action type lets you use a script to run an external command from the server that processed the command. For example, if the command is run from the Initiation page, the target command must be on the same server as the Windows scheduler that runs the command. When the command runs as a result of a rule evaluation, the action runs on the computer hosting the Windows scheduler server, as a result of running a job.

The Details section of the action definition contains the following fields:

Command Line

Specifies the command or substitution string to run. Alternatively, you can use automatically generated text or text extracted from event messages.

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Set Health State

The Set Rule Health State action type lets you set the health state of a system to any one of the following: Warning, Minor, Major, or Critical.

The Details section of the action definition contains the following fields:

Health State

Specifies one of the following actions:

- Warning
- Minor
- Major
- Critical

Require Approval

Select to specify that the ticket requires approval by a third party.

Note: CA SDM must be configured to use this option.

Auto close ticket on approval/rejection

Select to close the ticket after it is approved or rejected.

Note: CA SDM must be configured to use this option.

Ticket Types

Select a valid ticket type from the drop-down list. Depending on your configuration, valid types include:

- Default
- Incident
- Problem
- Request

Note: CA SDM must be configured to use this option.

Templates

Specifies the template to use to create a ticket. Select a template from the drop-down list. Depending on the ticket type selected, the form is populated with corresponding values.

Note: CA SDM must be configured to use this option.

Create a Custom Action

You can create customized action types by defining substitution parameters. Your custom action types are added to the Action Types drop-down list with the predefined action types.

Follow these steps:

1. In the Explore pane, select the Data Center node.
2. Click Resources, Policy, and then click the Custom Action Types tab.

The Custom Action Types page appears.

3. Click + (Add).

The Custom Action Types: Add New section appears.

4. Complete the following fields to define a new action type and a substitution parameter, and then click Save:

Action Type Name

Defines the name of the new action type.

Command

Defines the command line structure for the action type. You can define substitution parameters such as %SERVER%, \$MYKEY\$, and so on, for replacement as part of a command. Substitution keys can only be used once per command. For example, the %SERVER% substitution key can only be used once in a command.

Substitution Key

Defines a unique string for the substitution key. The substitution key name must match what is defined in the command. When defining multiple substitution keys, define each substitution key individually.

Prompt

Defines the argument name associated with the substitution parameter to input when creating actions.

Default Value

Defines the default substitution key value.

The new parameter appears in the substitution parameter list.

5. Select Save from the Actions drop-down list.

The custom action type is saved.

Define an Action Sequence

You can define action sequences for your rules. If the conditions for a rule evaluate to true, the action sequence that you defined runs. You can also create custom conditions and build them into your sequence.

Note: The action sequence can also be scheduled as a job or can be run using the `dpmpolicy runaction` CLI command.

Follow these steps:

1. In the Explore pane, select the Data Center node.
2. Click Resources, Policy, and then click the Action tab.
The Actions page appears.
3. Click + (Add new action).
The Action Definition: New section appears.
4. Type a meaningful name for the action sequence, and then select Run Action Sequence from the Type drop-down menu.
The Condition Logic section appears.
5. Leave the Restart if Interrupted check box selected to restart the sequence after an abnormal termination. The sequence restarts the last action that was executed and continues. Clear the check box to prevent the sequence from continuing after an abnormal termination.
6. Click + (Add action) in the Action Sequence pane to add actions to the action sequence. Add Action adds a new action at the end of the action sequence. If you want to insert an action in the middle of the sequence, remove all actions after the desired position. Insert the new action, and then redefine the actions that you removed.
7. Select a condition to build your condition logic for the action sequence. New condition logic can only be added to the end of the condition logic sequence. If you want to insert new condition logic in the middle of the sequence, remove all condition logic after the desired insertion point. Insert the new condition logic, and then redefine the condition logic that you removed.

8. Select the type of condition logic evaluation for each additional condition logic sequence. Output Types include the following:

ReturnCode

Evaluates the action return code.

Note: Valid comparison operators for Return Code evaluation are: ==, !=, >, <, >=, <=

STDOUT

Searches the standard output for a specific string.

STDERR

Searches the standard error for a specific string.

Note: Valid comparison operators for STDOUT and STDERR are "Contains" and "Does Not Contain".

Note: You can use the Logic OP field (AND/OR) to link conditions. Logic OP is set to NOOP automatically for the final condition.

The new condition logic is added to the sequence.

9. When you complete your conditions, click Save Condition.

The condition is saved.

10. Click Save in the Action Sequence pane.

The action is saved.

For testing purposes, you can run the action from the Actions page by selecting the action and clicking the Run action icon.

Define a Schedule

You can schedule actions to run at predefined times. For example, you can use the default Windows scheduler to schedule actions that must be performed every day, or actions that are performed periodically, such as maintenance tasks.

Follow these steps:

1. In the Explore pane, select the Data Center node.
2. Click Resources, Policy, and then click the Scheduled Actions tab.

The Scheduled Actions page appears.

3. Complete the following fields:

Name

Defines a name for the scheduled action.

Pre Notification

Specifies whether to generate an event before the scheduled action runs. The event appears in the dashboard.

Post Notification

Specifies whether to generate an event after the scheduled action runs. The event appears in the dashboard.

Frequency

Specifies how often the scheduled action runs: once, daily, weekly, monthly (day), or monthly (day of week).

Date

Defines a date on which to start the scheduled action.

Time

Defines a time of day to run the scheduled action.

Note: You do not need to enter seconds because they are not used for scheduling jobs.

Type

Specifies the action type used for the action you are scheduling.

Note: The scheduler does not support an action that contains substitution parameters (the only exceptions are %AutoIncrement(0)% and %AutoDecrement(0)%). You can run the actions that contain substitution parameters only through Policy rule evaluation.

Action

Lists the actions that have already been created for each action type.

Note: The list does not include actions that specify a help desk approval requirement.

4. Select Save from the drop-down list.

A message confirms that your action is scheduled. The scheduled action appears in the list of Scheduled Jobs in the Scheduled Actions list.

Note: Actions that specify a help desk approval requirement cannot be used for action scheduling. If you need the same action for a scheduled action, create a second action that does not include the help desk approval requirement.

Create Automation Policy

You can use the Create Automation Policy wizard to create automation rules based on two predefined policy types:

- Virtual Machine Dynamic Resource Brokering – CPU and memory allocation is dynamically changed based on defined utilization thresholds.
- Overall Utilization Metric Threshold Monitoring – Health state is set according to overall utilization.

Follow these steps:

1. Open the Manage pane, and click Create Automation Policy.

The Create Automation Policy wizard appears.

2. Select a Policy Type, and click Next to select target resources and set conditions for rules.

The Policy Summary displays the result.

3. Click Finish.

The policy is confirmed and the corresponding rules are created.

Use Cases for Policies

The following scenarios demonstrate some use cases for implementing policies.

More information:

[Use Case: Adding a Server to a Service](#) (see page 653)

[Use Case: Adding a New Rule to a Service](#) (see page 654)

[Use Case: Defining an Action](#) (see page 654)

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 656)

[Configure Data Collection for a Data Center](#) (see page 658)

[Configure Data Collection for a Server](#) (see page 659)

[Configure Data Collection for a Virtual Resource](#) (see page 661)

[Configure Performance Thresholds](#) (see page 663)

[Configure the Metric Filter](#) (see page 664)

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 Systems Performance LiteAgent or with the SystemEDGE agent on the remote computer to collect the specified system metrics. The CA Systems Performance LiteAgent 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 Systems Performance LiteAgent 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.

Install CA Systems Performance LiteAgent or the SystemEDGE agent on any server from which you want to collect the base system metrics. Alternatively, use CA NSM performance agents or SystemEDGE agent. If CA NSM performance agents or SystemEDGE agents are present, then the CA Systems Performance LiteAgent is not required. If necessary, you can install the SystemEDGE agent using the product user interface. All performance metrics are stored in the Performance DB.

- How is overall utilization calculated? Overall utilization is an aggregate calculation of all the metrics that are currently being collected for servers that are managed by CA Server Automation. The calculation is based on the value of the metrics and the user-defined thresholds that define the parameters for normal operation.

Note: Select Include for Overall Calculation in the Policy, Metrics, Thresholds section of the user interface to include new metric in the overall utilization calculation. If you include the metric, CA Server Automation provides up-to-date results when evaluating the state of the servers.

- How metric evaluations affect the overall utilization? The metric details provided in the tables help you understand how CA Server Automation evaluates the different metrics. Each metric has a method property that is set to either *exact* or *complement*. A higher exact value is a worse scenario than a lower exact value because it indicates an increase in overall utilization. A higher complement value is a positive scenario because it indicates a decrease in overall utilization. Generally, a high exact value negatively impacts overall utilization and a low exact value positively affects overall utilization. By contrast, a high complement value positively impacts overall utilization, and a low complement value negatively affects overall utilization. For example, if the value of Memory: Percentage Committed Bytes In Use increases, overall utilization of the system increases. If the value of Memory: Available MB increases, overall utilization decreases.

What are the default metrics? The default metric definitions are located in the metric list in the Filter section for all supported platforms. You can find the default metrics indicator on the metric list with the value Yes in the Default column. CA Server Automation uses this list to obtain the metric definitions when you add a server. You can configure platforms, types, subtypes, instances, and the type of data to collect in the Filter section. The metric filter and definitions for each server are stored in the Performance DB.

- Is performance data currently available for my systems? By default, if data cannot be collected, CA Server Automation does not negatively affect server state. The lack of data does not reflect server criticality. By reviewing the Events list or selecting a specific system, you can determine whether metric data is being collected. However, sometimes a more immediate means of determining if collected data is available is needed, or performance data is critical. Configure CA Server Automation 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>
```

CA Server Automation modifies the value surrounded by the value XML elements to one of the other supported values such as 5 or 10, for OK or Warning respectively. These changes reflect the desired state when performance data cannot be collected. For example:

```
<property name="CONFIG_KEY_DEFAULT_HEALTH_STATE">
    <!-- Valid values: 0 (Unknown); 5 (OK); 10 (Warning); 15 (Minor Failure);
20 (Major Failure); 25 (CriticalFailure) -->
    <!-- Changes the value of HealthState for the CA_CollectionState object
associated to the CA_ComputerSystem -->
    <!-- If set to 30, CE will not set the HealthState. -->
    <value>10</value>
    <displayName>Default node health state when problem encountered in metric
or data collection</displayName>
</property>
```

Because `<value>` was changed to "10", systems that do not have performance data available are displayed in a warning state in the CA 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 PowerVM Server
- IBM Logical Partition
- Solaris Zones Server
- Solaris Zone

To configure data collection for a virtual resource

1. Click Resources, and open the Explore pane.
2. Expand the Data Center or MS Cluster Service folder, then any subfolder, and select the object that you want to configure.

Subtabs for that object appear in the right pane.

Note: If you select a top-level folder (such as VMware vCenter Server) or an object for which no data is collected (such as a vCenter cluster), you must select the specific object contained within the folder or object for which to configure data collection.

Note: If you select MS Cluster Service as the top-level folder, then you see clusters and their nodes.

3. Right-click and select Policy, Configure Server Metrics Collection.

Note: If you select the top-level folder for Solaris Zones, the Hardware Class column in the System section always shows the value Other.

4. Select the metrics that you want to monitor from the Available Metrics section and then click the down arrow.

The metrics you select move to the Selected Metrics to Collect section.

Note: If you disable the default metrics (CPU and memory) and enable others, you will not see an overall utilization until you modify the thresholds of the newly selected metrics.

5. Click Save to apply the selected metrics.
6. Right-click the resource, and select Policy, Configure Collection Settings.
7. Complete the following fields in the Collection Setting section:

Use Default

Specifies the data center level as the default when selected. If you leave the check box cleared, the values that you specify are used instead.

Data recording interval (seconds)

Defines how often the data is collected and stored in the Performance DB.

Default: 300 seconds

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 nth 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.

Chapter 11: Provisioning Resources

This section contains the following topics:

[Imaging Services](#) (see page 667)
[Service Provisioning](#) (see page 668)
[CA Software Delivery](#) (see page 711)
[Amazon EC2 Provisioning](#) (see page 719)
[Bare Metal Provisioning to a Cisco UCS Blade](#) (see page 728)
[LPAR Provisioning for IBM AIX](#) (see page 728)
[IBM AIX Provisioning with NIM](#) (see page 729)
[Rapid Server Imaging](#) (see page 740)

Imaging Services

CA Server Automation can provision new physical and virtual computers and also reimagine 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.
- Hyper-V integration for VM provisioning that is based on local templates on Hyper-V servers out of the box.
 - Integration with Microsoft System Center Virtual Machine Manager (SCVMM) for reuse of existing SCVMM image libraries.
- Solaris Zones integration for zones provisioning.
- Citrix XenServer integration for VM provisioning.
- VMware vCloud Director integration for VM provisioning.
- IBM AIX Provisioning with NIM
- 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.

Service Provisioning

This section contains the following topics:

[How to Provision Services](#) (see page 668)

[How to Deploy a Wiki Web Page](#) (see page 690)

[How to Deploy Oracle WebLogic Server](#) (see page 704)

How to Provision Services

As a *Service Administrator*, you use CA Server Automation to manage and provision services in your physical and virtual server environments. You want to provide *Service Consumers* with a simple way to provision the applications they require, hosted on the machines that are required to run them.

For Service Administrators, service provisioning consists of defining applications and service templates. Service templates consist of a set of applications, deployed in a defined order, with a set of application actions and machine templates that determine how and where to deploy the service.

Service templates provide Service Consumers with one-click service provisioning. The consumer selects the service template to provision and enters any additional information that is required to deploy the template as a running service.

CA Server Automation deploys an instance of the service to provisioned resources which match the requirements in the service template.

For example, for a two-tier web application and database service:

Service Administrator

1. Defines the web application and the database application, and specifies the execution actions and resources that are required to install and run them.
2. Creates a service template with the web and database applications and configures the resources that are required to host them.

Service Consumer

1. Provisions an instance of the service template.

CA Server Automation deploys the applications to the available server resources and informs the Service Consumer that the web service is ready for use at a specified location.


CA Server Automation provides a set of ready-to-use application definitions and service templates, including the following service:

- Apache HTTP Server for Windows

CA Server Automation supports service provisioning to the following environments:

- VMware vCenter
- CA AppLogic

The following process describes a high-level overview of how Service Administrators combine applications to create service templates that enable one-click service provisioning for Service Consumers:



1. [Configure Resources for Service Provisioning](#) (see page 672)

Specify the environment and server resources available to host provisioned services.

- (Optional) [Configure Machine Templates for Service Provisioning](#) (see page 673)

Specify default machine configurations to use in service templates.

- (Optional) [Configure Dynamic Specifications for vCenter](#) (see page 673)

Specify and create custom licensing and domain configurations to apply to virtual machines in service templates.

2. [Define Applications](#) (see page 674)

Specify applications, any prerequisites and restrictions, and the resource and operating system requirements for the applications.

Define any of the following types of actions to specify the application execution:

- [Execute Application Installation](#) (see page 676)

Specify an application file from the Staging Directory and any execution action options that are required to deploy it.

- [Execute Commands](#) (see page 677)

Specify any additional commands that are required to enable the successful deployment of an application.

- [Create and Update Files](#) (see page 679)

Create or modify any configuration or property files that are required for the successful deployment of an application. For example, silent installation response files.

- [Deploy CA ITCM Software Packages or Groups](#) (see page 681)

Deploy CA ITCM software packages, either as standalone application definitions, or as part of the deployment of another application.

- [Execute CA Process Automation Processes](#) (see page 682)

Execute a process workflow defined in CA Process Automation, either as a standalone application definition, or as part of the deployment of another application.

3. [Create Service Templates](#) (see page 683)

Specify a set of applications and other service templates, and the machines that are required to host them, as a deployable service template.

4. (Optional) [Share Application Definitions and Service Templates](#) (see page 688)

Export application definitions and service templates to files, and then import them to other instances of CA Server Automation.

5. [Capture Services from vCenter](#) (see page 688)

Create service templates and application definitions from service instances running in VMware vCenter.

6. [Import Templates from AppLogic](#) (see page 689)

Create service templates from AppLogic templates.

7. [Provision Services](#) (see page 690)

Deploy an instance of a service to your server environment.

Configure Resources for Service Provisioning

To enable CA Server Automation to provision services, specify the environment resources available to host services.

CA Server Automation supports service provisioning to the following environments:

- VMware vCenter
- CA AppLogic

Note: CA Server Automation uses default OnDemand resource pools for service provisioning. To configure these resource pool or to specify other resource pools for service provisioning, use the administration features of Reservation Manager.

Follow these steps:

1. Click Resources, and in the Explore tree, select OnDemand Services.
2. In the OnDemand Services toolbar, click + (Add VMware vCenter Server).
3. Select the vCenter Server or AppLogic Server to use and click OK.

The dialog adds the vCenter or AppLogic server and its respective ESX host servers or AppLogic grids to the list of environment resources available for service provisioning.

4. (Optional) Select ESX host servers or AppLogic grids and click - (Delete) to make them unavailable for service provisioning.
5. Click Actions, Save.

CA Server Automation updates the available resources for service provisioning.

More information:

[How to Configure the vCenter Server Management Components](#) (see page 278)

Configure Machine Templates for Service Provisioning

CA Server Automation enables you to select VM templates for specified operating systems. Service templates can then use these templates to automate the provisioning of the machines that are required during service provisioning.

Follow these steps:

1. Click Resources, and in the Explore tree, select OnDemand Services.
2. In the OnDemand Services page, click the tool icon for the vCenter Server containing the templates to use.

The Machine Template Configuration panel opens.

3. Select an operating system from the list, select a Template from the list of available VM templates, and click Set As Default. Repeat this step for each operating system that requires VM templates.

CA Server Automation associates the selected VM template with the specified operating system.

4. (Optional) Click an operating system group folder to view and manage the list of all VM template settings for that OS family.
5. Click OK to exit Machine Template Configuration.

Configure Dynamic Specifications for vCenter

VMware vCenter uses custom specifications to store licensing and domain information settings to apply to virtual machines. CA Server Automation enables you to specify existing custom specifications or to create new specifications to apply to the machines that are used in service templates.

Follow these steps:

1. Click Administration, and in the Configuration, Provisioning menu, select Service Provisioning.
2. Click + (New) in the Dynamic Custom Specification toolbar.
3. Specify a Name, the Operating System it applies to, and indicate whether this specification is the Default Specification.
4. To use an existing VMware specification or create one, take one of the following actions:
 - Select an existing VMware Customization Specification from the drop-down list, and click Finish.
 - Click Next, specify the owner, licensing, domain, and other details for a new specification, and click Finish.

The specification is added to the list of specifications available to apply to machines in service templates.

Define Applications

The initial step in service provisioning defines the set of applications available to construct service templates and the actions that are required to execute them.

Note: CA Server Automation uses a default Staging directory in its installation folder to store executable files for service provisioning.

Follow these steps:

1. Copy any required application files to a new folder in the Staging directory.

Important! Licensing restrictions prohibit CA Technologies from providing licensed application files. Provide your own licensed versions.

Note: For some use cases, such as CA ITCM package deployment, application files are not required.

2. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Application.

The Define an Application wizard opens in the Application Details panel.

3. Specify the Name, Description, Version, and Vendor for the application.
4. (Optional) Specify a semi-colon separated list of Tags to enable you to organize applications into user-defined groups in the Explore tree.
5. Specify the File Location which is the folder in the Staging Directory containing the executable files for the application.
6. (Optional) Specify any prerequisite applications that must also execute to enable the successful deployment of this application.

Note: Define prerequisite applications using this procedure.

7. (Optional) Specify any restrictions for the application. For example, one instance per server or one instance per service.
8. Click Next.

The System Requirements panel opens.

9. Specify the resource requirements for the application, the supported operating systems, and click Next.

The Installation Actions panel opens.

10. (Optional) Click + to define the actions that are required to deploy the application during service provisioning.

The Define an Action wizard opens.

Use the Action input and refer to the referenced procedure to specify any of the following actions:

[Execute Application Installation](#) (see page 676)

Specify an application file from the Staging Directory and any execution action options that are required to deploy it.

Note: Specifying an application file may not be required in the following use cases:

- You want to deploy CA ITCM software packages.
- The application is a captured VM template that already contains a pre-installed application.

[Execute Commands](#) (see page 677)

Specify any additional commands that are required to enable the successful deployment of an application.

[Create and Update Files](#) (see page 679)

Create or modify any configuration or property files that are required for the successful deployment of an application.

[Deploy CA ITCM Software Packages or Groups](#) (see page 681)

Deploy CA ITCM software packages, either as standalone application definitions, or as part of the deployment of another application.

[Execute CA Process Automation Processes](#) (see page 682)

Execute a process workflow defined in CA Process Automation, either as a standalone application definition, or as part of the deployment of another application.

Note: Only applications with defined actions execute during provisioning. The referenced procedures assume that you skip setting actions during application definition and modify the definition later.

11. (Optional) If CA Server Automation is integrated with CA Configuration Automation, click Next and specify any Blueprints to associate with the application. Blueprints provide benchmark software deployments to compare against current deployments.

Note: For more information about blueprints, see the CA Configuration Automation documentation.

12. Click Finish.

The wizard adds the application to the list in the Applications pane. The application is now available for use in service template creation.

Execute Application Installation

To run an application during service provisioning, specify the application file to execute and specify action options to define how to execute the application.

Follow these steps:

1. Click Resources, and in the Explore tree, select Applications. In the Applications pane, click the tool icon for the application to Edit.
2. Select the Configure Installation Actions tab.
3. Click + (Add).

The Define Action wizard opens in the Select Installation Action panel.

Note: If you are adding actions during [Define Applications](#) (see page 674), start from this step.

4. Select Execute Program from the Action drop-down list.
5. Select the Program Name to execute from the drop-down list. The File Location that is specified for the application determines the available executable files.

Note: You can directly input the command and options that are required to execute the application file in the Program Name field. However, this option does not enable you to modify the application options.

6. (Optional) Click Username to specify user credentials to use when executing the action.

Note: This option enables user-restricted actions during service provisioning. The primary use case is to enable clustered application deployment using domain user credentials.

7. Add a description, modify the supported operating systems for the specified executable application as required, and click Next.

The Define Installation Action Options panel opens with the Action Preview displaying the application command to execute when the application is deployed during service provisioning.

Note: Specifying supported operating systems for actions as well as applications enables you to define separate OS-specific actions. An action only executes when an application deploys to the specified operating systems for the action. For example, define an application that supports Windows and Linux, and specify separate .bat and .sh actions for Windows and Linux respectively.

8. (Optional) Click the Add icon.
 - a. Specify an Action Parameter to use to modify the execution of the application.
 - b. Specify a Description, Data Type, and Default Value for the action parameter.
 - c. (Optional) Select User Editable to indicate that the end user can input a value for the action parameter during the execution of the application. Specify a Label as the prompt the end user receives when an input is required.
 - d. (Optional) Select Mandatory to specify that a user input is required.
 - e. Click Finish.

The wizard adds the action option to the Options list in the Define Action Options panel and updates the Action Preview.

Note: Click the Parse icon to input an entire executable command, with all its options. The wizard parses the command and creates each individual action option. Edit each action option individually before continuing.

9. Click Finish.

The wizard adds the Action to the Action list in the Installation Actions panel.

10. (Optional) Repeat steps 3 through 9 to add additional actions to apply during application execution. Use the up and down arrows to specify the order in which to execute actions.
11. Click OK.

Note: If you are adding actions during [Define Applications](#) (see page 674), click Finish.

CA Server Automation modifies the application definition to execute the defined actions when it is deployed during service provisioning.

Execute Commands

In some cases, the successful deployment of an application during service provisioning requires additional actions. For example, another service must stop before an application can deploy, and then restart after deployment. The application definition enables you to specify direct command line inputs to perform these additional actions.

Follow these steps:

1. Click Resources, and in the Explore tree, select Applications. In the Applications pane, click the tool icon for the application to Edit.
2. Select the Installation Actions tab.

3. Click + (Add).

The Define Action wizard opens in the Select Installation Action panel.

Note: If you are adding actions during [Define Applications](#) (see page 674), start from this step.

4. Select Execute Program from the Action drop-down list.
5. Input the executable command in the Program File field.

Note: If you require user definable inputs and configurable options, only input the basic command.

6. (Optional) Click Username to specify user credentials to use when executing the action.

Note: This option enables user-restricted actions during service provisioning. The primary use case is to enable clustered application deployment using domain user credentials.

7. Add a description, modify the supported operating systems for the specified executable command as required, and click Next.

The Define Installation Action Options panel opens with the Action Preview displaying the command to execute when the application is deployed during service provisioning.

Note: Specifying supported operating systems for actions as well as applications enables you to define separate OS-specific actions. An action only executes when an application deploys to the specified operating systems for the action. For example, define an application that supports Windows and Linux, and specify separate .bat and .sh actions for Windows and Linux respectively.

8. (Optional) Click the Add icon.
 - a. Specify an Action Parameter to use to modify the execution of the command.
 - b. Specify a Description, Data Type, and Default Value for the action parameter.
 - c. (Optional) Select User Editable to indicate that the end user can input a value for the action parameter during the execution of the command. Specify a Label as the prompt the end user receives when an input is required.
 - d. (Optional) Select Mandatory to specify that a user input is required.
 - e. Click Finish.

The wizard adds the action option to the Options list in the Define Action Options panel and updates the Action Preview.

Note: Click the Parse icon to input an entire executable command, with all its options. The wizard parses the command and creates each individual action option. Edit each action option individually before continuing.

9. Click Finish.

The wizard adds the Action to the Action list in the Installation Actions panel.

10. (Optional) Repeat steps 3 through 9 to add additional actions to apply during application execution. Use the up and down arrows to specify the order in which to execute actions.
11. Click OK.

Note: If you are adding actions during [Define Applications](#) (see page 674), click Finish.

CA Server Automation modifies the application definition to execute the defined actions when it is deployed during service provisioning.

Create and Update Files

Many applications reference configuration files when they execute. The application definition enables you to create and update files as part of the definition.

Follow these steps:

1. Click Resources, and in the Explore tree, select Applications. In the Applications pane, click the tool icon for the application to Edit.
2. Select the Installation Actions tab.
3. Click + (Add).

The Define Action wizard opens in the Select Installation Action panel.

Note: If you are adding actions during [Define Applications](#) (see page 674), start from this step.

4. Select Create File or Update File from the Action drop-down list.
5. Input the file name, and for the Create File option specify the File Format.
6. (Optional) Click Username to specify user credentials to use when executing the action.

Note: This option enables user-restricted actions during service provisioning. The primary use case is to enable clustered application deployment using domain user credentials.

7. Add a description, modify the supported operating systems for the specified file as required, and click Next.

The Define Installation Action Options panel opens.

Note: Specifying supported operating systems for actions as well as applications enables you to define separate OS-specific actions. An action only executes when an application deploys to the specified operating systems for the action. For example, define an application that supports Windows and Linux, and specify separate .bat and .sh actions for Windows and Linux respectively.

8. To add parameters to a file for the create file action, click the Add icon:
 - a. Specify an Action Parameter to add to the file.
 - b. Specify a Default Value for the action parameter.

Note: Use the Value List to maintain a set of values for the parameter.
 - c. (Optional) Select User Editable to indicate that the end user can input a value for the action parameter during the execution of the command. Specify a Label as the prompt the end user receives when an input is required.
 - d. (Optional) Select Mandatory to specify that a user input is required.
 - e. Click Finish.

The wizard adds the parameter and value to the properties file. Repeat this step for each line in the file.

Example: To Add an Editable Password to a Key-Value Properties File

- Specify the Action Parameter *password*
- Specify the Default Value *changeit*
- Select User Editable, and specify Label *Input Password*.

CA Server Automation adds *password=changeit* to the properties file. When users provision the service, CA Server Automation prompts them to input a password.

9. To edit parameters in a file for the update file action, click the Add icon:
 - a. Specify the text to replace in the Action Parameter input.
 - b. Specify the new text in the Default Value input.

Note: Use the Value List to maintain a set of values for the parameter.
 - c. (Optional) Select User Editable to indicate that the end user can input a value for the action parameter during the execution of the command. Specify a Label as the prompt the end user receives when an input is required.
 - d. (Optional) Select Mandatory to specify that a user input is required..
 - e. Click Finish.

The wizard replaces the specified text in the file with the new values. Repeat this step for each text string to replace.

Example: To Specify a Wiki Name in a Generic Wiki Properties File

- Specify the Action Parameter *###WIKINAME###*
- Specify the name of your Wiki site as the Default Value

CA Server Automation replaces instances of *###WIKINAME###* in the properties file with the name of your Wiki site.

10. Click Finish.

The wizard adds the Action to the Action list in the Installation Actions panel.

11. (Optional) Repeat steps 3 through 10 to add additional actions to apply during application execution. Use the up and down arrows to specify the order in which to execute actions.
12. Click OK.

Note: If you are adding actions during [Define Applications](#) (see page 674), click Finish.

CA Server Automation modifies the application definition to execute the defined actions when it is deployed during service provisioning.

Deploy CA ITCM Software Packages or Groups

CA Server Automation integrates with CA ITCM to enable you to deploy software packages and package groups during service provisioning. You can define an application to deploy a software package, or add software package deployment as an additional action to apply as part of the definition of another application.

Follow these steps:

1. Click Resources, and in the Explore tree, select Applications. In the Applications pane, click the tool icon for the application to Edit.
2. Select the Installation Actions tab.
3. Click + (Add).

The Define Action wizard opens in the Select Installation Action panel.

Note: If you are adding actions during [Define Applications](#) (see page 674), start from this step.

4. Select ITCM Package or ITCM Package Group from the Action drop-down list.
5. For Packages, specify the ITCM Server, the Package, and the Procedure to use to deploy the package.

For Package Groups, specify the Package Group to use.

Note: Only Selected Packages and Package Groups that are set up for use in CA Server Automation are available for selection.

6. Add a description, modify the supported operating systems for the specified package or group as required, and click Finish.

The wizard adds the Action to the Action list in the Installation Actions panel.

Note: Specifying supported operating systems for actions as well as applications enables you to define separate OS-specific actions. An action only executes when an application deploys to the specified operating systems for the action. For example, define an application that supports Windows and Linux, and specify separate .bat and .sh actions for Windows and Linux respectively.

7. (Optional) Repeat steps 3 through 6 to add additional actions to apply during application execution. Use the up and down arrows to specify the order in which to execute actions.
8. Click OK.

Note: If you are adding actions during [Define Applications](#) (see page 674), click Finish.

CA Server Automation modifies the application definition to execute the defined actions when it is deployed during service provisioning.

Execute CA Process Automation Processes

CA Process Automation enables you to combine series and parallel operations into complex workflows. CA Server Automation integration with CA Process Automation enables you to specify process workflows as part of an application definition.

Note: For more information about process workflows, see the CA Process Automation documentation.

Follow these steps:

1. Click Resources, and in the Explore tree, select Applications. In the Applications pane, click the tool icon for the application to Edit.
2. Select the Installation Actions tab.
3. Click + (Add).

The Define Action wizard opens in the Select Installation Action panel.

Note: If you are adding actions during [Define Applications](#) (see page 674), start from this step.

4. Select Process Automation Process from the Action drop-down list.
5. Specify the Start Request Form for the process from the drop-down list of available processes.

Note: Click Open Process to view and edit the process workflow in CA Process Automation.

6. Add a description, modify the supported operating systems for the specified process as required, and click Next.

The Define Installation Action Options panel opens showing each defined step in the process as an action option.

Note: Specifying supported operating systems for actions as well as applications enables you to define separate OS-specific actions. An action only executes when an application deploys to the specified operating systems for the action. For example, define an application that supports Windows and Linux, and specify separate .bat and .sh actions for Windows and Linux respectively.

7. (Optional) Click the tool icons to edit actions, and use the up and down arrows to modify their execution order.

8. Click Finish.

The wizard adds the Action to the Action list in the Installation Actions panel.

9. (Optional) Repeat steps 3 through 8 to add additional actions to apply during application execution. Use the up and down arrows to specify the order in which to execute actions.

10. Click OK.

Note: If you are adding actions during [Define Applications](#) (see page 674), click Finish.

CA Server Automation modifies the application definition to execute the defined actions when it is deployed during service provisioning.

Create Service Templates

Service provisioning is the creation of a working instance of a service that is based on a service template. A service template is the set of applications it consists of, with their associated actions, and the machine definitions that are required to host the service.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Service Template.

The Create Service Template panel opens.

2. Input a name and description for the service template and select the number of machines that are required to host the service.

3. (Optional) Specify a Service Information File to attach an endpoint template file to the service template.

The Service Information File is the panel that appears to the user when a service is successfully provisioned. The HTML format file specifies the connection details required to access the service and any other information that is required to use the service.

Note: For examples of service information files, review the service templates that are provided with CA Server Automation.

4. (Optional) Specify a semi-colon separated list of Tags to enable you to organize service templates into user-defined groups in the Explore tree.
5. (Optional) Select Allow Additional Applications to let the end user add dependent applications when they provision a service.

Note: Use this option to create base templates which the end user can then build on by adding specific dependent applications. For example, create a template that provisions an application server with supporting infrastructure and allows the end user to specify any additional dependent applications.

6. (Optional) Edit the configuration and requirements for the machines that are required to provision the service:

Note: By default, service templates automatically apply VM templates to any machines required for service provisioning. [Configure Machine Templates for Service Provisioning](#) (see page 673) to specify any default VM templates to use for specified operating systems. If there are no default templates and you do not select one, CA Server Automation requests an appropriate template from the vCenter VMware server.

- a. Click Details.

The Target Machine Details panel opens displaying the default configuration for each machine.

- b. Click the tool icon to edit the configuration for a single machine.

Note: Select multiple machines and click the right-arrows to edit multiple machine configurations.

The Machine panel opens displaying the configuration and resource settings for the machine.

- c. (Optional) Specify a VMware vCenter Template to use.

- d. Specify the CPU and memory resources that are required for the machine and whether the machine is a cluster of machines.

Note: For a cluster, the machine configuration settings apply to each machine in the cluster.

- e. (Optional) Specify or create a Dynamic Specification to use for VMware vCenter provisioning.

- f. Click the Disks tab to view the view and modify the disk storage configuration for a machine.

Note: If you are modifying multiple machines, adding disks creates shared storage for the selected machines.

- g. Click the NICs tab to view and modify the network configuration for a machine.

- h. Click OK

- i. Repeat steps b through h for each machine as required, and click Close to return to the Service Template panel.

7. (Optional) If more than one application requires the same parameters, you can use global substitution variables instead of setting the parameter in each individual application. For example, if multiple applications require the same database connection and credentials, you can create substitution variables specifying the default values, and then specify the substitution variable in the individual applications.

In the Options tab, click the Add icon to add global substitution variables:

- a. Enter the substitution variable in the Action Parameter input. Use # to start and end the variable.

Example: #DB_URL#

- b. Enter a Description, Label, and Data Type for the variable.
- c. Enter a Default Value for the variable.

Note: Click Value List to maintain a set of values for the variable.

- d. (Optional) Select User Editable to indicate that the end user can input a value for the action parameter during the execution of the application. Specify a Label as the prompt the end user receives when an input is required.
- e. (Optional) Select Mandatory to specify that a user input is required.
- f. Click OK.

CA Server Automation adds the variable to the list of Options for the service template.

- g. Edit application action options in the service template as required, replacing the Default Value for an action parameter with the substitution variable.

Note: Use the Options tab to modify configuration parameters for templates imported from AppLogic.

8. To add applications to the service template:
 - a. Click + in the Applications and Service Templates tab, select Application, and click OK.
 - b. Specify the machine to host the application.
 - c. Select the application from the list of available defined applications and click Next.

The Installation Actions panel opens.

- d. (Optional) Click the tool icon to edit actions for the application.

Note: These action modifications only apply to the service template. The application definition is not modified.

- e. Click Next.

The Confirm Application Configurations panel opens.

- f. Click Finish.

The Create Service Template panel updates with the details of the application.

9. To provision other service templates as part of this service template:
 - a. Click + in the Applications and Service Templates toolbar, select Application, and click OK.
 - b. Select the service template from the list of available templates and click OK.

The Create Service Template panel updates with the details of the service template.

10. Repeat steps 8 and 9 to add additional applications and service templates. Use the up and down arrows to specify the order in which to execute them.

Note: To deploy multiple instances of an application to the same or different machines, repeat the addition of the application for each machine.

11. Click OK.

CA Server Automation adds the template to the list of service templates in the Explore tree. The service template is ready for provisioning.

Note: CA Server Automation automatically creates a reservation template for use in Reservation Manager based on the service template.

Share Application Definitions and Service Templates

CA Server Automation enables you to share application definitions and service templates across instances of CA Server Automation.

To export application definitions or service templates:

1. Click Resources, and in the Explore tree, select Applications or Service Templates.
2. Select the applications or templates to export and click the Export icon.
3. Click Save, specify a file name and location for the export file, and click Save.

CA Server Automation creates an XML file containing the application definitions or service templates.

To import application definitions or service templates:

1. Click Resources, and in the Explore tree, right-click Applications or Service Templates, and select Import.
2. Specify the XML file containing the application definitions or service templates and click Import.

CA Server Automation creates the application definitions and service templates that the file specifies.

3. Add the appropriate executable files for new application definitions to your Staging directory.
4. Edit the application definitions and service templates as required to enable appropriate provisioning to your environment.

The new application definitions and service templates are ready to use for service provisioning.

Capture Services from vCenter

CA Server Automation enables you to capture an existing service, running on VMware vCenter, and create a service template and the applications it consists of.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click a service node under Data Center, and select Management, Capture Service.
The Capture Service to Service Template panel opens.
2. Input a name and description for the service template.
3. Use the up and down arrows to specify the order of application execution for the applications in the template.

4. Click Change to modify the application name and template name to use for each template.
5. Click OK.

CA Server Automation stops the machines and captures them as applications and creates a service template that is based on the service.

Each captured application consists of the VM template for the machine and the software it contains. Modify the applications and the service template as required to enable their deployment to different machines.

Note: CA Server Automation automatically creates a reservation template for use in Reservation Manager based on the service template.

Import Templates from CA AppLogic

CA Server Automation enables you to import AppLogic application templates, and create equivalent service templates.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click an AppLogic grid, and select Management, Import AppLogic Template to Service Template.
2. Select the template to import, optionally edit the Service Template Prefix, and click OK.

CA Server Automation imports the AppLogic template and creates a service template, named with the prefix and template name.

Edit the service template as required, using the Options tab to configure the AppLogic properties of the service.

Note: CA Server Automation automatically creates a reservation template for use in Reservation Manager based on the service template.

Provision Services

The creation of service templates enables one-click service provisioning for *Service Consumers*.

Note: This procedure provisions the service immediately. If you want to provision the service at a specified time, use Reservation Manager to make a reservation using the reservation template for the service.

Follow these steps:

1. Click Resources, and in the Explore tree, select Service Templates. In the Service Templates pane, click the Provision Service icon for the service to provision.
2. Input a name for the provisioned service instance.
3. (Optional) Click Additional Applications and select from the available applications to deploy them with the service template.

Note: Additional applications deploy and execute according to their application definition. Only applications that are dependent on applications already in the service template are available for selection.

4. (Optional) Provide any additional inputs that the application definitions in the service template specify.
5. Click OK.


CA Server Automation provisions the requested service instance and adds it to the list of resources under OnDemand Resources in the Explore tree. Track the status of the provisioning in the Jobs pane.

How to Deploy a Wiki Web Page

This example shows how to prepare a template that enables end users to deploy a MediaWiki web page. First, you define application executables and their execution actions. Then you combine these applications in a service template to define their execution order and the machines that are required to host them.

The following process describes a high-level overview of how System Administrators combine applications to create service templates that enable one-click service provisioning for Service Consumers:

How to Deploy a Wiki Web Page



1. [Configure Machine Templates for Wiki](#) (see page 692)
Specify default machine configurations to use in service templates.
2. [Define Applications for Wiki](#) (see page 692)
Specify executable applications, prerequisites, and the resource and operating system requirements for the applications.
3. [Create Wiki Service Template](#) (see page 703)
Specify the set of applications, and the machines that are required to host them, as a deployable service template.
4. [Deploy a Wiki](#) (see page 704)
Deploy an instance of the service to your server environment.

Configure Machine Templates for Wiki

CA Server Automation enables you to specify VM templates for specified operating systems. These templates can then be automatically applied to machines required during service provisioning.

Follow these steps:

1. Click Resources, and in the Explore tree, select OnDemand Services.
2. On the right-hand panel, click the tool icon for the vCenter Server containing the templates that you want to use.

The Machine Template Configuration panel opens.

3. Select an operating system from the list, select a Template from the list of available VM templates, and click Set As Default.

CA Server Automation associates the selected VM template with the specified operating system.

4. (Optional) Click an operating system group folder to view and manage the list of all VM template settings for that OS family.
5. Click OK to exit Machine Template Configuration.

Define Applications for Wiki

The initial step in service provisioning is to define the set of applications available to construct a service template and the actions that are required to execute them.

Follow these steps:

1. In the Staging directory in the CA Server Automation installation directory, create a folder for each of the following applications:
 - Apache HTTP Server
 - PHP
 - MySQL
 - Wiki Database
 - MediaWiki Content
2. Copy all the application files to the respective folders.
3. Define new applications in CA Server Automation. See the following chapters for details.

Define Apache HTTP Server

Apache HTTP Server is the host on which the wiki will be running. You install the web server as a service.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Application.

The Define an Application wizard opens in the Application Details panel.

2. Specify the Name, Version, and Vendor for the application.
3. Specify the location of the application files.
4. Click Next.

The System Requirements panel opens.

5. Set the following recommended values:

- Number of CPUs – 1
- Memory – 512 MB
- Disk Space – 5 GB
- Operating System – All Microsoft Windows Server versions

6. Click Next.

The Configure Installation Actions panel opens. You will define four installation actions: Disable Windows Firewall, Disable Windows Firewall on Windows 2003, Install Apache Server, and Reboot.

7. Click + to define the Disable Windows Firewall action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
netsh advfirewall set currentprofile state off
```

- b. Click Next.
- c. On the Define Installation Action Options panel, only check the preview and click Finish.

The action is saved.

8. Click + to define the Disable Windows Firewall on Windows 2003 action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
netsh advfirewall set opmode disable & if "%errorlevel%"=="1" exit /b 0
```

- b. Click Next.
- c. On the Define Installation Action Options panel, only check the preview and click Finish.

The action is saved.

9. Click + to define the Install Apache Server action.

- a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
cmd.exe /c start /w msixec.exe
```

- b. Click Next.

- c. On the Define Installation Action Options panel, add the following four options:

```
/i %CD%\httpd-2.2.22-win32-x86-openssl-0.9.8t.msi
```

```
INSTALLDIR=C:\Apache
```

```
SERVERADMIN=admin@localhost.com
```

```
SERVERNAME=%LOCALHOST%
```

```
AgreeToLicense=1
```

```
ALLUSERS=1
```

```
RebootYESNo=No
```

- d. Check the action preview and click Finish.

The action is saved.

10. Click + to define the Reboot action.

- a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
cmd.exe /c shutdown -r -t 20 & exit /b 1641
```

- b. Click Next.

- c. On the Define Installation Action Options panel, only check the preview and click Finish.

The action is saved.

11. Click Next, and then Finish.

The Apache HTTP Server application is saved.

Define PHP

PHP is a scripting language that is used for creating dynamic web pages, such as wiki.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Application.

The Define an Application wizard opens in the Application Details panel.

2. Specify the Name, Version, and Vendor for the application.
3. Specify the location of the application files (absolute path or relative to the Staging folder).
4. Specify Apache HTTP Server as a prerequisite for this application.
5. Click Next.

The System Requirements panel opens.

6. Set the following recommended values:
 - Number of CPUs – 1
 - Memory – 512 MB
 - Disk Space – 5 GB
 - Operating System – All Microsoft Windows Server versions
7. Click Next.

The Configure Installation Actions panel opens. You will define three installation actions: Install PHP, Stop Apache, and Start Apache.

8. Click + to define the Install PHP action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:
`cmd.exe /c start /w msixexec.exe`
 - b. Enter a description for the action.
 - c. Click Next.
 - d. On the Define Installation Action Options panel, add the following five options:
`/i %CD%\php-5.3.13-Win32-VC9-x86.msi`
`APACHEDIR=C:\Apache\conf`
`/qn`
`AgreeToLicense=YES`
`ADDLOCAL=ext_php_mysql,ext_php_mysqli,apache22`
 - e. Click Finish.

The action is saved.

9. Click + to define the Stop Apache action.

- a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
net stop apache2.2
```

- b. Click Next.

- c. On the Define Installation Action Options panel, only check the preview and click Finish.

The action is saved.

10. Click + to define the Start Apache action.

- a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
net start apache2.2
```

- b. On the Define Installation Action Options panel, only check the preview and click Finish.

The action is saved.

11. Click Next, and then Finish.

The PHP application is saved.

Define MySQL

MySQL is the server on which your wiki database will be running.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Application.

The Define an Application wizard opens in the Application Details panel.

2. Specify the Name, Version, and Vendor for the application.
3. Specify the location of the application files (absolute path or relative to the Staging folder).
4. Click Next.

The System Requirements panel opens.

5. Set the following recommended values:
 - Number of CPUs – 1
 - Memory – 512 MB
 - Disk Space – 5 GB
 - Operating System – All Microsoft Windows Server versions

6. Click Next.

The Configure Installation Actions panel opens. You will define three installation actions: Install MySQL, Configure MySQL, and Grant Permissions to Root.

7. Click + to define the Install MySQL action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
msiexec
```

- b. Enter a description for the action.
- c. Click Next.

- d. On the Define Installation Action Options panel, add the following three options:

```
/i %CD%\mysql-5.5.25-winx64.msi
```

```
/passive
```

```
INSTALLDIR=C:\MySQL
```

- e. Make the INSTALLDIR option user editable.
- f. Click Finish.

The action is saved.

8. Click + to define the Configure MySQL action.

- a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
cmd
```

- b. Enter a description for the action.
- c. Click Next.

- d. On the Define Installation Action Options panel, add the following options:

```
/C C:\MySQL\bin\MySQLInstanceConfig.exe -i -q "-lc:\mysql_install_log.txt"
"-nMySQL Server 5.5" "-pC:\MySQL" -v5.5 "-tC:\MySQL\my-template.ini"
"-cC:\mytest.ini" ServerType=DEVELOPMENT DatabaseType=MIXED
ConnectionUsage=DSS Port=3306 ServiceName="MySQLD"

RootPassword=pass
```

- e. Make the RootPassword option user editable.
- f. Click Finish.

The action is saved.

9. Click + to define the Grant Permissions to Root action.

- a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
cmd
```

- b. Enter a description for the action.
- c. Click Next.

- d. On the Define Installation Action Options panel, add the following options:

```
/C C:\MySQL\bin\mysql -u root -e "GRANT ALL PRIVILEGES ON *.* TO
'Root'@'localhost' IDENTIFIED BY 'pass';"

--password=pass
```

- e. Make the password option user editable.
- f. Click Finish.

The action is saved.

10. Click Next, and then Finish.

The MySQL application is saved.

Define MediaWiki Database

The MediaWiki database is the place where the content of the wiki web page is stored.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Application.

The Define an Application wizard opens in the Application Details panel.

2. Specify the Name, Version, and Vendor for the application.
3. Specify the location of the application files (absolute path or relative to the Staging folder).
4. Specify MySQL as a prerequisite for this application.
5. Click Next.

The System Requirements panel opens.

6. Set the following recommended values:
 - Number of CPUs – 1
 - Memory – 512 MB
 - Disk Space – 5 GB
 - Operating System – All Windows options
7. Click Next.

The Configure Installation Actions panel opens. You will define three installation actions: Install Database, Stop MySQL, and Start MySQL.

8. Click + to define the Install Database action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
cmd /C "C:\MySQL\bin\mysql
```

- b. Enter a description for the action.
- c. Click Next.
- d. On the Define Installation Action Options panel, add the following options:

```
--user=root
```

```
--password=pass
```

```
< %CD%\wiki_db.txt > %CD%\output.txt"
```

- e. Make the password option user editable.
- f. Click Finish.

The action is saved.

9. Click + to define the Stop MySQL action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
net stop MySQLD
```
 - b. On the Define Installation Action Options panel, only check the preview and click Finish.
The action is saved.
10. Click + to define the Start MySQL action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:

```
net start MySQLD
```
 - b. On the Define Installation Action Options panel, only check the preview and click Finish.
The action is saved.
11. Click Next, and then Finish.
The MediaWiki Database application is saved.

Define MediaWiki Content

These files are the application files that the wiki needs for running.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Application.

The Define an Application wizard opens in the Application Details panel.

2. Specify the Name, Version, and Vendor for the application.
3. Specify the location of the application files (absolute or relative path to the Staging folder).
4. Specify PHP and MediaWiki Database as prerequisites for this application.
5. Click Next.

The System Requirements panel opens.

6. Set the following recommended values:
 - Number of CPUs – 1
 - Memory – 512 MB
 - Disk Space – 5 GB
 - Operating System – All Windows options
7. Click Next.

The Configure Installation Actions panel opens. You will define two installation actions: Copy Files and Update DB Server.

8. Click + to define the Copy Files action.
 - a. On the Select Installation Action panel, in the Action field, input the command to execute:
`xcopy`
 - b. Enter a description for the action.
 - c. Click Next.
 - d. On the Define Installation Action Options panel, add the following two options:
`%CD%\* C:\Apache\htdocs`
`/S`
 - e. Click Finish.
The action is saved.
9. Click + to define the Update DB Server action.
 - a. On the Select Installation Action panel, select Update File from the Action drop-down list.

- b. Enter a description for the action.
- c. In the File Name field, enter the path to the configuration file:
C:\Apache\htdocs\LocalSettings.php
- d. Click Next.
- e. On the Define Installation Action Options panel, add the following five options:

CONFIG_FILE_ACTION=FILEUPDATE

CONFIG_FILE_NAME=C:\Apache\htdocs\LocalSettings.php

#DB_SERVER#=%DEPENDINGHOST%

#BLOG_TITLE#=<Wiki Title>

#DBPASSWORD#=pass

- f. Make the #DB_SERVER#, BLOG_TITLE, and #DBPASSWORD# options user editable.
- g. Click Finish.

The action is saved.

- 10. Click Next, and then Finish.

The MediaWiki Content application is saved.

Create Wiki Service Template

Service provisioning is the creation of a working instance of a service that is based on a service template. In this procedure, you create a service template for Wiki deployment that consists of the applications that you have already defined, and the number of machines that are required to host the service.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Service Template.

The Create Service Template panel opens.

2. Input a name and description for the service template and select the number of machines that are required to host the service. The number of machines depends on the expected load of the wiki web page and the amount of content that you intend to store on it. In this example, we install all application on one server.
3. Click Details to display the configuration of the machine. Leave all the values set to Auto. With this setting, the supported OS types of all applications in the template are checked during provisioning, and the most convenient VMware template is automatically selected.
4. Specify a Service Information File. This HTML file is displayed to the end users at the end of the provisioning process informing them how to access the service.
5. Click + (New) to add the MediaWiki Content application to the service template.

The Service Template Application Configuration Wizard opens in the Select Applications panel.

6. Select MediaWiki Content from the list of available defined applications and click Next.

The Installation Actions panel opens. Do not make any changes here.

7. Click Next.

The Confirm Application Configurations panel opens.

8. Click Finish.

The Create Service Template panel updates with the details of the application and machine to host it.

9. Repeat steps 5 through 8 to add the MediaWiki Database application. Due to the applications prerequisites, the Apache HTTP Server, PHP, and MySQL are added automatically during provisioning.

10. Click OK.

CA Server Automation adds the template to the list of service templates in the Explore tree. The service template is ready for provisioning.

Deploy a Wiki

The creation of service templates enables one-click service provisioning for end users.

Follow these steps:

1. Click Resources, and in the Explore tree select Service Templates under OnDemand Services.
2. In the Summary tab, select the created Wiki service template, and click the Provision Service icon.
3. Input a name for the provisioned service instance.
4. Provide a title for the Wiki web page. This user input has been specified in the definition of the MediaWiki Content application.
5. Click OK.

CA Server Automation provisions the requested service instance and adds it to the list of resources under OnDemand Resources in the Explore tree. Track the status of the provisioning in the Jobs pane.


6. After the provisioning has completed, a page with a link to the Wiki web page is displayed. Follow the link to verify that the provisioning was successful.

How to Deploy Oracle WebLogic Server

This example shows how to prepare a template that enables end users to deploy an instance of Oracle WebLogic application server. First, you define application executables and their execution actions, then you create a service template containing this application.

The following process describes a high-level overview of how System Administrators create a service template that enables one-click service provisioning for Service Consumers:

How to Deploy Oracle WebLogic Server



1. [Configure Machine Templates for WebLogic](#) (see page 706)
Specify default machine configurations to use in service templates.
2. [Define the WebLogic Application](#) (see page 707)
Specify the WebLogic application and the resource and operating system requirements for the application.
3. [Create WebLogic Service Template](#) (see page 710)
Specify the machine that is required to host the application and set up a deployable service template.
4. [Deploy WebLogic Server](#) (see page 711)
Deploy an instance of the service to your server environment.

Configure Machine Templates for WebLogic

CA Server Automation enables you to specify VM templates for specified operating systems. These templates can then be automatically applied to machines required during service provisioning.

Follow these steps:

1. Click Resources, and in the Explore tree, select OnDemand Services.
2. On the right-hand panel, click the tool icon for the vCenter Server containing the templates that you want to use.

The Machine Template Configuration panel opens.

3. Select an operating system from the list, select a Template from the list of available VM templates, and click Set As Default.

CA Server Automation associates the selected VM template with the specified operating system.

4. (Optional) Click an operating system group folder to view and manage the list of all VM template settings for that OS family.

Click OK to exit Machine Template Configuration.

Define the WebLogic Application

The initial step is to define the WebLogic application to be used in the service template and the actions that are required to execute the application.

Follow these steps:

1. In the Staging directory in the CA Server Automation installation directory, create a folder for the WebLogic application.
2. Copy all the application files to the folder.
3. In the CA Server Automation user interface click Resources, and in the Explore tree, right-click OnDemand Services, and select New Application.

The Define Application wizard opens in the Application Details panel.

4. Specify the Name, Version, and Vendor for the application.
5. Specify the location of the application files (absolute path or relative to the Staging folder).
6. Click Next.

The System Requirements panel opens.

7. Set the following recommended values:
 - Number of CPUs – 1
 - Memory – 512 MB
 - Disk Space – 5 GB
 - Operating System – All Microsoft Windows Server versions.
8. Click Next.

The Configure Installation Actions panel opens. You will define four installation actions: Install WebLogic, Create Files for Domain, Create Domain, and Start WebLogic Service.

9. Click + to define the Install WebLogic action.
 - a. On the Select Installation Action panel, select Execute Program from the Action drop-down list.
 - b. Enter a description for the action.
 - c. In the Program Name field, select the following executable file:


```
server103_win32.exe
```
 - d. Click Next.
 - e. On the Define Installation Action Options panel, add the following three options:


```
-mode=silent
```

```
-silent_xml=C:\SA\Weblogic-10.3\silent.xml
```

-log=C:\CA_SA_Weblogic_silent.log

- f. Click Finish.

The action is saved.

10. Click + to define the Create Files for Domain action.

- a. On the Select Installation Action panel, select Execute Program from the Action drop-down list.
- b. Enter a description for the action.
- c. In the Program Name field, input the batch file to execute:

C:\SA\Weblogic-10.3\CreateWeblogicFile.bat

- d. Click Next.
- e. On the Define Installation Action Options panel, add the following option:

-domiannname CA_SA_Weblogic

Make the option editable for end users.

- f. Click Finish.

The action is saved.

11. Click + to define the Create Domain action.

- a. On the Select Installation Action panel, select Execute Program from the Action drop-down list.
- b. Enter a description for the action.
- c. In the Program Name, input the batch file to execute:

C:\bea\WLS_WLI_WLP_103_silent\wlserver_10.3\common\bin\wlst.cmd

- d. Click Next.
- e. On the Define Installation Action Options panel, add the following option:

C:\SA\Weblogic-10.3\CreateDomain.py

- f. Click Finish.

The action is saved.

12. Click + to define the Start WebLogic Service action.

- a. On the Select Installation Action panel, select Execute Program from the Action drop-down list.
- b. Enter a description for the action.
- c. In the Program Name field, input the batch file to execute:

C:\SA\Weblogic-10.3\InstallStartWeblogicSvc.bat

- d. Click Next.

- e. On the Define Installation Action Options panel, only check the preview and click Finish.

The action is saved.

13. Click Next, and then Finish.

The WebLogic application is saved.

Create WebLogic Service Template

Service provisioning is the creation of a working instance of a service that is based on a service template. In this procedure, you create a service template for WebLogic Server deployment that consists of the application that you have already defined, and the number of machines that are required to host the service.

Follow these steps:

1. Click Resources, and in the Explore tree, right-click OnDemand Services, and select New Service Template.
The Create Service Template panel opens.
2. Input a name and description for the service template and select the number of machines that are required to host the service. The number of machines depends on the expected load of the server.
3. Click Details to display the configuration of the machine. Leave all the values set to Auto. With this setting, the supported OS types of all applications in the template are checked during provisioning, and the most convenient VMware template is automatically selected.
4. Specify a Service Information File. This HTML file is displayed to the end users at the end of the provisioning process informing them how to access the service.
5. Select the Additional Applications check-box. This enables you to add applications to WebLogic Server during service provisioning.
6. Click + (New) to add the WebLogic application to the service template.
The Configure Service Template Application wizard opens in the Select Applications panel.
7. Select WebLogic from the list of available defined applications and click Next.
The Configure Installation Actions panel opens. Do not make any changes here.
8. Click Next.
The Confirm Application Configurations panel opens.
9. Click Finish.
The Create Service Template panel updates with the details of the application and machine to host it.
10. Click OK.
CA Server Automation adds the template to the list of service templates in the Explore tree. The service template is ready for provisioning.

Deploy WebLogic Server

The creation of service templates enables one-click service provisioning for end users.

Follow these steps:

1. Click Resources, and in the Explore tree select Service Templates under OnDemand Services.
2. In the Summary tab, select the created WebLogic service template, and click the Provision Service icon.
3. Input a name for the provisioned service instance.
4. Provide a WebLogic domain name. This user input has been specified in the definition of the WebLogic application.
5. Click Additional Applications and select from the available applications to deploy them with the service template.

Note: Only applications that are dependent on applications already in the service template (in this case on WebLogic) are available for selection. Set WebLogic as a prerequisite for all the applications that you want to make available during provisioning.

6. Click OK.

CA Server Automation provisions the requested service instance and adds it to the list of resources under OnDemand Resources in the Explore tree. Track the status of the provisioning in the Jobs pane.

7. After the provisioning has completed, a page with a link to the WebLogic Server Administration Console is displayed. Follow the link to verify that the provisioning was successful.

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 ITCM interface and the other job is queued. If you view the jobs from the CA ITCM 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.

CA Patch Manager

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 ITCM. 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 whether 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 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 the Software Delivery adapter to synchronize 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 the Software Delivery adapter to synchronize 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 the Software Delivery adapter to synchronize 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 the Software Delivery adapter to synchronize 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 ITCM 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 that are permitted to run simultaneously.

Default: 20

CONFIG_KEY_SDA_Img_Job_Rrtry_Delay

Specifies the delay for retrying an 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 dpmsd CLI.

Default: 60

Limits: minutes

CONFIG_KEY_SDA_PROVISIONING_TIMEOUT

Specifies the default timeout value for OSIM job.

Default: 120

Limits: minutes

CONFIG_KEY_SDA_SCREG_RETRY_MAX_COUNT

Specifies the maximum number of attempts to register the Software Delivery service 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 CA 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

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 Amazon Web Service (AWS) accounts
- 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 Images (AMIs) and instances, 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
- Amazon EC2 instance types
 - t1.micro
 - m1.small
 - c1.medium
 - m1.large
 - m1.medium
 - 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 and Provision Amazon EC2 Resources

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 and provisioning Amazon EC2 resources in CA Server Automation.

How to Configure and Provision Amazon EC2 Resources



1. [Configure the Amazon Web Service \(AWS\)](#) (see page 723)
2. [Configure the Amazon Virtual Private Cloud](#) (see page 725)
3. [Configure and Validate the Amazon EC2 Server](#) (see page 726)
4. [Provision EC2 Instances](#) (see page 727)

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	Multiple accounts are 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

Select Region

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.

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 <i>Amazon Virtual Private Cloud Network Administration Guide</i> .	Subnet: Subnet IP Address: Domain Name:
Create the VPN connection	See the <i>Amazon Virtual Private Cloud Network Administration Guide</i> .	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-XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX.pem
 - **Private Key File Full Path:** The Amazon Private Key file (.pem).
Example: C:\ec2_account\pk-YYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYY.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
 - **Default Account:** Whether the account is enabled by default.
Example: Enabled
 - **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)

How to Provision EC2 Instances

In CA Server Automation user interface, use the following right-click popup wizard to provision EC2 instances: Provisioning, Provision Amazon EC2 Instance. The following table provides additional information.

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.

Bare Metal Provisioning to a Cisco UCS Blade

For bare metal provisioning to a Cisco UCS blade, you can use any of the following methods in the CA Server Automation user interface:

- Cisco UCS provisioning wizard
- Rapid Server Imaging (RSI) provisioning
- CA ITCM

You also can use either of the following command line provisioning methods:

- `dpmucs` command
- `dpmrsi` command

Consider the following information to decide which method to use:

- Cisco UCS blade provisioning can take significant time due to the blade power cycle and the time required to associate the blade with a specific service profile.
- If you are using an RSI or ITCM provisioning method, you must ensure that the service profile is already associated to the blade *before* running the operation.
- For ITCM deployment, you must retrieve the MAC address from the service profile to be used for PXE boot. After the provisioning job is submitted to CA ITCM, you must run blade powercycle from the CA Server Automation user interface.
- 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.

Note: Cisco UCS provisioning does not require Wake-On-LAN because it uses the Cisco UCS Manager for power control.

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.

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 Domain Manager where the operation should be performed. This name is optional when only one Software Delivery adapter or CA ITCM domain manager is configured. This parameter is valid only for CA Server Automation.

Scalability Server

Serves as the primary interface for the agent and distributes the CA ITCM workload across multiple hosts. CA Software Delivery 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 CA ITCM Domain Manager where the operation should be performed. This is optional when only one SD adapter or CA ITCM domain manager is configured. This parameter is valid only for CA Server Automation.

Scalability Server

Serves as the primary interface for the agent and distributes the CA ITCM workload across multiple hosts. CA Software Delivery 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.

Using the MKSYSB Utility

More information:

[Prerequisite Knowledge](#) (see page 734)

[How to Provision an IBM AIX Image Using MKSYSB](#) (see page 734)

Prerequisite Knowledge

Verify that you know the following information before using CA Server Automation to deploy an MKSYSB image:

- You are familiar with the IBM PowerVM and LPAR operating environments, including the NIM MKSYSB utility.
- You have a basic understanding of the CA Server Automation user interface and how to provision resources.


How to Provision an IBM AIX Image Using MKSYSB

As a System Administrator, you use CA Server Automation to provision an IBM AIX system image using the MKSYSB utility.

IBM AIX images can be provisioned to virtual machines (LPARs) or to physical IBM computers that already run an AIX operating system. The IBM NIM MKSYSB utility provides one-step IBM AIX image provisioning that can include both the operating system and additional software. This provisioning approach is both faster and more efficient than default run-time processing; however, it requires that you have prepared MKSYSB images or resources available.

The following flowchart shows a set of procedures required to deploy the IBM AIX system image.

How to Provision an IBM AIX Image Using MKSYSB



Use the following tasks to provision an IBM AIX image using MKSYSB:

1. [Complete Prerequisite Tasks](#) (see page 736)
2. [Install NIM Adapter on AIX NIM Server](#) (see page 737)
3. [Update the Hashed Password Variable](#) (see page 737)
4. [Increase the size of the /tmp and /opt filesystems](#) (see page 737)
5. [Start the NIM Adapter Daemon](#) (see page 738)
6. [Configure NIM Master Server](#) (see page 738)
7. [Provision the MKSYSB Image](#) (see page 739)

Complete Prerequisite Tasks

Verify the following prerequisites before using CA Server Automation to deploy an MKSYSB image:

- You have obtained prepared MKSYSB images or NIM MKSYSB resources.
- You have determined which resources and IBM servers to use for imaging.
- You have administrator credentials for CA Server Automation and the IBM PowerVM and NIM servers.
- You can access the CA Server Automation user interface.

Install NIM Adapter on AIX NIM Server

You can install the NIM adapter using a graphical interface or a command line text console.

Follow these steps:

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.

Installation completes, and a summary page appears.
7. Click OK to exit the installer.

Update the Hashed Password Variable

When a NIM client is imaged, the IBM AIX installer creates an empty root password. To specify a root password, update the `ca_post_install.sh` script file. Set the hashed password (DES format), and update the `HASH_PASSWORD` variable for NIM clients to use.

Note: The `ca_post_install.sh` script file has a default hash password that translates to the plaintext value of **admin**.

Follow these steps:

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 /tmp and /opt File Systems

Default file system sizes are not always large enough for IBM image provisioning. If there is insufficient space in the `/tmp` or `/opt` file systems, agent deployment to NIM clients can fail. If the selected defaults are too small for the agents to install, increase them to at least 400 MB (for `/tmp`) and 700 MB (for `/opt`). If you do not have NIM scripts that increase the file system sizes, uncomment the `chfs` commands in the `ca_post_install.sh` script. `chfs` commands enable the NIM adapter to increase the file system size by using this script as a NIM script resource following NIM client imaging.

Note: Only enable these lines if your scripts do not already use them.

Follow these steps:

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 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. Do not set them lower than the default.

3. Save the file and exit.

Start or Stop the NIM Adapter Daemon

The NIM adapter daemon starts automatically after installation or when the system is booted up.

To start the NIM adapter daemon manually, run the following command:

```
install-path/imaging/bin/canimstart.sh start
```

To stop the NIM adapter daemon manually, run the following command:

```
install-path/imaging/bin/canimstart.sh stop
```

Configure NIM Master Server

You can configure a NIM master server after CA Server Automation installation. The NIM master server provides the MKSYSB resources.

Follow these steps:

1. Log in to the CA Server Automation user interface, and click Administration.
2. Click Configuration, and in the left pane of the Configuration page, click NIM Master Server.
The NIM Servers page appears.
3. Click + (Add).
4. Enter the NIM administrator credentials, and click OK.
5. Click Validate to verify the connection status.

Provision the IBM AIX Image

When the IBM NIM environment is configured with MKSYSB resources, use the CA Server Automation user interface to provision the IBM AIX image.

Follow these steps:

Note: Verify that you are logged in to the CA Server Automation user interface as an administrator.

1. Click Resource, and open the Explore pane.
2. Right-click an IBM PowerVM resource, and select Provisioning, Provision NIM.
The Provision AIX with NIM dialog appears.
3. Complete the following fields:

NIM Master

Specify the computer where you configured your NIM environment. There can only be one active NIM master server in an environment.

Machine Resource Name

Specify the target computer for NIM installation and update operations.

4. In the Installation Type drop-down, select mksysb.
5. In the Resource Type drop-down, select Resource Group.
6. Select an MKSYSB resource group.
7. In the System Attributes pane, provide the administrator credentials to log in to the host NIM system.

Note: The administrator user ID and password must match the ones specified in the *install_path/imaging/etc/ca_post_install.sh* script during IBM AIX imaging configuration.

8. Click Add Computer.

The new IBM AIX system image is created and ready for use. By having the MKSYSB resource group already available, provisioning did not require a multistep build process.

Rapid Server Imaging


As a *System Administrator*, you use CA Server Automation to manage and provision your physical and virtual server environments. You require a reliable imaging mechanism to provide server images for provisioning, backup, disaster recovery, and migration across your server environments.

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 can deploy the captured images to a different managed server or to a bare-metal system. You can capture and deploy, or directly migrate, 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.

Note: For more information about installing and configuring RSI, see the RSI documents in the CA Server Automation bookshelf.

The following diagram provides an overview of RSI functionality in CA Server Automation:



1. Verify [RSI Prerequisites](#) (see page 741)
 - [RSI Prerequisites for Hyper-V](#) (see page 743)
2. [Register the RSI Server](#) (see page 743)
3. [Capture RSI Images](#) (see page 744)
 - [Capture Windows Driver Sets](#) (see page 744)
 - [Capture Windows Driver Collections](#) (see page 745)
4. [Deploy RSI Images](#) (see page 746)
 - [Deploy RSI Images to Bare-Metal Servers](#) (see page 747)
 - [Deploy RSI Images to AppLogic](#) (see page 748)
5. [Migrate Servers with RSI](#) (see page 748)
6. [Manage RSI Images and Drivers](#) (see page 750)
7. [Automate RSI Actions](#) (see page 751)

Combine RSI functionality with other CA Server Automation functionality to support the following specific business goals:

- [How to Deploy RSI Using ITCM](#) (see page 751)
- [How to Set Up RSI for AppLogic](#) (see page 759)
- [How to Backup and Restore with RSI](#) (see page 761)
- [How to Perform Disaster Recovery with RSI](#) (see page 762)

RSI Prerequisites

Before you can register an RSI server with CA Server Automation and can use RSI functionality, configure your RSI environment.

Follow these steps:

1. Install the RSI server.
Note: For more information about installing RSI servers, see the *RSI Server Installation Guide*.
2. Configure the RSI server.
Note: For more information about configuring RSI servers, see the *RSI Server Administration Guide*.
3. Configure hypervisors.
Note: For more information about configuring hypervisors, see the *RSI Server Administration Guide*.

4. (Optional) If you want to manage additional servers outside the RSI network environment, configure external networks.

Note: For more information about configuring external networks, see the *RSI Server Administration Guide*.

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: 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.

Note: CA Server Automation supports alternative RSI deployment methods. For more information, see [How to Deploy RSI Using ITCM](#) (see page 751) and [How to Set Up RSI for AppLogic](#) (see page 759).

RSI Prerequisites for AppLogic

Using Rapid Server Imaging (RSI) with AppLogic requires the deployment of an RSI server as an AppLogic application in the AppLogic grid.

Follow these steps:

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*.

RSI Prerequisites for Hyper-V

When working with virtual machines in a Microsoft Hyper-V/RSI environment, on any supported operating system, verify the following prerequisites:

- Each VM must have at least one Legacy Network Adapter that is connected to one of the DynaCenter boot networks. Using only one Legacy Network Adapter is recommended. If you use more than one adapter, ensure 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.

Note: For a complete list of supported operating systems, see the support matrix in the *RSI Server Installation Guide*.

- Each VM must have one or more para-virtualized network adapters (for use by the OS when it is deployed to the VM).

Register the RSI Server

To enable RSI functionality across environments managed by CA Server Automation, register the RSI server and optional environment settings.

Follow these steps:

1. Click Administration, and in left Configuration panel, locate Provisioning, and click Rapid Server Imaging.
2. In the RSI Servers panel, click the + (Add) icon, add the RSI Server, and validate that the connection status is green.
3. (Optional) If you use virtual machines, in the RSI Registered Hypervisors panel, click + (Add) to add hypervisors.

Note: Hypervisor credentials are validated.

4. (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.

5. (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.

CA Server Automation connects to the RSI server and enables RSI functionality across your server environments.

Capture RSI Images

When CA Server Automation is integrated with an RSI server, you can capture RSI images of physical and virtual servers for later deployment. Use stored images to provision new physical or virtual servers, as backup images to restore servers, or for disaster recovery.

Follow these steps:

1. Click Resources, right-click the machine to capture from in the Explore tree, and select Provisioning, Capture Image.
2. Select the RSI Server and Depot to use, and set the Image Name.
3. Specify the Target MAC and OS Type of the machine to capture.
4. (Optional) Set a Description for the image, and a Profile URL of an external file containing capture parameters.
5. (Optional) Specify File Systems to Exclude from the image.
6. (Optional) Select Live Capture to perform an online capture of the running server. Disable Live Capture to shut down the server first, and capture an image after reboot.
7. (Optional) If multiple networks are registered and Live Capture is disabled, select the network to use for capture.
8. Click OK.

The panel submits a request to capture the image from the target machine and stores it in the selected depot. Track the progress of the deployment in the Jobs pane.

Important! For Windows images, ensure that you capture drivers that are required for successful deployment of the image.

Capture Windows Driver Sets

For the successful deployment of RSI images on Windows, it is often necessary to deploy the drivers associated with software in the image. After RSI is successfully installed and configured, use CA Server Automation to capture driver sets from target servers for deployment with RSI images.

Follow these steps:

1. Click Resources, right-click Data Center in the Explore tree, and select Management, RSI Resources.
The RSI Resources panel opens.
2. Select the Driver Set tab.
The Driver Set list displays the current set of captured driver sets.

3. Click + (Capture Driverset) in the toolbar.
4. Enter the Target Server to capture from, and a Driver Set Name.
5. Specify the Target MAC of the machine to capture from, and the RSI Server to use.
6. (Optional) Specify a Description for the driver set.
7. (Optional) If multiple networks are registered, select the network to use for capture.
8. Click OK.

The dialog submits a request to capture the driver set from the target server. Track the progress of the capture in the Jobs pane.

Capture Windows Driver Collections

For the successful deployment of RSI images on Windows, it is often necessary to deploy the drivers associated with software in the image. After RSI is successfully installed and configured, use CA Server Automation to capture driver collections directly from driver media for deployment with RSI images.

Follow these steps:

1. Click Resources, right-click Data Center in the Explore tree, and select Management, RSI Resources.
The RSI Resources panel opens.
2. Select the Driver Collection tab.
The Driver Collection list displays the current set of captured driver collections.
3. Click + (Collect Driver Import) in the toolbar.
The Collect Driver Import dialog opens.
4. Enter the media to capture from, a description for the driver collection, the source and vendor IDs and click OK.

The dialog submits a request to capture the driver collection from the target media. Track the progress of the capture in the Jobs pane.

Deploy RSI Images

After RSI is installed and configured, you can deploy previously captured RSI images to physical servers and virtual machines.

Note: This procedure is the basic RSI image deployment to a server or virtual machine already set up with an operating system and RSI agent. For more complex image deployments see [Deploy RSI Images to Bare-Metal Servers](#) (see page 747) and [Deploy RSI Images to AppLogic](#) (see page 748).

Follow these steps:

1. Click Resources, right-click the machine to deploy to in the Explore tree, and select Provisioning, Deploy Image.

The Deploy Image wizard opens.

2. Select the RSI Server and Depot to use, and the Images to deploy.
3. (Optional) If multiple networks are registered, select the network to use for deployment.
4. Enter the Target MAC of the machine to host the image.
5. (Optional) Select Scale to deploy only the parts of the image containing data, and set a Profile URL of an external file containing deployment parameters.
6. (Windows Only) Select a Driver Set to deploy.

Note: After you successfully deploy a Windows image, Windows runs chkdsk. The target system automatically restarts, then Windows runs PnP (plug-and-play) and installs drivers.

7. (Optional) Specify a Host Name to apply to the machine hosting the image and select the Network Interfaces to use.

Note: Setting a host name overrides the host name in the image or in a profile.

8. Click OK.

The wizard submits a request to deploy the image to the target machine. Track the progress of the deployment in the Jobs pane.

Deploy RSI Images to Bare-Metal Servers

RSI enables you to deploy a complete image, including operating system and RSI agent, to a physical server or virtual machine that currently contains no software. Bare-metal deployment enables you to add a physical server, such as a UCS blade, or add a new virtual machine to your environment, and then provision it with a complete image containing all your software requirements.

Follow these steps:

1. Click Resources, right-click the machine or environment to deploy to in the Explore tree, and select Provisioning, Add a new Windows/Linux System (RSI).

The Add Bare Metal wizard opens.

2. Select the RSI Server and Depot to use, and the Images to deploy.
3. (Optional) If multiple networks are registered, select the network to use for deployment.
4. Specify the Machine Type.
 - Specify the Target MAC for a Physical machine.
 - Specify the Hypervisor and Target System ID for a Virtual machine.
5. (Optional) Select Scale to deploy only the parts of the image containing data, and set a Profile URL of an external file containing deployment parameters.
6. (Windows Only) Select a Driver Set to deploy.

Note: After you successfully deploy a Windows image, Windows runs chkdisk. The target system automatically restarts, then Windows runs PnP (plug-and-play) and installs drivers.

7. (Optional) Specify a Host Name to apply to the machine hosting the image and select Network Interfaces to use.

Note: Setting a host name overrides the host name in the image or in a profile.

8. Click OK.

The wizard submits a request to deploy the image to the target machine. Track the progress of the deployment in the Jobs pane.

Deploy RSI Images to AppLogic

After RSI is installed and configured, you can deploy previously captured RSI images as new applications in an AppLogic grid. CA Server Automation combines AppLogic operating system application deployment with RSI image deployment to provide a complete instance meeting your software requirements.

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.

CA Server Automation submits a request to deploy the image to a new instance of the OS application template in the selected grid with the specified resource allocations. Track the progress of the deployment in the Jobs pane.

Migrate Servers with RSI

When CA Server Automation is integrated with an RSI Server, you can migrate server images directly from a physical or virtual machine to another physical or virtual machine. Optionally, you can migrate a server to a bare metal server or a new instance of a virtual machine.

Important! Before migrating Windows images, ensure that you capture drivers that are required for successful deployment of the image.

Follow these steps:

1. Click Resources, right-click the machine to capture from in the Explore tree, and select Provisioning, RSI Migration.

The Perform Migration wizard opens in the Migrate Source panel.

2. Select the RSI Server and Depot to use, and optionally edit the Image Name.

3. Specify the Target MAC and OS Type of the machine to capture.
4. (Optional) Set a Description for the image, and a Profile URL of an external file containing capture parameters.
5. (Optional) Specify File Systems to Exclude from the image.
6. (Optional) Select Live Capture to perform an online capture of the running server. Disable Live Capture to shut down the server first, and capture an image after reboot.
7. (Optional) If multiple networks are registered and Live Capture is disabled, select the network to use for capture.
8. (Optional) Select Delete Image to delete the image from the depot after migration is complete.
9. Click Next.

The Migrate Destination panel opens.

10. Select the Provisioning Type.
 - Select Default to deploy to an existing machine.
 - Select Bare Metal to provision to a new server or virtual machine that does not currently have a deployed RSI agent or operating system installed.
 - Select New VM to provision a new instance of a virtual machine to host the migrated server image.

The VMware vCenter Provisioning wizard opens. Follow the procedure described in [Add a Virtual Machine \(vCenter Server\)](#) (see page 321), starting at step 2. Completing the wizard returns to the Migrate Destination panel.

Note: New VM provisioning does not use VM templates. The disk and network configuration steps are not optional.

Important! Specify the same datastore for all disks. If you create multiple disks with different datastores, new VM provisioning fails.

11. Select the RSI server to use for deployment.

Note: If the RSI servers for capture and deploy are different, they must access the same depot.

12. Specify the Machine Type.
 - Specify the Target MAC for a Physical machine.
 - Specify the Hypervisor and Target System ID for a Virtual machine.
13. (Optional) Select Scale to deploy only the parts of the image containing data, and set a Profile URL of an external file containing deployment parameters.

14. (Windows Only) Select a Driver Set to deploy.

Note: After you successfully deploy a Windows image, Windows runs chkdsk. The target system automatically restarts, then Windows runs PnP (plug-and-play) and installs drivers.

15. (Optional) Specify a Host Name to apply to the machine hosting the image and select Network Interfaces to use.

Note: Setting a host name overrides the host name in the image or in a profile.

16. Click Finish.

The wizard submits a request to migrate the image to the target machine. Track the progress of the migration in the Jobs pane.

Note: For migration to a physical bare metal server which cannot be network booted using WakeOnLan from CA Server Automation, manual intervention may be required to reboot the server during migration.

Manage RSI Images and Drivers

CA Server Automation enables you to manage captured images, driver sets, and collections.

Follow these steps:

1. Click Resources, right-click Data Center in the Explore tree, and select Management, RSI Resources.

The RSI Resources panel opens.

2. Select the Images, Driver Set, or Driver Collection tab.
3. Click + to add driver sets and collections, and click - to delete selected images, driver sets, and collections.

Note: Deleted images and drivers are also removed from the depot and are no longer available for deployment from RSI Server.

Automate RSI Actions

CA Server Automation enables you to automate RSI actions using Policy Management functionality. Each action duplicates the corresponding manual functionality with the addition of policy management inputs such as approval.

The following RSI actions are available to automate using policy management:

- Capture Image corresponds to [Capture RSI Images](#) (see page 744).
- Deploy Image corresponds to [Deploy RSI Images](#) (see page 746).

Policy Management enables you to automate actions on a regular scheduled basis or in response to event monitoring. For more information, see [Using Rules and Actions](#) (see page 561).

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 product CA ITCM (which includes CA 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 configuration, set the Domain Manager and Scalability Server settings to match your CA ITCM deployment.
3. Verify that the following CA Server Automation services are installed and configured to support RSI deployment using CA ITCM:

- Provisioning Manager
- Deployment and Configuration Distribution (Scalability) Server
- SDAdapter
- CA Integration to 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, ensure that you have the following access rights:

- Access to DVD3 of the CA Server Automation installation media that contains the RSI packages.
- 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 754)
2. [Customize the RSI Server Package](#) (see page 755)
3. [Deploy the RSI Server](#) (see page 757)
4. [Register the RSI Server](#) (see page 743)
5. [Register the RSI Agent Packages](#) (see page 758)
6. [Deploy the RSI Agent Packages](#) (see page 758)

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 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 on the CA Server Automation 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
- 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*.

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 that was 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, scalability server name, operating system type, and click OK.

The panel submits a request to install the RSI Server package to the selected server.

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 that were 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, scalability server name, 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 Set Up Rapid Server Imaging for AppLogic

As a *System Administrator*, you use CA Server Automation to manage and provision your physical and virtual server environments. Your environment includes CA AppLogic as either an internal or external cloud. You require a reliable mechanism for deploying images to your AppLogic environment.

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.

The following process provides a high-level overview of how to set up and use CA Server Automation to simplify RSI image deployment to AppLogic:

How to Set Up and Use RSI for AppLogic




1. Verify [RSI Prerequisites for AppLogic](#) (see page 742).
 2. (Optional) If the AppLogic server is outside your internal network, [Add Proxy Servers](#) (see page 475).
 3. [Add AppLogic Servers](#) (see page 479).
 4. [Register the RSI Server](#) (see page 743) 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.
5. [Deploy RSI Images to AppLogic](#) (see page 748).

How to Backup and Restore with RSI

As a *System Administrator*, you use CA Server Automation to manage and provision your physical and virtual server environments. You require a reliable mechanism for backup and restore.

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.

The following process provides a high-level overview of how to use RSI to backup and restore your managed servers:



1. Verify [RSI Prerequisites](#) (see page 741)
 - [RSI Prerequisites for Hyper-V](#) (see page 743)
2. [Register the RSI Server](#) (see page 743).
3. Backup servers by [Capturing RSI Images](#) (see page 744) and storing them in the RSI depot.
 - (Windows Only) [Capture Windows Driver Sets](#) (see page 744)
 - (Optional) [Automate RSI Actions](#) (see page 751). [Define a Schedule](#) (see page 651) to capture RSI images on a periodic basis.
4. Restore servers by [Deploying RSI Images](#) (see page 746) from the RSI depot.


How to Perform Disaster Recovery with RSI

As a *System Administrator*, you use CA Server Automation to manage and provision your physical and virtual server environments. You require a reliable mechanism for disaster recovery.

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.

The following process provides a high-level overview of how to use RSI to perform disaster recovery for your managed servers:

How to Perform Disaster Recovery with RSI



1. Verify [RSI Prerequisites](#) (see page 741)
 - [RSI Prerequisites for Hyper-V](#) (see page 743)
2. [Register the RSI Server](#) (see page 743).
3. Back up servers by [Capturing RSI Images](#) (see page 744) and storing their images in the RSI depot.
 - (Windows Only) [Capture Windows Driver Sets](#) (see page 744)
 - (Optional) [Automate RSI Actions](#) (see page 751). [Define a Schedule](#) (see page 651) to capture RSI images on a periodic basis.
4. Recover server images to a new server by [Deploying RSI Images](#) (see page 746) from the RSI depot.
 - (Optional) [Automate RSI Actions](#) (see page 751). [Create a Rule](#) (see page 564) to deploy an image automatically, based on monitored server events.
 - [Deploy RSI Images to Bare-Metal Servers](#) (see page 747) to recover system images, including a new operating system, to empty servers.
 - [Deploy RSI Images to AppLogic](#) (see page 748) to recover system images into a base AppLogic OS application.

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.

Chapter 12: Setting Up Reservation Manager

Reservation Manager provides the capability to reserve physical and virtual machines, create reservation templates, view inventory, and manage reservations.

After Reservation Manager installation, complete the following activities to prepare 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.

This chapter describes post-installation tasks 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:

[Reservation Manager Prerequisites](#) (see page 765)

[Setup and Configuration](#) (see page 768)

[User Management](#) (see page 803)

[Administration](#) (see page 809)

[Chargeback](#) (see page 816)

[Customization](#) (see page 824)

Reservation Manager 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 766)

[Prepare CA Server Automation for Reservation Manager](#) (see page 768)

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:
 - 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. Identify the virtual platforms you will support and confirm the corresponding server support. The following are examples of servers and platforms to consider:

Amazon Elastic Compute Cloud

Provides imaging functionality for Windows and Linux provisioning to the Amazon Elastic Compute Cloud (EC2).

CA ITCM 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.

6. Identify all users to grant access to Reservation Manager. Confirm the following:
 - Reservation Manager administrator users
 - Reservation Manager end users
7. 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 test 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. (Optional) 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.

Setup and Configuration

Proper setup and configuration of Reservation Manager is outlined in this section.

Predefined Content and Configuration for Service Provisioning

In this release, some service provisioning functions performed in CA Server Automation are effective in Reservation Manager. Service templates are predefined in the CA Server Automation Core UI. Reservation Manager reservation templates are created based on the service templates. You can make reservations for services from the Reservation Manager UI or the CA Server Automation portal interface.

To provide predefined service provisioning, the following actions are performed:

- The OnDemand resource pools are predefined for AppLogic and VMware Virtual Center supported platforms. When a vCenter server is configured in the CA Server Automation portal, the ESX server resources are added to the OnDemand Resources pool.
- The Service Administrators org unit is added and configured to access the OnDemand Resources pools.
- The Reservation Templates are automatically created from any service templates that are predefined in CA Server Automation.

Note: Navigate to Reservation Templates and use the wizard to view the detail of the predefined reservation templates.

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.

Prerequisites for Supporting Reservations of Virtual Machines

Before Reservation Manager deploys virtual machines for use, identify one or more VMware ESX servers or clusters of VMware ESX servers for use. Next, 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, 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 the context of the vSphere Datacenter at the time they are added. If the VMware data center 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.
The list of resource pools appears.
2. Select Add Virtual Pool from the Actions menu in the upper right corner of the Resource Pools list and follow the instructions in the wizard.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.

2. Open the Approvals area, change the following values and click OK.

- Open Help Desk Ticket
- Help Desk Ticket Type
- Help Desk Ticket Template
- Automatically Close Help Desk Ticket Upon Approval

The configuration change takes effect when the next reservation is made.

Manage Snapshots in VMWare Resource Pools

Administrators can give users permission to take snapshots of VMWare virtual machines. They can also specify how many snapshots are allowed, and indicate whether to quiesce the file system.

To manage snapshots in resource pools

1. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The list of resource pools appears.

2. Click a resource pool.

The Resource Pool Details pane opens to the Properties tab.

3. 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.

EMC CLARiiON storage and NetApp storage are both supported.

Storage can be provisioned on EMC CLARiiON storage systems, automatically attached to an ESX server, and added to a resource pool in a single operation.

The Storage Provisioning Manager for NetApp must be configured before this feature can be used.

To provision and attach storage to VMware hosts in a VMware Resource Pool

1. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The list of resource pools appears.

2. Select a VMware resource pool.

The Resource Pool Details pane opens.

3. Click Resource Pool Details.

A page with a list of Datacenters opens.

4. Select a Datacenter.

The Datacenter is highlighted.

5. Click the Actions drop-down and select Edit.

The Edit Datastores page appears.

6. Click the Actions drop-down and select Provision Datastore...

7. 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.

8. 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The list of resource pools appears.

2. 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your reservation templates.

The Resource Pools page opens.

2. Select Create from the Actions menu in the upper right corner and follow the instructions in the wizard.

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 variables instead of or in combination with a name or prefix for a Hyper-V or VMware virtual machine, or for the folder name in the Reservation Manager Virtual Resource Pool. There is no limit to the number of characters when you enter substitution variables. These substitution variables 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.

%RESERVATIONNOTES%

Identifies notes that are associated with the reservation..

%RESOURCEPOOL%

Identifies the Reservation Manager resource pool name.

%TENANT%

Identifies the name of the tenant submitting the reservation.

%TENANTID%

Contains the value of the Tenant ID configuration setting. Serves as a short or abbreviated name for the tenant, suitable for use in virtual machine names. The tenant-specific setting is used if the user is a member of a tenant, otherwise the global configuration setting is used.

%USERNAME%

Identifies the name of the user who logged in to Reservation Manager.

%VCSERVERNAME%

Identifies the name of the VC server on which the reserved system resides..

When a prefix is used, a numeric suffix is appended to the name to 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 create a reservation for a virtual machine, as an administrator on the Reservation Manager home page click the Create Virtual Machine option and complete the wizard instructions.

Specify a Prefix VMWare for Virtual Machine Names

You can specify a prefix for VMWare 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The Resource Pools page opens.

2. Double-click a resource pool.

The Resource Pool details page opens, with the Properties tab displayed.

3. Select values for the following fields.

VM Name Generation

Name/Prefix

4. 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 776)

The following sections describe how to set up the Reservation Manager to support Hyper-V reservations.

Enable Hyper-V Virtual Machine Reservation

From the Reservation Manager, an administrator enables a user to reserve Hyper-V systems by setting up a resource pool and template.

- In the Reservation Manager, Manage your reservation pools option, select Add Hyper-V Pool from the Actions drop-down menu. A wizard guides you through the process to set up a pool.
- In the Reservation Manager, Manage your reservation templates option, select Create in the Actions drop-down menu. A wizard guides you through the process to set up a template.

Make Amazon Machine Images Available from EC2

Prerequisites


Review the following prerequisites before configuring Amazon Machine Images (AMIs) for reservations:

- CA Server Automation is configured with Reservation Manager
- You can access the Reservation Manager administration user interface
- You have CA Server Automation administrator credentials
- The Amazon Elastic Compute Cloud (EC2) virtual computing environment is configured in CA Server Automation
- You have general knowledge of the EC2 environment and AMIs

Configure Reservation Manager for Amazon EC2

As a System Administrator, you can configure Reservation Manager so end users can reserve Amazon Machine Images (AMIs) from the Amazon Elastic Compute Cloud (EC2). This process assumes that you are already logged in to the Reservation Manager administrator user interface.

The following diagram illustrates the process for making AMIs available for reservations:



Follow these steps:

1. [Create an EC2 Resource Pool](#) (see page 779)
2. [Add Amazon Machine Images to Inventory](#) (see page 780)
3. [Troubleshoot Adding Amazon Machine Images \(if Necessary\)](#) (see page 781)
4. [Create a Template for Amazon Machine Images](#) (see page 781)
5. [Decide How to Manage Help Desk Requests](#) (see page 782)
6. [Configure Help Desk to Manage Approvals](#) (see page 782)

Create an EC2 Resource Pool

Create an EC2 resource pool for AMI reservations.

Follow these steps:

1. Log in to the Reservation Manager using CA Server Automation administrator user credentials.
2. Select Administer Your Reservation Manager, Manage Your Resource Pools.
3. Click Actions and select Add EC2 Pool from the drop-down menu.
4. Enter a pool name and (optionally) a pool description, then click Next.
5. Use the Create Resource Pool wizard to set up your pool.

Key Pair Name

Defines the key pair name that is used to access the AMI instance. Select one from the drop-down list.

Note: Provide the private key file that is associated with the selected key pair 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 instances are to be deployed to Amazon's public cloud. Optionally, select an availability zone. The EC2 best practice is **not** to select a specific zone, but if desired you can select one from the list.

Private

Indicates a private area in the cloud; a Virtual Private Cloud (VPC). Select a subnet.

6. Specify limits on the duration and number of AMIs a user can request when submitting a reservation:

Maximum Days

Limits reservation length to the specified number of days. Choose Unlimited if you do not want to place restrictions on reservation length.

Maximum Systems

Limits how many instances a user can reserve from this pool at one time.

7. Select the organizational units whose members are permitted to start AMIs in the context of the selected EC2 pool, and click Finish.

Note: (Optional) Grant access to yourself before granting users access so that you can test reservations from this pool. Skip this step if you want to grant access later.

Add Amazon Machine Images to Inventory

Add AMIs to the system images inventory so users can reserve AMIs from EC2.

Follow these steps:

1. Log in to the Reservation Manager user interface using CA Server Automation administrator credentials.
2. Select Administer Your Reservation Manager, Manage Your System Image Inventory.
3. Select Add AMI Image from the Actions drop-down menu and make the following selections:
 - An operating environment for the AMI or Unknown to view unclassified images
 - An image from the filtered list that is based on the operating environment you selected
 - (If you specified Unknown for the operating environment) An operating environment from Update OS to classify an unknown image
 - A default Tier name (instance type) to use when starting 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
 - Any user data that you want available when this AMI is started
4. Move the security groups that are to be used to configure access to instances of this AMI, to the Selected Security Groups area. Select the User Override Mode check box to allow users to change the default security group settings when submitting reservations.
5. Move the organizational units that are permitted to install this image to the Selected Organizational Units area, and click Finish.

Troubleshoot Adding Amazon Machine Images

Symptom:

When starting the Add AMI Image wizard, no AMIs are available for selection even though the EC2 server is connected.

Solution:

By default, AMIs are available only to the owner specified in the EC2 server connection. You can change the default so that public AMIs, such as those provided by Amazon, are available to users. Run `dpmutil` from a Command Prompt on the Reservation Manager server to reconfigure EC2 adding published public AMIs to the inventory.

Create a Template for Amazon Machine Images

Create one or more reservation templates for AMIs so that users can make reservations in the Amazon Elastic Compute Cloud (EC2).

Follow these steps:

1. Log in to the Reservation Manager using CA Server Automation administrator user credentials.
2. Select Administer Your Reservation Manager, Manage Your Reservation Templates.
3. Select Create from the Actions menu.
4. Use the Create Reservation Template wizard to create an AMI template.

Specify Name page

Indicates a name and optional description.

Specify System Image page

Lists the system images available for deploying new systems. Select an AMI image from the System Images list.

Specify Requirements page

Indicates the system requirements. Specify the number of instances to create. (Optional) Change the security groups and select a different instance type (tier of service) if desired

Specify Actions page

(Optional) Select actions to run before provisioning a system, or after provisioning is completed or expired.

Specify Access Policy page

Select the organizational units whose members are to be granted access to the template.

Decide How to Manage Approval Requests

Reservation requests can be automatically approved or sent to an administrator for manual approval. If you configure the Reservation Manager to approve requests automatically and the resource is available, the reservation is approved immediately. No email is sent. If you configure the Reservation Manager not to approve reservation requests automatically, an email notification is sent to you after users submit new requests. The status of those requests is Pending Approval and require you to approve or reject the request manually.

Automatic approval of requests is enabled per resource pool. To configure automatic approval of requests for a specific resource pool, set the "Automatically approve reservation requests" option in the Access Policy tab of the Resource Pool Details page.

Configure Help Desk to Manage Approvals

Optionally, configure a help desk system to manage reservation approvals.

Note: This procedure applies to all platforms.

Follow these steps:

1. Log in to the Reservation Manager using CA Server Automation administrator user credentials.
2. Select Administer Your Reservation Manager, Manage Your Configuration Settings.
3. Configure the following settings in the Approvals area:

Automatically Close Help Desk Ticket Upon Approval

Specifies whether help desk tickets are immediately closed when a reservation is approved or rejected.

Default: false (tickets are not closed)

Help Desk Ticket Template

Specifies an optional ticket template that is used when opening tickets. Templates let you define default values for ticket attributes such as priority.

Default: None

Help Desk Ticket Type

Specifies the type of ticket that is opened.

Limits: Request | Incident | Problem

Default: Request

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

Enable AppLogic Application Reservation

From the Reservation Manager, an administrator enables a user to reserve AppLogic applications by setting up a resource pool and template.

- Select the Add AppLogic Pool option, in the Actions drop-down menu of the Reservation Manager, Manage your reservation pools option. A wizard guides you through the process to set up a pool.
- Select Create, in the Actions drop-down menu of the Reservation Manager, Manage your reservation templates option. A wizard guides you through the process to set up a template.

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 784)

[Edit a Resource Pool for IBM PowerVM Logical Partitions](#) (see page 784)

[Create a Template for IBM PowerVM Logical Partitions](#) (see page 786)

[Static IP Addresses for IBM PowerVM Logical Partitions](#) (see page 786)

[Define Network Address Pools](#) (see page 795)

[Configure Chargeback by Tier for IBM PowerVM Logical Partitions](#) (see page 821)

[Select a Chargeback Tier for IBM PowerVM Logical Partitions](#) (see page 822)

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. As an administrator on the Reservation Manager home page, click Manage your resource pools.

The Resource Pools page opens.

2. 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.

3. Complete the instructions in the wizard with the following when prompted:

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.

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. As an administrator on the Reservation Manager home page, click Manage your resource pools.

The Resource Pools page opens, showing existing resource pools in a table.

2. 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

- Maximum Systems
- Maximum Days
- Domain Manager
- Scalability Server

- Automatically approve reservation requests
- Allow users to manage logical partition power state

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 field:

Allow users to manage logical partitions power state

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. As an administrator on the Reservation Manager home page, 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.

2. Select Create from the Actions menu and complete the wizard.

Static IP Addresses for IBM PowerVM Logical Partitions

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 cancellation.

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 cancellation.

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 795)

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. As an administrator on the Reservation Manager home page, click Manage your configuration settings.

The Configure Settings page opens.

2. Edit the settings in the IBM PowerVM area.

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 by Importing Inventory

You can set up a resource pool quickly by linking it to a CA Server Automation Service.

To create a resource pool by importing inventory

1. As an administrator on the Reservation Manager home page, click Manage your resource pools.

The Resource Pools page opens.

2. Select Import Inventory from the Actions drop-down menu in the upper right corner of the Resource Pools list.

The Import Inventory 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.

3. Complete the Select Service wizard page, and click Next.

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. As an administrator on the Reservation Manager home page 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.

2. 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

3. Click Properties and modify the settings as necessary.
4. Click the Resource Pools tab and modify the list of selected resource pools for that system if necessary.
5. 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 788)

[Prerequisites for Supporting Reservations of Physical Systems](#) (see page 787)

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. As an administrator on the Reservation Manager home page, 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.

2. Select Add from the Actions menu and follow the wizard instructions.

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. As an administrator on the Reservation Manager home page, 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.

2. Select the Add Image item from the Actions drop-down menu and follow the wizard instructions.

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 788)

Reserve a System

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, click Reserve a system from the Reservation Manager user interface home page and complete the wizard instructions.

Make Services Available to Users

Use the following process to make services available to users for reservation:

1. Identify and establish user access to resource pools.
2. Identify and establish user access to service templates.

The following sections describe how to use the Reservation Manager to support service reservations.

Configuring Service Resource Pools

Reservation Manager uses resource pools when provisioning services.

When provisioning a service from the Server Automation interface, the following default OnDemand Resources pools are used.

OnDemand Resources

A resource pool containing virtual resources. Resources are automatically added to this pool when the Virtual Center is configured for Server Automation usage.

OnDemand EC2 Resources

A resource pool containing EC2 resources. Resources are automatically added to this pool when EC2 is configured for Server Automation usage.

By default, the Service Administrators organization unit has access to these pools. The administrator can remove resources from these pools and can modify the access policy as needed.

The default resource pools are automatically available in Reservation Manager. You can create additional resource pools in Reservation Manager with customized services and user access to meet your needs.

The end user must have access to at least one resource pool to make a service reservation. Org unit membership determines user access to resource pools.

Create a Template for a Service

Service reservations are created exclusively through reservation templates. Service parameters are defined at the time the service template is created. The administrator can override these parameter values when making a reservation with Reservation Manager. The parameters are used to configure services properly when systems are allocated.

Service templates for Reservation Manager are automatically created when they are created in the Server Automation interface. You can also create or modify service templates with the Reservation Manager interface.

Note: You do not need to create system images for Services.

To create a service template in Reservation Manager

1. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your reservation templates.
2. Select Create Service from the Actions drop-down menu.

The Create Service Reservation Template wizard opens. Follow the instructions in the wizard.

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 793)

[Create a Template for VMware Virtual Machines](#) (see page 774)

[Create a Template for IBM PowerVM Logical Partitions](#) (see page 786)

Create a Template for the Image Type

Administrators create one or more templates so that users can make reservations. You can 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, click Create from the action menu in the upper-right corner of the Manage your reservation templates option of the Reservation Manager, and complete the wizard.

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 799)). When a reservation is made using a group template, users are not allowed to specify the password.

To create group templates, complete the Create Group wizard in the Actions menu of the Manager your reservation templates option of Reservation Manager.

Use 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 795)

[Create VMware Customization Specification](#) (see page 796)

[Enable Static IP Addresses](#) (see page 796)

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`

Note: The default port is 8443.

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: 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 809) 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.

Enable Static IP Addresses

Enable static IP addresses by setting the Network Selection Allowed field value to True. Go to Administer Your Reservation Manager, Manage your configuration settings, and open the Reservations area to find the Network Selection Allowed field.

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 795)

[Create VMware Customization Specifications and Templates](#) (see page 796)

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, complete the settings in the Notifications area of the Manage your configuration settings option of the Reservation Manager.

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, complete the following settings in the Notifications area of the Manage your configuration settings option of Reservation Manager:

- Reservation Ready Text
- Reservation Ready Text Contains Password

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, go to Manage your configuration settings in Reservation Manager and set the Reservation Expiration Warning Times settings.

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. As an administrator on the Reservation Manager home page, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Notifications area, change the value and click OK:

Notify End User On Job Failure

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. As an administrator on the Reservation Manager home page, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Notifications area, specify a value and click OK.

Task Status Update Timeout

The configuration change takes effect when the next reservation is made.

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. As an administrator on the Reservation Manager home page, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the General area:

Announcements

3. 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.

4. Save and close the file.

The announcement change is made when you restart the browser.

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, select the desired settings in the Reservation Manager, Manager your resource pools option. Double-click the desired resource pool and select or clear the fields

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.

User Management

You must configure and specify the limits and capabilities of your users.

This sections contains procedures for user configuration.

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.
- By default, all users are members of the Public org unit. Do not explicitly add users to the Public org unit. Do add resource access to the Public org unit in order to enable all CA EEM global users to make reservation requests. Public org unit privileges are in addition to specific org unit privileges.

Important! You can disable the default Public org unit. If you disable the Public org unit, you must explicitly specify reservation privileges for each org unit.
- 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.

Note: The LIKE operator includes an implicit wildcard search. Use the EQUAL operator to get an exact match.

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.

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 are used only by members of that tenant.

- Grant access to these objects to the tenant in Reservation Manager so the tenant administrator can 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. However, when the super administrator creates a tenant, Reservation Manager automatically creates a single organizational unit for the tenant. This new default organizational unit is given the same name as the tenant. All resources that are assigned to the tenant during the tenant creation interview are assigned to the default organizational unit.

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, complete the wizard in the Add option of the Action menu of the Manager your tenants option of Reservation Manager.

Edit a Tenant

The super administrator can edit tenant properties, administrators and end users, and access to resource pools.

To edit a tenant, click the tenant in the Manager your tenants option of the Reservation Manager, and update the information in the tabs presented

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 cannot modify some pool settings. Tenant administrators cannot modify CA ITCM (ITCM Domain Manager and Scalability Server, vCenter (Folder placement), and Amazon Cloud (Keypair Name and Network Selection). However, tenant administrators can view these settings.

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.

Administration

You can perform several operations to configure, manage, 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.

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.

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 SDM 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 SDM 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 843)

[Set Automatic Cancellation of Unapproved Reservations](#) (see page 843)

Extending Reservations

You must extend current reservations before the reservation end date or the resources will be returned to the resource pool.

You extend a reservation by selecting Extend... from the reservation details Actions menu.

If the resources you requested are not available or your maximum allowed systems count has been reached, your request is denied. Your Administrator can override extension denials and allow your extension.

Administrators are notified when availability checking determines that sufficient resources are not available to extend reservations beyond end dates. The administrator determines if it is safe to extend the reservation and override the availability checking warnings.

Resource Allocation and Forecast Charts

Resource allocation and forecast charts allow Reservation Manager administrators to view resource allocations for analysis.

- Allocation charts show the current usage of resources.
- Forecast charts show anticipated usage which is based on scheduled reservations for a specified time frame.

To view resource allocations

1. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The list of resource pools appears.

2. Mark the checkbox for a resource pool and click Resource Allocation from the Actions drop-down menu.

A Resource Allocation chart appears.

Note: Click the Show Chart or Show Table button in the table header to display the chart or the table.

3. Show the chart and select the criteria that you want from the Current Chart View drop-down options
4. Hover over or click a bar in the chart.

The host information appears.

Note: In the table view, hover over or select the host for the same information.

For forecasts, 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 allocation forecasts

1. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The list of resource pools appears.

2. Mark the checkbox for a resource pool and click Allocation Forecast from the Actions drop-down menu.

A Resource Allocation Forecast chart appears.

3. 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.

4. Click Refresh Chart.

The chart displays resource allocation, which is based on the values you chose.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The list of resource pools appears.

2. 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.

3. Fill in the fields and press OK.

The report opens in a separate window.

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 you allow enough time for all reservation tasks that are in progress to be completed before the services are stopped. This action ensures that 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

Block web access (maintenance mode)

4. 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click View all reservations.

The View Reservations page opens. The Job Status column shows that a job is suspended.

2. Click the check box next to the reservation you want to restart, click the Actions menu, and click Details.
 3. Click the link in the Job Status column to resume the job.
 4. Select the suspended job, click the Actions menu, and select Restart Selected Task.
- The job is restarted.

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 818).

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 844).

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 818)

[Configure Chargeback for Storage Tiers](#) (see page 820)

[Configure Chargeback by Tier for Amazon EC2](#) (see page 821)

[Configure Chargeback by Tier for IBM PowerVM Logical Partitions](#) (see page 821)

[Select a Chargeback Tier for IBM PowerVM Logical Partitions](#) (see page 822)

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. 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

Chargeback Calculation Frequency

Chargeback Retention

Chargeback Is Enabled

The configuration changes take effect when you log out and log back in.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your chargeback models, select a model, and then select Details from the Actions menu.

A pane with fields for resource chargeback appears.

2. 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 820).

Note: Base Disk (GB) is not supported for AppLogic applications. See [Configure Chargeback by Resource for AppLogic Applications](#) (see page 823).

3. (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 844).

Note: Surcharges for Additional Disk are not supported for AppLogic applications. For more information, see [Configure Chargeback by Resource for AppLogic Applications](#) (see page 823).

4. 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 844).

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your chargeback models, select a model, and then select Details from the Actions menu.

A pane with fields for resource chargeback appears.
2. Enter values in the following fields:
 - **Base Hourly Rate** (example, \$0.15)
 - **Base Memory (GB)** (example, 1)
 - **Base CPU** (example, 1)
3. (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)
4. Enter hourly rates for disk space on storage tiers.
5. Click OK.

More information:

[Allow Users to Select Storage Tiers](#) (see page 844)

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your chargeback models, and Amazon Cloud Instances.

A pane with fields for tier chargeback appears.

2. 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
3. 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your chargeback models, and IBM Logical Partitions.
A pane with fields for tier chargeback appears.
2. 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
3. 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your reservation templates.
The Reservation Templates page opens. This page lists the public templates that users can select when reserving systems.
2. Double-click a template for logical partitions.
The Reservation Templates Details page opens.
3. 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 821).

Configure Chargeback by Resource for AppLogic Applications

Reservations on AppLogic applications use *chargeback by resource*.

When assigning chargeback rates, 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 is 0 (no chargeback).

To configure chargeback by resource for AppLogic applications

1. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your chargeback models, and AppLogic applications.

A pane with fields for tier chargeback appears.

2. Provide base hourly rate and memory resource thresholds as described in the [Configure Chargeback by Resource section](#) (see page 818).

Note: CPU resources are measured in "shares" as opposed to number of CPUs. Choose a base amount to 2 decimal places (for example, 0.01). The surcharge rate will be applied for each 0.01 increment above the base amount.

Note: Charging for disk resources is not supported for AppLogic Applications.

Note: Bandwidth is measured in MBs. The surcharge rate will be charged for each MB above the base amount.

3. Click OK.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Reservations area.
Display Chargeback Cost Information
3. Enter a value and click OK.
The change takes effect when you restart the browser.

Customization

In addition to standard configuration and setup, additional customization is available.

This section contains procedures on how to customize your services and resources.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.
2. 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

Administrator Contact Url

The configuration change is made after you restart the browser.

Use the Reservation Manager Mobile App

The CA Reservation Manager mobile app enables you to perform reservation tasks from your mobile device.

You can download the mobile app to your iPhone or iPad.

Search for “CA Mobile Reservation Manager” in the Apple App Store, and download the application for free.

The mobile app gives you the following features:

My Systems

Lists the systems you have reserved for now and in the future.

Recents

Shows past reservations you have made and allows you to create new reservations.

Announcements

Shows the Reservation Manager home announcement screen information.

Settings

Contains configuration settings for the mobile app.

Follow these steps:

1. Start the mobile app
2. Supply the following in the Settings option:
 - The Reservation Manager server name. This is the machine where Reservation Manager was installed.
 - The Port Number. Default is 443.
 - The Reservation Manager User ID and Password.
 - Your email address.
3. Tap Test Connection to verify the information.
4. If you are a member of more than one organizational unit, tap and select the organizational unit you would like to use.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the General area:

End User Help

3. Enter a value and click OK.

The change takes effect when you restart the browser.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.
2. 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

Reserve VM Access

Create AMI Access

View Template Access

View Reservation Access

View Machine Access

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Quick Start area.

Use Short Descriptions

3. Change the value and click OK.

The configuration change takes effect when you restart the browser.

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 827)

[Templates and Email Types](#) (see page 828)

[Substitution Variables](#) (see page 828)

[Email Types and Categories](#) (see page 831)

[Conditional Substitution](#) (see page 833)

[Configure Parameters for Email Notification](#) (see page 783)

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, such as the following:

<p>System Name: %SYSTEM_NAME%</p>

Note: All substitution variable values can be set at the global level and some can be set at the tenant level. If a value is specified at the tenant level, the global value is ignored.

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.
%TENANT%	The name of the tenant.
%TENANTID%	An alternate tenant identification (short or abbreviated)
%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.
%VCSERVERNAME%	Name of the VC Server that is hosting the reserved system.
%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 833)

Email Types and Categories

The following table shows email notification types and the message categories associated with them. For additional explanation of message categories and associations, see the [Substitution Variables section](#) (see page 828).

Yes in the table below, indicates that the email notification type is associated with the message category.

Email Notification 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_EXTENDED	Yes	No	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
REVERT_TO_SNAPSHOT_FOR_RESERVED_SYSTEM_COMPLETED	Yes	Yes	No	No
SCHEDULER_PROCESSING_RESUMED	No	No	No	Yes
SCHEDULER_PROCESSING_SUSPENDED	No	No	No	Yes
SNAPSHOT_FOR_RESERVED_SYSTEM_CREATED	Yes	Yes	No	No
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 blank field:

```
<td>VM Console URL</td>
<td>@%VMCONSOLEURL%</td>
```

Enable or Disable Inheritance of Resources from the Public Org Unit

Administrators can specify whether to make the resources that are assigned to the Public org unit available to all users.

A value of True for this configuration setting indicates that all organizational units are automatically granted access to all resources to which the Public organizational unit has access.

To set the Public org unit option

1. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the General area:

Organizational units are granted access to public resources

3. Change the setting in the Value field as desired and click OK.

The configuration change is made when you restart the browser.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the General area:

Home Welcome Text

3. Change the setting in the Value field and click OK.

The configuration change is made when 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the General area:

Data Timeout

3. Enter a value and click OK.

The change takes effect when you restart the browser.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. 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

Virtual CPU Limit

Virtual Memory Limit

Virtual Disk Limit

Virtual Disk Space Limit

The configuration changes take effect when the next reservation is made.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The Resource Pools page opens, and lists existing pools.

2. Double-click an ESX server or cluster.

The Resource Pool details page opens, with the Properties tab displayed.

3. Select the following field and enter a percentage:

Allow memory overcommitment. Percent:

4. Click OK.

Memory is overcommitted.

Specify a Folder for VMware 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your resource pools.

The Resource Pools page opens, and lists existing pools.

2. Double-click a VMware resource pool.

The Resource Pool details page opens, with the Properties tab displayed.

3. 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.

4. Click OK.

New VMs are placed in the specified folder.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Virtual Machines area, set the value, and click OK.

Virtual Network Interface Limit

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.
2. Click the following link in the Virtual Machines area, set the value, and click OK:.

Add Virtual Machine To Service

The configuration change takes effect when the next reservation is made.

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. Project ID can be used to 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.

2. 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

Network Selection Allowed

Project Id Is Mandatory

Specify Requirements Allowed

The configuration changes take effect when the next reservation is made.

More information:

[Allow Alternate Selection](#) (see page 838)

[Configure Chargeback Display](#) (see page 823)

[Override Automatic Selection](#) (see page 839)

[Specify Memory and CPU Selections](#) (see page 839)

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Reservations area:
Specify Alternate System Specifications Allowed
3. Enter a value and click OK.

The configuration change takes place when the next reservation is made.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.
2. Click the following link in the Reservations area:
Machine Selection Allowed
3. 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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.

2. 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

Reservation CPU Selections

The changes take effect when the next reservation is made.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. 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

Service Upper Threshold

Service Lag

Service Priority

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.
2. 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

Filesystem Growth Allowance for Snapshots

Manage Snapshot Reserve

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.
2. Click the following link in the CA ITCM Integration area, and click OK.

Software Deployment Enabled

The change takes effect when Reservation Manager is restarted.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.
The Configure Settings page opens.
2. Click the following links in the Physical Provisioning area. Each one opens a dialog in which you can change the value and click OK.

Cpu Count Reservation Weighting

Disk Space Reservation Weighting

Memory Reservation Weighting

The configuration change takes effect when the next reservation is made.

More information:

[How the Low Cost Algorithm Works](#) (see page 842)

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.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Approvals area:

Approval Required Notification Time

3. Change the setting in the Value field and click OK.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. 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

Automatically Cancel Unapproved Reservations Time Delay

The configuration change takes effect when the next reservation is made.

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. As administrator in Reservation Manager, under Administer Your Reservation Manager, click Manage your configuration settings.

The Configure Settings page opens.

2. Click the following link in the Virtual Machines area:

Allow Data Store Tier Selection

3. 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 820)

Chapter 13: Scalability Best Practices

This section contains the following topics:

[Scalability Overview](#) (see page 845)

[Hardware Specifications](#) (see page 846)

[Database Considerations](#) (see page 846)

[Network Considerations](#) (see page 847)

[Remote Deployment and Policy Configuration Overview](#) (see page 847)

[Scalability Recommendations](#) (see page 849)

Scalability Overview

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
- Monitoring of IBM LPAR Environments
- Monitoring of Oracle Solaris Zones Environments
- Deployment of SystemEDGE and other Monitoring Software
- Initial and on-going configuration of SystemEDGE

The following sections are included:

1. [Remote Deployment and Policy Configuration Overview](#) (see page 847)
2. [Hardware Specifications](#) (see page 846)
3. [Database Considerations](#) (see page 846)
4. [Network Considerations](#) (see page 847)
5. [Scalability Recommendations and Limitations](#) (see page 849)
6. [Scalability Use Cases](#) (see page 856)

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.
- LPAR AIM Monitoring Server: Dual Core Processor: 2.6-GHz CPU, 4-GB RAM, 100-GB disk.
- Solaris Zones AIM Monitoring Server: Dual Core Processor: 2.6-GHz CPU, 4-GB RAM, 100-GB disk.
- Target System: 1-GHz Single Core/Processor/Virtual Processor, 512-MB RAM, 2-GB disk, Single 100-Mb/sec Ethernet.

Note: 50 percent is the maximum usage allowed for CPU and memory.

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 CA Server Automation 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 that are collocated with the target systems.


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 that are based on 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 of 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.



Scalability Recommendations

This section provides information about scalability recommendations and limitations.

Consider the following information:

- [Monitoring of VMware Environments](#) (see page 849)
- [CA Server Automation Management of VMware Environments](#) (see page 850)
- [Remote Deployment and Policy Configuration Recommendations](#) (see page 854)
- [Domain Server Recommendations](#) (see page 856)
- [Distribution Server Recommendations](#) (see page 856)
- [Scalability Use Cases](#) (see page 856)

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 CPU usage:

- The dynamic nature of the vCenter Servers being monitored.

The level of activity of vCenter Servers influences 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 the 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 that is 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 20,000 VMs.

Scalability Limitation in Terms of Monitored Servers

The vCenter AIM functions in a multiple vCenter Server environment. The vCenter AIM framework results in a slight CPU usage reduction with fewer VMs per vCenter Server. For example, CPU usage is lower for three vCenter Servers, each with 2,000 VMs, than for a single vCenter Server with 6,000 VMs.

Sometimes 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 CA Server Automation manager does far more than simply monitor and publish vCenter data. The manager stores and manages historical data, performs active operations against the vCenter, runs automation policy, 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

Because most 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 that is required to cache the data and handle messaging increases rapidly. We recommend increasing the memory of the 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 an additional CPU (3.2 GHz or larger) for moderately large managed environments to improve the responsiveness of automated processes.
- **Operating System Limitations:** Most 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 a 64-bit processor and an 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: 2,500 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: 8,000 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 actions 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 the 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, are considerably quicker. For example, 2,500 Computer Systems are typically published through vCenter AIM and the CA Server Automation manager caches them in approximately 5 minutes.

LPAR AIM Monitoring Recommendations

The LPAR AIM (Application Insight Module) is a pluggable component that is implemented within the SystemEDGE framework. As such, the data that it publishes is available to multiple SNMP Managers. Products such as CA Server Automation, eHealth, and Spectrum IM can leverage this data.

The following section specifies the scalability recommendations for LPAR AIM monitoring.

Scalability Recommendation for LPAR AIM Monitoring

One LPAR AIM can handle P systems environment configurations in the following range:

- Number of HMC Servers: 1 to 4
- Number of Power Systems per HMC Server: 2 to 10
- Number of Virtual I/O Servers per Power System: 1 to 2
- Number of LPARs per Power System: 10 to 100

General recommendations for LPAR AIM monitoring present a typical configuration of LPAR environment:

- Monitored Power Systems: up to 20
- Monitored VIO Servers: up to 30
- Monitored LPARs: up to 300

The AIM consumes CPU and Memory which is roughly a function of the number of LPARs. With a maximum of 300 LPARs the CPU consumption of the sysedge process would be less than 10 percent.

Note: The presented CPU consumption is valid for a dedicated SystemEDGE agent not running any other AIM.

With 300 LPARs added to the LPAR AIM, the sysedge process memory consumption increases by approximately 8 MB.

Solaris Zones AIM Monitoring Recommendations

The Solaris Zones AIM (Application Insight Module) is a pluggable component that is implemented within the SystemEDGE framework. As such, the data that it publishes is available to multiple SNMP Managers. Products such as CA Server Automation, eHealth, and Spectrum IM can leverage this data.

The following section specifies the scalability recommendations for Solaris Zones AIM monitoring.

Scalability Recommendation for Solaris Zones AIM Monitoring

General recommendations are to have one Zones AIM monitor the following configuration:

- Monitored Zones Servers: Up to 20
- Monitored Zones: Up to 1000

The AIM consumes CPU and Memory proportionally to the number of Zones. However the initial consumption of memory for the first couple of zones is higher and decreases with more zones being added.

With 1000 Zones the CPU consumption of the sysedge process is less than 5 percent.

Note: The presented CPU consumption is valid for a dedicated SystemEDGE agent not running any other AIM.

With 1000 Zones being added to the Zone AIM, the memory consumption of the sysedge process increases by approximately 20 MB.

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 in a 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 (slower than 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 action 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 then be applied directly to a Service.

- Test Deployment

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 10,000, we recommend running multiple instances of CA Server Automation.

Distribution Server Recommendations

A deployment and configuration distribution (scalability) server ensures 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 2,000 target systems for a typical 100-Mb/sec network environment.

Note: If you use Policy Configuration and you do not use Remote Deployment, each Distribution Server can scale up to 3,000 servers.

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 that best matches your environment.

Departmental Data Center

In this use case, monitoring is required for 1,000 systems, all contained within one data center. All systems are in one location, within a 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 completes, 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, we suggest aligning the monitored servers into Services, with approximately 200 servers per Service. You can align services 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, 10,000 systems, spread across multiple Data Centers, are being managed. The number of systems per data center varies from 500 through 2,500. The Data Centers are geographically distributed across multiple locations. The links between the Data Centers include 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 2,000 servers, we recommend installing a second Distribution Server.

Note: If you use Policy Configuration and you do not use Remote Deployment, each Distribution Server can scale up to 3,000 servers.

- **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 2,000 to 2,500 systems.

If the base policy is required to be changed, we recommend applying the policy to systems in batches of 2,000 to 2,500 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 21,000 Agents are being managed. These agents are spread across three data centers, with each data center containing from 2,000 through 10,000 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, so that each instance manages up to 10,000 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 and 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 that the SystemEDGE is configured to request this policy on registration.

For application-specific monitors, create templates that are 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 2,000 to 2,500 systems.

If the base policy is required to be changed, we recommend applying the policy to systems in batches of 2,000 to 2,500 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 you 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.

Appendix A: FIPS 140-2 Encryption

This section contains the following topics:

[FIPS Overview](#) (see page 863)

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: Configure AIMs with NodeCfgUtil

This section contains the following topics:

[NodeCfgUtil Overview](#) (see page 865)

[Configure AIMs with NodeCfgUtil in Dialog Mode](#) (see page 867)

[Configure AIMs with NodeCfgUtil in Command Mode](#) (see page 870)

NodeCfgUtil Overview

To configure the AIMs and discover virtual environments, do *one* of the following actions:

- Open the Administration tab from the user interface, change to Configuration, Provisioning, and select the appropriate server type to add credentials and configure the AIM. CA 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 that are 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.
- Citrix XenServer AIM can only connect to pool masters or standalone Citrix XenServers. Otherwise, the AIM does not work.
- To manage VMware vSphere, enter the credentials for the corresponding vCenter Servers.
- Install VMware Tools on all VMware VMs for the features to work properly. VMware Tools optimize the virtualization of VMs. Many features are available only if these tools are installed. For further information, see the VMware documentation.

To discover supported environments from an AIM Server:

1. 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.
2. Run the NodeCfgUtil.exe utility on a Windows AIM Server to update the configuration data for the corresponding AIMs.

The NodeCfgUtil.exe utility stores the data for each AIM in a file (for example, zone.cfg, vc.cfg, ...).
3. Open the user interface on the Manager Server and click Resources, Data Center in the navigation pane.
4. Right-click and select Management, Discover.

The discovery options appear.
5. Select one of the following actions:
 - Discover a system
 - Discover a network
The corresponding dialog opens.
6. Enter a system name of server that you want to manage. Alternatively, you can enter network properties for the discovery process. Click OK.

CA 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 278)

Configure AIMs with NodeCfgUtil in Dialog Mode

NodeCfgUtil.exe lets you modify the AIM configurations for IBM LPAR, IBM HACMP, Solaris Zones, VMware vCenter, VMware vCloud, Microsoft Clusters, Cisco UCS, and Citrix XenServer. The utility writes a configuration file for the corresponding AIM to the *sysedge_InstallPath\plugins\AIPCommon* directory. You can also use the NodeCfgUtil utility to edit or remove existing entries.

Use the utility in dialog mode to configure which nodes the appropriate AIMs manage.

Note: Run NodeCfgUtil.exe as Windows Administrator.

Follow these steps:

1. Log in as Administrator and open Windows Explorer on the computer on which the AIM is installed.
2. Change to the *SystemEDGE_InstallPath\plugins\AIPCommon* directory, and start NodeCfgUtil.exe.

NodeCfgUtil discovers and lists the installed AIMs in subsequent dialogs.

3. Enter 1 to add a new managed node.
4. Follow the on-screen instructions to complete the configuration. Each node requires a valid user name and password for authentication.

After the configuration, enter 0 to return to previous menus, or to exit the utility.

NodeCfgUtil writes a configuration file for Solaris Zones (zone.cfg), vCenter Server (vc.cfg), vCloud Director (vcloud.cfg), Microsoft Clusters (mscs.cfg), XenServer (xen.cfg), UCS (ucs.cfg), LPAR (lpar.cfg), or HACMP (hacmp.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.

Examples

The following example shows the Install Managed Node dialog for the myvc5 server that has been successfully added to the configuration of the vCenter AIM. The AIM is now ready to manage the vCenter Server. The vCenter AIM is a multi-instance AIM. So you can repeat this procedure and can add more vCenter Servers that you want to manage with this AIM.

```
***** Main MENU *****
1. Install Managed Node
2. Modify Managed Node

3. Remove Managed Node

0. Exit

*****

Enter choice:

**** Choose Managed Node ****
1. IBM LPAR
2. Solaris Zones
3. Citrix Xen
4. VMware vCenter
5. Cisco UCS
6. Microsoft Cluster
7. IBM HACMP
8. vCloud Director
0. Go Back to Previous Menu
*****

Enter choice: 4
Enter following information for the VMware vCenter Node...

(At any point to go back to the previous menu, Enter 'CTRL Q').

1. Server Name: myvc5
2. User Name: administrator
3. Password: *****
4. Port [default=443]:
5. Protocol [default=https]:

CAAC1016 Authenticating, please wait...
CAAC1019 Authentication SUCCESSFUL.
CAAC1023 Added Node Successfully.

Press any key to continue...
```

The following example shows the Managed Node dialog for the HMC1 server that has been successfully added to the configuration of the LPAR AIM. After the AIM discovers all Virtual I/O Servers that are related to the HMC server, they are visible in NodeCfgUtil and each one can be modified to specify its credentials. The AIM uses the credentials of the first fully configured VIO Server as default credentials for all VIO Servers not yet configured. Thus, it is sufficient to specify the credentials for only one VIO Server if all of them share credentials. Otherwise, it is necessary to configure each VIO Server with different credentials. The AIM is now ready to manage the HMC Server.

```
**** Choose Managed Node ****
1. IBM LPAR
0. Go Back to Previous Menu
*****
Enter choice: 1
List of existing entries...
1. hmc: HMC1.company.com
2. vio: ibm101.company.com
Select the entry to be modified (0 to go back to the previous menu): 2
Enter following information for the IBM LPAR Node...
(At any point to go back to the previous menu, Enter 'CTRL Q').
1. Server Name: ibm101
2. User Name: admin
3. Password: *****

CAAC1016 Authenticating, please wait...
CAAC1019 Authentication SUCCESSFUL.
CAAC1024 Modified Node Successfully.

Press any key to continue...
```

Configure AIMs with NodeCfgUtil in Command Mode

NodeCfgUtil.exe lets you modify the AIM configurations for IBM LPAR, IBM HACMP, Solaris Zones, VMware vCenter, VMware vCloud, Microsoft Clusters, Cisco UCS, and Citrix XenServer. The utility writes a configuration file for the corresponding AIM to the *sysedge_InstallPath\plugins\AIPCommon* directory. You can also use the NodeCfgUtil utility to edit or remove existing entries.

When you use the utility in command mode, you can only add managed nodes to an AIM configuration.

Note: Run NodeCfgUtil.exe as Windows Administrator.

This command has the following format:

- (1) nodecfgutil -help
- (2) nodecfgutil {lpar|zone|mscs} -u *user* -p *password* -h {pvmname|hostname|cluster_name}
- (3) nodecfgutil {vc|ucs} -u *user* -p *password* -h *hostname* -t *port* -c *protocol*
- (4) nodecfgutil {xen|vcloud} -u *user* -p *passwd* -h *hostname*
- (5) nodecfgutil hacm -u *user* -p *password* -h {cluster_name|hostname} [-t *port*]

-help

Displays usage information about the console.

lpar|ucs|vc|zone|mscs|xen|vcloud|hacmp

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 cluster.

-t *port*

(Optional) Specifies the port number.

-c protocol

(vCenter, UCS only) Specifies the protocol (HTTP, https).

Return codes: 0 success, -1 failure

-o option

Specifies the option for providing management.

0

Specifies the entire domain for monitoring.

1

Specifies a specific host of the domain for monitoring.

Follow these steps:

1. Open a command prompt on the system on which the AIM is installed.

The command prompt appears.

2. Enter *one* of the following commands:

(1) `nodecfgutil -help`

(2) `nodecfgutil {lpar|zone|mscs} -u user -p password -h {pvmname|hostname|mscs}`

(3) `nodecfgutil {vc|ucs} -u user -p password -h hostname -t port -c protocol`

(4) `nodecfgutil {xen|vcloud} -u user -p passwd -h hostname`

(5) `nodecfgutil hacmp -u user -p password -h {cluster_name|hostname} [-t port]`

(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 vCenter or Cisco UCS.

(4) Authenticates and stores the passed credentials for Citrix XenServer or VMware vCloud.

(5) Authenticates and stores the passed credentials for IBM HACMP.

Appendix C: Troubleshooting

This section contains the following topics:

[CA Server Automation Troubleshooting](#) (see page 873)

CA Server Automation Troubleshooting

This section contains troubleshooting topics for CA Server Automation.

Note: If you receive a security certificate request, bypass it and continue. To eliminate these messages, acquire a certificate from the vendor of your choice and apply it to the server. For information about installing security certificates, see the Apache Tomcat website.

More information:

[Adjusting Poll Interval Settings for Solaris Zones Environments](#) (see page 875)
[Attributes Show a Value of Zero](#) (see page 875)
[Browsers Do Not Display Consecutive Spaces in Events](#) (see page 875)
[CA NSM Fails to Receive Traps on Windows](#) (see page 876)
[Cisco UCS Folder Does Not Display in UI](#) (see page 876)
[DB Transaction Log Sizes Increase Unexpectedly](#) (see page 876)
[Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero](#) (see page 877)
[Domain Server is not available](#) (see page 877)
[Empty Task ID for the dpmvc virtualswitch Command](#) (see page 878)
[Local and Remote Monitors Do Not Show the Same Values](#) (see page 878)
[Remote Deployment to Solaris Lists SPARC and x86 Systems](#) (see page 879)
[Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear](#) (see page 879)
[Resetting the vCenter Server Password Causes Data Collection to Fail](#) (see page 879)
[Solaris Zones AIM Reset if a Monitored System is Down](#) (see page 880)
[Status Icon of Component Shows Not Configured](#) (see page 880)
[Unable to Connect to Microsoft SQL Server](#) (see page 880)
[Upgrading SystemEDGE](#) (see page 881)
[User Interface is Unresponsive on Provisioning and Policy Screens](#) (see page 881)
[vCenter Server AIM Attributes Show Zero](#) (see page 881)
[VM Usage Values Do Not Update Immediately After Power Down](#) (see page 882)
[Blank Query Results Tab after Upgrade](#) (see page 883)
[Accessing the CA Process Automation Server Requires Credentials After Configuration](#) (see page 884)
[CA DSM Functionality is Not Fully Supported](#) (see page 884)
[CA Configuration Automation Agent Stops During Installation](#) (see page 885)
[Deleted OS Images from CA ITCM are Not Deleted from CA Server Automation](#) (see page 885)
[Discovering Large Networks](#) (see page 886)
[Discovery Does Not Identify Operating System](#) (see page 886)
[Duplicated Zone Entries in the Managed Folder](#) (see page 887)
[Error When Installing CA DSM Agent and Asset Management Plug-in](#) (see page 887)
[ESX Job Status is Current But OS Installation Not Complete](#) (see page 888)
[ESX Machines Fail to Discover](#) (see page 888)
[Log In with Different User Credentials on the Same Computer Using IE8](#) (see page 888)
[No Cisco UCS Manager in Explore Pane](#) (see page 889)
[New System Name is not Displayed](#) (see page 890)
[OpenSSL Software Compatibility Issues](#) (see page 890)
[Password Changes May Cause Authentication Errors](#) (see page 890)
[Rapid Server Imaging \(RSI\) Troubleshooting](#) (see page 893)
[Reservation Manager Troubleshooting](#) (see page 899)
[Scheduled Jobs do not Run](#) (see page 905)
[CA SDM Exception Error](#) (see page 906)
[Software Delivery Adapter Errors](#) (see page 907)
[SSP - The Home content does not display in Internet Explorer 9](#) (see page 908)
[vCenter Server Folder Does Not Display in UI](#) (see page 908)

[VM Reservation Fails: Could Not Find Computer UID for Software Delivery](#) (see page 909)

[VMs Not Being Discovered](#) (see page 909)

Adjusting Poll Interval Settings for Solaris Zones Environments

Symptom:

I do not know how to adjust poll interval settings for Solaris Zones environments.

Solution:

Increase the poll interval of the Solaris Zones AIM if the number of systems and zones increases. For example, if the host and zone count is greater than 100, set the default poll interval to 240.

Attributes Show a Value of Zero

Symptom:

Attributes show a value of zero.

Solution:

SystemEDGE rounds values down to zero, if they are smaller than one.

Note: The zoneAimStatHostDiskSvc MIB attribute always shows a value of zero.

Browsers Do Not Display Consecutive Spaces in Events

Symptom:

Browsers do not display more than one consecutive space character in event descriptions.

Solution:

Browsers do not display more than one consecutive space, because additional spaces are truncated according to the HTML specification. Use caution when cutting and pasting events from the browser into rules as the event descriptions can differ.

CA NSM Fails to Receive Traps on Windows

Symptom:

After a few days of receiving traps from CA 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>.

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 nodecgtutil.exe to verify access information for the UCS Manager node. Use a MIB Browser to verify data polling from UCS Manager. If UCS access information is not populated, review the sysedge log for additional information.

DB Transaction Log Sizes Increase Unexpectedly

Symptom:

In data centers with numerous managed objects, configuration changes, and metrics data collection activities, the Management DB and Performance DB transaction logs can increase unexpectedly. This issue can cause disk space to become low in environments with limited resources.

Solution:

To resolve this issue, see the KB article on the Microsoft Support website <http://support.microsoft.com/kb/873235>

The transaction log files, aom2.ldf and dpm.ldf, are located in the directory C:\Program Files\Microsoft SQL Server\...\MSSQL\Data in default Microsoft SQL Server installations.

Note: If the database log file is reduced in size, restart the Apache service to improve performance.

Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero

Symptom:

Some Solaris Zones AIM MIB values always show N/A or zero.

Solution:

These MIB attributes of the Solaris Zones AIM are deprecated and remain for backward compatibility. The deprecated MIB attributes are:

- zoneAimStatHostDiskMode
- zoneAimStatProcessorSetContainerList
- zoneAimStatProcessorSetResourcepoolId
- zoneAimStatProcessorSetResourcePoolIdList
- zoneAimStatProcessorSetResourcepoolName
- zoneAimStatProjectFSSEnabled
- zoneAimStatResourcePoolContainerList

Domain Server is not available

Symptom:

Domain server is unavailable, stopped, nonfunctional or not servicing request and Service Controller (SC) shows that the component is up and running.

Solution:

This behavior is due to a database connection failure or an AIM password expiry and can potentially impact the behavior of Policy Configuration and Remote Deployment components. Monitoring Support Service Web Service (ISM) monitors the functionality of the Domain Server periodically. When ISM identifies an unexpected behavior, the user is notified about the changes in the status with the message: *The CA SM Domain Service is down or not responding.*

You can monitor the status with the following command:

```
Caaipscutil /status /id=ISM /user=<user> /password=<password>
```

Administration panel indicates the Domain Server status to make sure the infrastructure state and functionality.

Empty Task ID for the dpmvc virtualswitch Command

Symptom:

When I run the dpmvc virtualswitch command, the result shows an empty task ID.

Solution:

This operation does not run asynchronously, and the result gets back immediately. However, the PMM treats the operation as a tasked operation. Therefore the response contains a task ID, but it is always an empty string ("").

For example, when you run the following command from the CLI, you get an empty task ID:

```
dpmvc virtualswitch -vs_add -vc_server MYVC5 -switch_name XYZ  
-esx_host_name MYESX -ws_user admin -ws_password ca_admin
```

CLI output:

```
...  
SC URL: https://VASManager/aip/sc  
VC URL: https://VASManager:443/aip/vc  
Task ID:  
Command execution successful
```

Other commands like dpmvc faulttolerance or dpmvc distributedswitch run asynchronously and you get a task ID.

Local and Remote Monitors Do Not Show the Same Values

Symptom:

Local and remote monitors do not show the same values for the same attributes at the same time.

Solution:

For seamless local and remote monitoring, the monitored object names have been chosen identical. However, the APIs used are different and can return different values.

SystemEDGE on a remote machine runs independently from the RM AIM on the server, and the start point of their poll schedulers cannot be synchronized.

The monitored metrics are highly volatile. For example, two samples can differ with a high probability.

Remote Deployment to Solaris Lists SPARC and x86 Systems

Symptom:

The computers listed in the Deployment UI are typically filtered to the chosen operating environment for which you are deploying. However, you can see computers other than the chosen operating environment listed under the following situations:

- When you deploy to either a Solaris x86 or a Solaris SPARC server, the servers listed are for all Solaris architectures regardless of whether you selected Solaris x86 or Solaris SPARC as the target operating environment.
- When you deploy to any computer that is unclassified.

Solution:

Verify that the target computer matches the chosen agent architecture for a successful deployment. If you proceed by selecting all computers listed, deployment succeeds for the matching architectures and fails on mismatched architectures.

Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear

Symptom:

When I remove a vCenter Server from management, objects of another managed vCenter Server disappear unexpectedly.

Solution:

To avoid product management issues, do not install the vCenter AIM on a VM that is managing another vCenter Server. If you remove the monitoring and management of the vCenter associated with that VM from CA Server Automation, it removes the objects associated with the vCenter including the VM system that is running the AIM.

Resetting the vCenter Server Password Causes Data Collection to Fail

Symptom:

After resetting the VMware vCenter Server password for the user that CA 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.

Status Icon of Component Shows Not Configured

Symptom:

After CA Server Automation installs a component, the status icon of this component shows  (Not configured). This status appears if CA Server Automation registered a component that is connected to an unconfigured server.

Solution:

To change the status of the component to ready, add the missing Server connection settings and validation.

Unable to Connect to Microsoft SQL Server

Symptom:

Attempts to authenticate credentials to a Microsoft SQL Server Evaluation Edition fail during product installation. The error message, Failed to establish connection to MSSQL displays.

Solution:

This issue occurs because TCP/IP is disabled by default on the Evaluation Edition. Enable TCP/IP.

Upgrading SystemEDGE

Symptom:

When I upgrade SystemEDGE to 5.7.0, the AIMs of the previous CA Server Automation release do not run.

Solution:

Upgrade Advanced Encryption and all AIMs to CA Server Automation Release 12.7. SystemEDGE 5.7.0 does not load AIMs of previous CA Server Automation releases.

User Interface is Unresponsive on Provisioning and Policy Screens

Symptom:

If the database server is restarted while you are on the Provisioning page or Policy page, the user interface goes blank or is unresponsive.

Solution:

Log out of the CA Server Automation user interface and log back in.

vCenter Server AIM Attributes Show Zero

Symptom:

vCenter Server Attributes show zero.

Solution:

The following object values are only retrievable when the vCenter Server AIM is installed on the local vCenter Server instance. When the AIM is remote, these parameters show zero (0).

- vmvcAimStatServerCPUUsage [1.3.6.1.4.1.546.16.52.2.2.12.0]
- vmvcAimStatServerMemUsage [1.3.6.1.4.1.546.16.52.2.2.17.0]
- vmvcAimStatServerTotalPhysMem [1.3.6.1.4.1.546.16.52.2.2.18.0]
- vmvcAimStatServerUsedPhysMem [1.3.6.1.4.1.546.16.52.2.2.19.0]

VM Usage Values Do Not Update Immediately After Power Down

Symptom:

VM usage values do not update immediately after power down.

Solution:

After VMs are powered off, usage values do not drop to 0 until the next successful poll. Polling can take up to 5 minutes, which is the default data collection and recording interval.

Blank Query Results Tab after Upgrade

Symptom:

Remote Monitoring Query results show blank values after the upgrade.

Solution:

RM PMM require remote system names to comply with in Fully Qualified Domain Name (FQDN) notation while adding systems. However, the RM AIM leaves existing system non-FQDN names. This name mismatch shows blank query results. You can fix the name mismatch as following:

Before Upgrade Conversion

- Log in to refresh UI and remove all non-FQDN systems from Remote Monitoring.
Deletes any associated systems, queries, instances, and monitors from both the manager (database); and the SystemEDGE agent with RM AIM plugin.
- Re-add these systems using FQDN notation and specify the same configuration sets to re-create associated queries, instances, and monitors.
- Upgrade.

After Upgrade Conversion

- Log in to the SystemEDGE agent machine and run the Refresh RM AIM plugin.
- Locate the `rmonwbem.cf` file containing the current Remote Monitoring configuration in the data directory path and make a copy of this file. For example, save as `rmonwbem-upgrade.cf`
- Log in to the Refresh UI and remove all non-FQDN systems from Remote Monitoring.
Deletes any associated systems, queries, instances, and monitors from both the manager (database) and the RM AIM agent machine.
- Now upgrade. On the agent machine, run the `rmonwatch add` command with `rmonwbem-upgrade.cf` as input file. This process re-adds all systems and associated queries, instances, and monitors with FQDN notation.

Note: The After Upgrade Conversion approach has the advantage to re-add the systems automatically, and configure systems from the UI.

Query results show the values after the upgrade conversion.

Accessing the CA Process Automation Server Requires Credentials 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.

CA DSM Functionality is Not Fully Supported

Symptom:

Not all CA ITCM functionality is available in CA Server Automation.

Solution:

Review the following unsupported features:

- GETIMAGE for ImageX
- Microsoft Automated Deployment Services (ADS)
- Cluster Fail Over
- Software Job Prioritization
- Platform Virtualization
- Windows Embedded for Point of Service (WEPOS)
- Backup and restore model computers
- Shut down computers when software delivery jobs are complete
- Custom administrator messages

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 CA ITCM are Not Deleted from CA Server Automation

Symptom:

After deleting all of the OS Images from CA ITCM, they no longer appear in CA ITCM, but they still display in CA Server Automation.

Solution:

To resolve this issue, delete the OS images from the Software Package Library as well as CA ITCM.

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.

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
|-- ...
```

Error When Installing CA DSM Agent and Asset Management Plug-in

Symptom:

I get the following error when I try to install the CA DSM Agent + Asset Management Plugin on a Windows server:

1619: This installation package could not be opened.

Solution:

This problem is a known limitation of CA ITCM. For the workaround, see the topic, Network Installation of MSI Package Fails in the CA Software Delivery Guide. The limitation for Windows operating systems is detailed in the Microsoft KB article: <http://support.microsoft.com/kb/2022222/>

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.

ESX Machines Fail to Discover

Symptom:

My ESX and ESXi machines that were provisioned using CA ITCM OSIM, fail to discover with the following error:

<machinename> failed discovery

Solution:

This issue is a limitation of VMware, which does not allow ESX or ESXi machines to take the assigned host name in the DHCP setting. To recover:

1. Using the VI client or the console, change the host name of ESX/ESXi machine to the name specified in CA Server Automation when the ESX/ESXi machine was provisioned.
2. Reboot the host.
3. Rediscover the host.

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 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.

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.

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.

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 user interface. 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 user interface. 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.

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 CA Software Delivery and Rapid Server Imaging (RSI) servers are installed, provisioning tasks (for example, offline capture, capture driverset, deploy) can fail with the following error:

Check server for boot configuration errors. (Is network boot enabled?)

This occurs when the RSI server needs to answer PXE boot requests, but the CA-Unicenter ManagedPC Boot Server answers them instead.

Solution:

To configure the CA-Unicenter ManagedPC Boot Server to ignore PXE boot requests, follow these steps:

1. Log on to the CA ITCM user interface.
2. From Explorer, navigate to Control Panel, Configuration, Configuration Policy.
3. Right-click on Default Computer Policy, and select Un-Seal.
4. Navigate to DSM Explorer, Scalability Server, OSIM, ManagedPC, Server.
5. In the right panel, select "User Answer Control List" and change the Value to 2 (two).
6. 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*.

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.

Listing OS Types, Depots, or Networks Returns No Results

Symptom:

After RSI installation, listing OS types, depots, or networks returns no results.

Solution:

Verify that the portmap and nfs services are running on the RSI server. If not, start the services and make sure that they are configured to start automatically on system boot. Uninstall RSI and run the installation again.

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\productname\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*.

Installation With Rapid Server Imaging Fails

Symptom:

When I upgrade an existing RSI server, it fails without an indication of the problem.

Solution:

If you are installing the Rapid Server Imaging (RSI) component and the installation fails, follow these steps:

1. Verify that you have a valid Racemi DynaCenter license.
Run "rshowlicense" on the RSI server to check the license.
2. Verify that the Apache HTTP Server service is running on the RSI server.

Reservation Manager Troubleshooting

This section contains troubleshooting topics for Reservation Manager.

More information:

[Amazon Machine Images Are Not Available for Selection](#) (see page 900)
[Chargeback Calculations Are Lower or Higher Than Reservation Amounts](#) (see page 900)
[Installation Target Cannot be Resolved](#) (see page 901)
[No Resources Available Message When Requesting a VM](#) (see page 901)
[Password Change Causes Error Message](#) (see page 902)
[Tier Label Changes on VMware Datastore](#) (see page 902)
[Unable to Find Package Entries for Personality AutoDeploy](#) (see page 902)
[Unable to Retrieve Information from vCenter](#) (see page 903)
[VM Resources are not Available for Dates Requested](#) (see page 904)
[VM Reservation Fails Because of CPU Limitation](#) (see page 904)
[VM Reservation Fails in a Clustered Environment](#) (see page 905)

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 818).

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 ITCM 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.

No Resources Available Message When Requesting a VM

Symptom:

After selecting a virtual machine template, a user sees the warning message “Organizational unit org_unit has not been granted access to a resource pool that can be used to provision the selected virtual template.” 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 714) 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 714)

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 27) 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 835).

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 835).

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.

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 ensure that the job runs. Symptoms of this issue include the following:

- Dashboard messages do not appear.
- 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 `dmpolicycli.log`, `dpmccmcli.log`, and so on.

Note: For more information about log files, see the *Reference Guide*.

CA SDM Exception Error

Symptom:

When CA SDM is down, the connection status in the Administration, Configuration page displays the following message indicating that CA SDM cannot be connected:

"ServiceDeskClientAdapter validateSDUser() exception: ; nested exception is: java.net.ConnectException: Connection refused: connect".

Solution:

To resolve this issue, restart CA SDM.

Software Delivery Adapter Errors

Symptom:

When I add a resource in the user interface, I get the following error:

Software Delivery Adapter service not registered for ITCM *servername*

Solution:

When the CA Software Delivery and the SQL server are installed remotely (not on the manager system), set the Apache service credentials to Administrator. To fix this problem:

1. Change the Apache service credentials to Administrator.
2. Restart the Apache service.

Symptom:

When I provision a Software Delivery server, I get the following error:

Failed to add computer <name>. Error: SDAadapter_addComputer (*failed*). *Failed to add computer to DSM.*

Solution:

The most common reason is that you have not configured each CA ITCM instance for management by a single CA Server Automation instance. Another reason is that you may have exceeded the maximum number of connections supported in <asdm>. To fix this problem:

1. Go to Administration tab, Software Delivery servers.
2. Delete all entries for other CA Server Automation instances.
3. Reset IIS and CA ITCM.
4. Restart the caf services on the <asdm> server.
5. Restart the Apache service.
6. Reprovision the server.

If the problem still exists, contact Customer Support.

SSP - The Home content does not display in Internet Explorer 9

Symptom:

The Home content of Self-Service Portal does not display in Internet Explorer 9.

Solution:

Disable the Internet Explorer Enhanced Security Configuration. Open Server Manager and click Configure IE ESC in the Security Information pane.

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.

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 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.

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.

Glossary

access control list

The access control list or ACLs specify a space separated list of IP addresses to restrict community usage to those addresses only. If you leave the list blank, the agent grants access to any system that uses the associated community name.

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 *automation 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.

capped logical partition (LPAR)

A *capped logical partition* is a logical partition that cannot use more processor power than its assigned processing units. The capped partition is assigned a maximum capacity and guarantees a capacity that cannot be exceeded and cannot affect the overall behavior of the physical system.

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.

cluster

A *cluster* consists of two or more independent computer systems that are linked together and work as a single entity. Clustering is used for parallel processing, load balancing, and fault tolerance.

cmdlet

A *cmdlet* is a command that must start with the first non-white character in a line. Because of this restriction they can only be used standalone and not as part of a broader JavaScript expression. In particular, they cannot be used as an rvalue (right hand side operand of an assignment operator).
? is an example for an AutoShell cmdlet.

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.

CPU shares (VMware)

Shares are specified as natural numbers and express a proportional weight to each virtual machine.

Specifying shares makes sense only with regard to sibling virtual machines, vApps, or resource pools which have the same parent in the hierarchy. When you assign shares to a virtual machine, you always specify the priority for that virtual machine relative to other powered-on virtual machines.

For example, when competition occurs, a virtual machine with 2000 shares receives more CPU time than a virtual machine with 1000 shares. Shares are configured relative to the other shares; thus, only the proportion of shares matters, not the values of the shares. Three virtual machines with share values of 1000, 2000, 3000 act the same as three virtual machines with share values of 1, 2, 3. You can use any number scheme you prefer.

datacenter (VMware)

A *datacenter* serves as a container for your hosts, virtual machines, resource pools, or clusters. If their virtual configurations meet the requirements of specific departments, datacenters can represent organizational structures, such as geographical regions or separate business functions. You can also use datacenters to create isolated virtual environments for testing or to organize your environment.

datastore (VMware)

A *datastore* specifies a virtual representation of combinations of underlying physical storage resources in a datacenter. These physical storage resources can be provided by local disks on a server, by SAN disk arrays, and so on.

dual HMC (LPAR)

A *dual HMC* is a redundant Hardware Management Console (HMC) management system that provides high availability.

dvPort group (VMware)

Each VMware vNetwork Distributed Switch has one or more *dvPort Groups* assigned to it. dvPort Groups group multiple ports under a common configuration and provide a stable point for VMs connecting to labeled networks. A unique network label identifies each dvPort Group. The network labels are unique to the current datacenter.

A dvPort Group specifies port configuration options for each member port on a vNetwork Distributed Switch. dvPort Groups define how a connection is made to a network.

dvUplink port (VMware)

Distributed Virtual Uplinks (dvUplinks) provide a level of abstraction for the physical NICs (vmnics) on the ESX Hosts. Each physical NIC is mapped to a dvUplink. For each host associated with a VMware vNetwork Distributed Switch, each physical NIC (uplink) is assigned to the vNetwork Distributed Switch through one uplink port.

dynamic reconfiguration connector index, DRC-index (LPAR)

Each slot in a physical system unit has a *DRC-index* assigned to it. The deploy process requires this number to perform the actual creation of the LPARs. The management console (HMC) and the system uses this index to identify uniquely each slot on the system. The DRC-index is not assigned to a slot until the unit is powered up.

EC2

See *Amazon Elastic Compute Cloud (EC2)*.

entitled pool capacity (LPAR)

The *entitled pool capacity* of a shared processor pool defines the guaranteed processor capacity that is available to the group of partitions in the processor pool.

ESX/ESXi host (VMware)

An *ESX or ESXi host* is a physical computer that uses ESX or ESXi Server virtualization software to run virtual machines. Hosts provide the CPUs and memory resources that virtual machines use and give virtual machines access to storage and network connectivity.

fair share scheduler, FSS (Solaris)

The *fair share scheduler (FSS)* specifies a scheduler class that allocates CPU time based on shares. Shares define the portion of the system's CPU resources allocated to a project.

Fibre Channel, FC

Fibre Channel is a standardized gigabit-speed technology for transmitting data between computer devices. Fibre Channel is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

funclet

Funclets maintain the verbose command like syntax with optional clauses, stringification and so on. Funclets are often used like cmdlets, that is, standalone in a single line. They can return a value that can be processed as part of a broader expression.

global zone (Solaris)

A *global zone* is a zone that is contained on every Solaris system. If non-global zones exist on the system, the global zone is the default zone for the system and for systemwide administration.

Hardware Management Console, HMC (LPAR)

The *Hardware Management Console (HMC)* is an external appliance that is used to perform management tasks on IBM PowerVM Systems. HMC can be used to create or change logical partitions, including dynamically assigning resources to a partition. The HMC communicates with the server firmware layers of POWER Systems, providing a single point of control in large PowerVM environments.

IBM High Availability Cluster Multiprocessing (HACMP)

IBM High Availability Cluster Multiprocessing (HACMP) is a solution for building high-availability clusters on the AIX UNIX and Linux for IBM system p platforms.

Integrated Virtualization Manager (IVM, LPAR)

The *Integrated Virtualization Manager (IVM)* is an enhancement of the Virtual I/O Server (VIOS) and allows you to manage a single POWER System. IVM lets you create and manage LPARs. IVM enables management of VIOS functions and provides a web-based user interface.

Internet Small Computer Systems Interface, iSCSI

iSCSI is used to facilitate data transfers over intranets and to manage storage over large distances. iSCSI encapsulates SCSI commands in IP packets, which are routed just like any other IP packet on the network. When the IP packet reaches its destination, the iSCSI device removes the encapsulation and interprets the SCSI command.

Lightweight process, LWP (Solaris)

Lightweight processes (LWP) belong to the Solaris 10 kernel thread model. LWPs form the execution context for a user thread by associating a user thread with a kernel thread. In the Solaris 10 kernel, kernel services and tasks run as kernel threads. When a user thread is created, the associated LWP and kernel threads are also created and linked to the user thread. Resource control allows to set bounds for LWPs.

Logical memory block, LMB (LPAR)

A *logical memory block (LMB)* specifies the granularity of physical and logical memory assigned to an LPAR (for example: 256 MB).

Logical partition, LPAR

A *Logical Partition (LPAR)* is a subset of hardware resources, virtualized as a separate system. A physical system can be partitioned into multiple LPARs, each providing a separate operating system and applications. The number of logical partitions depends on the hardware configuration of the system. LPARs are typically used for different environments, such as databases, web servers, and so on. LPARs communicate as separate systems in the network.

Management Information Base (MIB)

A *Management Information Base (MIB)* is a data store that describes properties of a resource. MIBs are written in ASN.1, which is a language specified by a management standard and complies with OSI's structure of management information (SMI) standards for defining SNMP MIBs.

MIB objects, MIB attributes

A *MIB object* is an entity defined in a MIB that represents one or more resource objects or data items. MIB objects include groups, tables, and individual attributes, and they must be defined in accordance with the structure for management information (SMI).

Multiple Shared-Processor Pools (MSPPs)

The *multiple shared-processor pools (MSPPs)* is a capability that is supported on Power6 and later servers. This capability enables the creation of multiple processor pools to make allocation of the CPU resource more flexible.

Multiple Virtual I/O Servers

Multiple Virtual I/O servers offer capability that increase application availability by enabling Virtual I/O server maintenance without a downtime for the client partitions.

network installation manager, NIM (LPAR)

A *Network Installation Manager (NIM)* provides a central point of management for installing and maintaining AIX images for LPARs and individual servers. It also facilitates the installation of all of those instances from the same master image, from different images, from installation media or from a previous mksysb of that instance. An instance refers to an OS image, regardless of whether it is an LPAR or on a physical server.

non-global zone (Solaris)

A *non-global zone* provides a virtualized operating system environment in a single instance of the Solaris operating system. The Solaris Zones software partitioning technology virtualizes operating system services.

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.

platform management module, PMM

A *Platform Management Module (PMM)* is a web service which is responsible for providing connection and operational support for the corresponding environment. Supported environments are for example: VMware vSphere, Microsoft Hyper-V, IBM PowerVM, Solaris Zones, Cisco UCS, or Microsoft Cluster Service. A PMM manages connections with the servers of these environments, performs environment-related operations, retrieves data from the corresponding AIM, and populates the CA Server Automation Management Database.

policy based configuration

Policy based configuration provides the ability to create agent configuration policy that you can deploy to sets of managed machines in one operation.

poll interval

The *poll interval* is the length of time between consecutive polls of a resource group.

POWER processors (LPAR)

POWER processors are RISC-based and used as the CPU in many of IBM servers, mini-computers, workstations, and supercomputers.

processor pools (LPAR)

A *processor pool* is a set of physical processors that can be shared across different logical partitions.

processor set, pset (Solaris)

Processor sets define disjoint groups of CPUs. Each processor set can contain zero or more processors. It is a resource element in the resource pools configuration.

project (Solaris)

A *project* defines a container associated with a host. It is an abstraction layer that helps to organize and manage the collection of physical system resources.

Projects are collections of tasks, which are collections of processes. A new task is started in a project when a new session is opened by a login, cron, newtask, setproject or su command. Each process belongs to only one task, and each task belongs to only one project.

Projects and tasks are the basic entities which are used to identify workloads in the Solaris 10 operating system. A project is associated with a set of users and a set of groups. Users and groups can run their processes in the context of a project they are a member of, but they can be members of more than one project. The project is the basic entity against which the usage of resources can be restricted. The task is the entity to which a process is associated and the project is associated with a set of tasks.

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.

regular expressions

Regular expressions are text patterns used for matching. Regular expressions are strings that include a mix of plain text and special characters to indicate the kind of matching required.

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 Area Network, SAN

A *storage area network (SAN)* is an architecture to connect remote computer storage devices to servers in such a way that the devices appear as locally connected to the operating system.

storage 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 I/O Server, VIOS (LPAR)

A *Virtual I/O Server (VIOS)* is a special logical partition that is configured to own all physical I/O resources and provides its virtualization capabilities to other LPARs. LPARs access disk, network, and optical devices through the Virtual I/O Servers as virtual devices. Each PowerVM system with virtualized input output devices has one or more Virtual I/O Servers.

virtual 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

(

- (Optional) Add the SCVMM Management Instance to the CA Server Automation Manager • 433
- (Optional) Apply Policy and Template Updates to Servers and Verify Updates • 194
- (Optional) Manage the Base Policy and Templates for One or More Servers • 192
- (Optional) Reindex Monitors from Templates or a Policy • 189
- (Optional) Specify Access Control Lists at the Policy Level • 96
- (Optional) Specify SNMP Settings and Access Control Lists at the Server Level • 97
- (Optional) Update the Policy or Templates • 190, 193

A

- About Packages • 114
- Access Control • 546
- access control list • 911
- Access the CA EEM User Interface • 31
- Access the CA Process Automation User Interface • 515
- Access the Reservation Manager User Interface • 810
- Access the User Interface • 27
- Accessing the CA Process Automation Server Requires Credentials After Configuration • 884
- Action Types • 569
- Actions • 318
- Activate Logical Partition • 385
- Active Directory • 29
- Active Directory Password Expiration Causes Log in Issues • 890
- Add a Cisco UCS Manager Server • 448
- Add a Cisco UCS to the Manager • 454
- Add a Citrix XenServer Connection to the Manager • 393
- Add a Logical Partition for an IBM AIX Computer • 377
- Add a Microsoft Cluster Service to the Manager • 528
- Add a Monitor To SystemEDGE Policy • 226
- Add a New Hyper-V Server Connection to the Manager • 430
- Add a New SCVMM Server Connection to the Manager • 435
- Add a New vCenter Server Connection to the Manager • 282
- Add a New vCloud Director Connection to the Manager • 340
- Add a Solaris Zone • 419
- Add a Solaris Zones Connection to the Manager • 411
- Add a Tenant • 807
- Add a Test to SRM Policy • 243
- Add a Threshold Definition To SRM Policy • 247
- Add a Virtual Machine (Hyper-V Server) • 439
- Add a Virtual Machine (vCenter Server) • 321
- Add Amazon Machine Images to Inventory • 780
- Add an HMC or an IVM Server Connection to the Manager • 361
- Add an IBM AIX Client System Using a Resource Group • 730
- Add an IBM AIX System Using an Individual Resource • 732
- Add AppLogic Servers • 479
- Add Disk
 - VMware vCenter • 572
- Add MIB Extensions to a Template or a Policy • 187
- Add Monitors to a Template or the Policy • 175
- Add Network Interface
 - VMware vCenter • 574
- Add New Virtual Machines to a Service • 837
- Add or Remove Virtual Disk • 292
- Add or Remove Virtual Network Interface • 293
- Add Packages to the Image • 498
- Add Patches to the Image • 499
- Add Proxy Servers • 475
- Add Remote Systems for Monitoring • 555
- Add Server to Service • 575
- Add Server-level SNMP Settings • 103
- Add the AIM Instance for the vCenter Server • 286
- Add the AIM Instance for the vCloud Server • 344
- Add the Discovered Citrix XenServer AIM Instance • 395
- Add the Discovered MSCS AIM Instance • 530
- Add the LPAR AIM Instance • 363
- Add the Zones AIM Servers • 414
- Add Users To An Organizational Unit • 804

- Adjusting Poll Interval Settings for Solaris Zones Environments • 875
- Administration • 809
- Advantages of Remote Monitoring • 545
- Agent Configuration • 68
- Agent Configuration Without Write Community • 147
- Agent Policy Dashboard Views • 159
- Agent Remote Deployment • 85
- Agent Visualization • 75
- Agent-less Monitored Systems • 545
- Agent-less Monitoring • 543
- AIM • 911
- AIP • 911
- Allow Alternate Selection • 838
- Allow Users to Select Storage Tiers • 844
- Amazon EC2 Provisioning • 719
- Amazon Elastic Compute Cloud (EC2) • 911
- Amazon Machine Images Are Not Available for Selection • 900
- AOM • 911
- application insight module, AIM • 911
- Application Insight Modules (AIMs) • 66
- AppLogic appliance • 911
- AppLogic application • 911
- AppLogic assembly • 912
- AppLogic component • 912
- AppLogic grid • 912
- AppLogic Management • 478
- AppLogic Prerequisites • 479
- AppLogic template • 912
- Apply Global SNMP Settings and Access Control Lists to Policies • 95
- Apply Policy and Templates to Servers and Verify Settings • 191
- Apply Policy to Machines • 257
- Apply Predefined Autowatchers • 199
- Apply Required Settings for Using Microsoft Hyper-V • 428
- Apply Required Settings for Using Microsoft SCVMM • 434
- Apply Templates to Machines • 224
- Apply the Package Wrapper SNMP Settings as Server-level Settings • 104
- Apply the Policy • 111
- Approve or Reject Reservation Requests • 811
- Architecture • 21, 547
- Assign External Directory User Groups to User Groups • 39

- Assign User Groups Access Rights to Services • 44
- Assign Users to Groups • 39
- Associate Service Profiles with Blades • 465
- Attributes Show a Value of Zero • 875
- Audit Trail • 132
- Automate RSI Actions • 751
- Automating Processes with CA Process Automation • 512
- Automation • 547
- automation integration platform (AIP) • 912
- automation object model (AOM) • 912
- autoshell • 912
- autoshell loadable module, ALM • 912
- Available Solaris Zones Actions • 424

B

- Back Up a UCS Manager Configuration • 466
- Bare Metal Provisioning to a Cisco UCS Blade • 728
- blade (UCS) • 913
- Blank Query Results Tab after Upgrade • 883
- Browsers Do Not Display Consecutive Spaces in Events • 875

C

- CA AppLogic • 913
- CA Configuration Automation Agent Stops During Installation • 885
- CA DSM Functionality is Not Fully Supported • 884
- CA EEM Password Change Causes Authentication Failure • 891
- CA HACMP SystemEDGE AIM Trap Types • 541
- CA IBM HACMP SystemEDGE AIM Traps • 541
- CA NSM Fails to Receive Traps on Windows • 876
- CA Patch Manager • 712
- CA Process Automation Prerequisites • 513
- CA SDM Exception Error • 906
- CA Server Automation Troubleshooting • 873
- CA Server Automation vCenter Management Recommendations • 850
- CA Software Delivery • 711
- CA Technologies Product References • 3
- calpara.xml File • 369
- Cancel Network Discovery • 53
- capped logical partition (LPAR) • 913
- Capture or Deploy Fails • 894
- Capture RSI Images • 744
- Capture Services from vCenter • 688
- Capture Windows Driver Collections • 745

Capture Windows Driver Sets • 744
Change Machine State
 Microsoft Hyper-V • 577
Change the CA EEM Administrator Password
 (EiamAdmin) • 33
Change the Database Administrator (sa) Password • 34
Change the Domain Server a Distribution Server
 Connects To • 117
Change the Preferred HMC for the Managed Power
 System • 365
Change the System User Password for Active
 Directory Security • 36
Change the System User Password for Native
 Security • 35
Changing Agent Versions • 718
Chargeback • 816
Chargeback Calculations Are Lower or Higher Than
 Reservation Amounts • 900
chassis (UCS) • 913
CIM • 913
Cisco Nexus 1000V Switch • 913
Cisco UCS • 80, 446
Cisco UCS Folder Does Not Display in UI • 876
Cisco UCS Management • 461
Cisco UCS Server • 452, 476
Cisco Unified Computing System (UCS) • 913
Citrix XenServer • 81, 388
Clone a Virtual Machine • 323
Clone a Zone • 423
Clone Machine
 Solaris Zones • 578
Clone vApp • 311
cluster • 913
cmdlet • 913
Collection Engine and Report Support for Remote
 Monitoring Metrics • 554, 556
common information model (CIM) • 914
Common Usage of Policy Configuration Functions • 208
Compatibility Libraries for Linux • 156
Complete Prerequisite Tasks • 736
Conditional Substitution • 833
Configuration • 114, 546
Configuration Overview • 158
Configuration Prerequisites • 550
Configure a CA Process Automation Process • 516
Configure AIMs with NodeCfgUtil • 865

Configure AIMs with NodeCfgUtil in Command Mode
 • 870
Configure AIMs with NodeCfgUtil in Dialog Mode • 867
Configure and Validate the Amazon EC2 Server • 726
Configure and View Applied Policies • 259
Configure Announcements • 798
Configure CA Network Automation • 507
Configure CA Process Automation for Single Sign-On
 • 513
Configure CA SDM • 562
Configure CA Server Automation to Forward Events • 521
Configure Chargeback • 818
Configure Chargeback by Resource • 818
Configure Chargeback by Resource for AppLogic
 Applications • 823
Configure Chargeback by Tier for Amazon EC2 • 821
Configure Chargeback by Tier for IBM PowerVM
 Logical Partitions • 821
Configure Chargeback Display • 823
Configure Chargeback for Storage Tiers • 820
Configure CPU • 382
Configure CPU/Memory
 IBM LPAR • 579
 Microsoft Hyper-V • 581
 VMware vCenter • 583
Configure Data Collection for a Data Center • 658
Configure Data Collection for a Server • 659
Configure Data Collection for a Virtual Resource • 661
Configure Dynamic Specifications for vCenter • 673
Configure Email Notifications • 797
Configure HACMP AIM with NodeCfgUtil in
 Command Mode • 540
Configure HACMP AIM with NodeCfgUtil in Dialog
 Mode • 539
Configure Help Desk to Manage Approvals • 782
Configure IBM PowerVM Logical Partitions • 787
Configure Machine Templates for Service
 Provisioning • 673
Configure Machine Templates for WebLogic • 706
Configure Machine Templates for Wiki • 692
Configure Memory • 382
Configure NIM Master Server • 488, 738
Configure Object Aggregation • 170, 216
Configure Online Help • 825
Configure Parameters for Email Notification • 783
Configure Performance Thresholds • 663

-
- Configure Power
 - Cisco UCS • 585
 - IBM LPAR • 586
 - Microsoft Hyper-V • 590
 - VMware vCenter/Adjust vApp Power • 592
 - Configure Reservation Manager for Amazon EC2 • 778
 - Configure Reservations • 837
 - Configure Resources for Service Provisioning • 672
 - Configure Service Profile
 - Cisco UCS • 594
 - Configure Services • 840
 - Configure Shares
 - VMware vCenter • 596
 - Configure Short Descriptions on the Home Page • 827
 - Configure Snapshots • 841
 - Configure SNMP Management Servers • 522
 - Configure SNMPv1 Traps by Editing the sysedge.cf File • 519
 - Configure Solaris DHCP Servers Using the dpmutil Utility • 493
 - Configure SSH • 539
 - Configure SSH for JumpStart • 505
 - Configure the CA Network Automation Server • 507
 - Configure the CA SDM Ticket Status Setting • 563
 - Configure the Cisco UCS AIM from the Command Line • 477
 - Configure the Contact Hyperlink • 824
 - Configure the Metric Filter • 664
 - Configure User Notification Email for Job Failures • 798
 - Configure Windows for SNMP • 518
 - Configuring CPU and Memory • 381
 - Configuring Data Collection • 655
 - Configuring Remote Monitor Systems • 551
 - Configuring Resources • 475
 - Configuring Service Resource Pools • 792
 - Considerations for Resource Pools and Templates • 796
 - Contact CA Technologies • 4
 - container (Solaris) • 914
 - Control Power Status for Logical Partitions • 383
 - Control Zone Status • 422
 - Conventions • 18
 - Convert a Template to a Virtual Machine • 325
 - Convert a Virtual Machine to a Template • 326
 - Convert Template to VM
 - VMware vCenter • 597
 - Convert the VM to a Template (XenServer) • 403
 - Convert VM to Template
 - VMware vCenter • 599
 - Copy a Monitor Within SystemEDGE Policy • 238
 - Copy a Package Wrapper • 134
 - Copy SRM Policy • 241
 - Copy SRM Test • 246
 - Copy SRM Test Definition Template • 252
 - Copy SRM Threshold Definition Template • 255
 - Copy SystemEDGE Monitoring Template • 222
 - Copy SystemEDGE Policy • 209
 - Copy the post_install.sh File • 495
 - CPU shares (VMware) • 914
 - Create a Custom Action • 649
 - Create a Deployment Job • 136
 - Create a Global SNMPv3 Object • 108
 - Create a History Monitor • 183
 - Create a Log File Monitor • 180
 - Create a New Package Wrapper • 133
 - Create a Policy • 109, 164
 - Create a Process Group Monitor • 185
 - Create a Process Monitor • 178
 - Create a Resource Pool by Importing Inventory • 788
 - Create a Resource Pool for IBM PowerVM Logical Partitions • 784
 - Create a Resource Pool for VMware Virtual Machines • 770
 - Create a Rule • 564
 - Create a Rule for CPU Metric to Decrease Allocation • 307
 - Create a Rule for CPU Metric to Increase Allocation • 307
 - Create a Service • 55
 - Create a Snapshot • 326
 - Create a Sub Organization • 467
 - Create a Template for a Service • 793
 - Create a Template for Amazon Machine Images • 781
 - Create a Template for IBM PowerVM Logical Partitions • 786
 - Create a Template for the Image Type • 793
 - Create a Template for VMware Virtual Machines • 774
 - Create a Threshold Monitor • 176
 - Create a UCS Pool • 469
 - Create a User Group • 38
 - Create a Windows Event Monitor • 182
 - Create Action and Rules • 443
 - Create an Action for CPU Metric • 306
-

- Create an EC2 Resource Pool • 779
- Create an Organizational Unit • 804
- Create and Apply an Autowatcher to a System • 200
- Create and Update Files • 679
- Create Automation Policy • 653
- Create CA EEM Users • 31
- Create Configuration Sets • 556
- Create Default User Groups • 32
- Create Event • 600
- Create Group Templates • 794
- Create Port Profile Network Topology • 473
- Create Port Profiles and Port Profile Clients • 473
- Create Report • 601
- Create Resource Pool • 421
- Create Service • 602
- Create Service Templates • 683
- Create Templates for Server Workload • 174
- Create VMware Customization Specifications and Templates • 796
- Create WebLogic Service Template • 710
- Create Wiki Service Template • 703
- Customization • 824
- Customize Reservation Ready User Notification Email • 797
- Customize the Home Page • 826
- Customize the RSI Server Package • 755

D

- Database Considerations • 846
- Databases • 25
- datacenter (VMware) • 914
- datastore (VMware) • 914
- DB Transaction Log Sizes Increase Unexpectedly • 876
- Decide How to Manage Approval Requests • 782
- Default Package Wrappers • 120
- Default Values • 372
- Define a History Monitor • 234
- Define a Log File Monitor • 231
- Define a Process Group Monitor • 236
- Define a Process Monitor • 229
- Define a Schedule • 651
- Define a Threshold Monitor • 227
- Define a Windows Event Monitor • 233
- Define an Action Sequence • 650
- Define Apache HTTP Server • 693
- Define Applications • 674
- Define Applications for Wiki • 692

- Define MediaWiki Content • 701
- Define MediaWiki Database • 699
- Define MIB Extensions • 187
- Define MySQL • 697
- Define Network Address Pools • 795
- Define New SRM Policy • 241
- Define New SRM Test Definition Template • 250
- Define New SRM Threshold Definition Template • 253
- Define New SystemEDGE Monitoring Template • 217
- Define PHP • 695
- Define SRM Control Settings • 249
- Define SystemEDGE Policy Control Settings • 165, 211
- Define the WebLogic Application • 707
- Define Traps and Communities • 171
- Define Your JumpStart Boot Servers • 791
- Delete a Monitor from SystemEDGE Policy • 239
- Delete a Network • 54
- Delete a Package Wrapper • 135
- Delete a Snapshot • 327
- Delete a System • 50
- Delete a UCS Pool • 471
- Delete a Virtual Machine • 328, 442
- Delete a Zone • 424
- Delete all Snapshots • 327
- Delete Logical Partition • 386
- Delete Machine
 - IBM LPAR • 603
 - Microsoft Hyper-V • 605
 - Solaris Zones • 606
 - VMware vCenter • 607
- Delete Managed Resources • 60
- Delete Monitors from Templates or a Policy • 190
- Delete Services • 58
- Delete SRM Policy • 242
- Delete SRM Test • 246
- Delete SRM Test Definition Template • 253
- Delete SRM Threshold Definition Template • 256
- Delete SystemEDGE Monitoring Template • 223
- Delete SystemEDGE Policy • 210
- Delete User Groups • 43
- Deleted OS Images from CA ITCM are Not Deleted from CA Server Automation • 885
- Departmental Data Center • 857
- Deploy a Virtual Machine from a Template • 328
- Deploy a Wiki • 704
- Deploy CA ITCM Software Packages or Groups • 681
- Deploy RSI Images • 746

- Deploy RSI Images to AppLogic • 748
- Deploy RSI Images to Bare-Metal Servers • 747
- Deploy the RSI Agent Packages • 758
- Deploy the RSI Server • 757
- Deploy WebLogic Server • 711
- Deploying/Installing SystemEDGE Agents Using Custom Ports • 142
- Deployment Components • 114
- Deployment Credential Restrictions • 132
- Deployment Dashboard Views • 115
- Deployment Jobs • 149
- Deployment Management Certificate on Linux or UNIX • 155
- Deployment Management Certificate on Windows • 155
- Deployment Package Configuration File • 131
- Deployment Package Library • 128
- Deployment Packages • 119
- Deployment Primer Installation on Linux or UNIX • 155
- Deployment Primer Installation on Windows • 154
- Deployment Restrictions • 132
- Deployment Sizing Key Factors • 118
- Deployment to Windows Vista, Windows 2008 and Windows XP Computers Running Firewall Software • 147
- Deprecated Solaris Zones AIM Attributes Always Show N/A or Zero • 877
- Device Management for VMs • 292
- Directory Structure • 827
- Disable Software Deployment • 841
- Discover a Network • 51
- Discover a System • 49
- Discover AppLogic Grids • 480
- Discover Host by Name • 608
- Discover Network • 609
- Discover the Servers • 437
- Discover the System Running SystemEDGE in Unmanaged Mode • 269
- Discover the System to Run SystemEDGE in Managed Mode • 272
- Discovering Large Networks • 886
- Discovering the Agents • 208
- Discovery • 49
- Discovery Does Not Identify Operating System • 886
- Distribute Policies to Server Groups • 98
- Distributed Virtual Switches • 314
- Distribution Server Recommendations • 856
- Domain Server is not available • 877

- Domain Server Recommendations • 856
- dual HMC (LPAR) • 914
- Duplicated Zone Entries in the Managed Folder • 887
- dvPort group (VMware) • 914
- dvUplink port (VMware) • 915
- Dynamic NIM Machine Resource Support • 489
- dynamic reconfiguration connector index, DRC-index (LPAR) • 915
- Dynamically Add or Remove Memory • 301
- Dynamically Add or Remove vCPU • 300

E

- EC2 • 915
- EC2 Linux Provisioning Password Configuration • 803
- EC2 Windows Provisioning Password Configuration • 802
- Edit a Tenant • 807
- Edit a Resource Pool for IBM PowerVM Logical Partitions • 784
- Edit a Service • 56
- Edit Startup and Shutdown Actions • 444
- Edit the ca_post_install.sh script File • 485
- Edit the cajmpst.cf File • 494
- Edit the Configuration File • 503
- Edit the Finish File • 504
- Edit the Order File • 500
- Edit the Package Table of Contents File • 501
- Edit the Profile File • 502
- Edit the Rules File • 502
- Edit VM CPU and Memory Allocation • 304, 445
- Email Customization • 827
- Email Types and Categories • 831
- Empty Task ID for the dpmvc virtualswitch Command • 878
- Enable AppLogic Application Reservation • 783
- Enable Hyper-V Virtual Machine Reservation • 777
- Enable Maintenance Mode • 71
- Enable or Disable Inheritance of Resources from the Public Org Unit • 833
- Enable Static IP Addresses • 796
- Enhanced Discovery and SNMP Information • 52
- Enhanced Search Functionality for Remote Deployment • 116
- Enter Home Page Welcome Text • 834
- entitled pool capacity (LPAR) • 915
- Error When Installing CA DSM Agent and Asset Management Plug-in • 887
- ESX Host Fault Tolerance Attributes • 296

- ESX Job Status is Current But OS Installation Not Complete • 888
- ESX Machines Fail to Discover • 888
- ESX/ESXi host (VMware) • 915
- Event Forwarding • 518
- Example for Three Server Groups • 99
- Execute Application Installation • 676
- Execute CA Process Automation Processes • 682
- Execute Commands • 677
- Extending Reservations • 811
- Extract an Installable Image from the Media • 497

F

- fair share scheduler, FSS (Solaris) • 915
- Fault Tolerance for Virtual Machines • 294
- Fault Tolerance Properties of Virtual Machines • 295
- Fault Tolerance Requirements • 295
- Features and Benefits • 545
- Fibre Channel, FC • 915
- Filter Displayed Data • 810
- FIPS 140-2 Encryption • 863
- FIPS Overview • 863
- funclet • 915

G

- General Recommendation for vCenter AIM Monitoring • 849
- Generic Autowatchers • 198
- Global and Server-level SNMP Settings • 89
- global zone (Solaris) • 915

H

- Hardware Classes • 788
- Hardware Management Console, HMC (LPAR) • 916
- Hardware Specifications • 846
- Hot-plug Support for VMs • 299
- How Autowatchers Work • 196
- How CA EEM Works with CA Server Automation • 30
- How the Low Cost Algorithm Works • 842
- How to Apply Policy and Layered Templates to Servers • 161
- How to Backup and Restore with RSI • 761
- How to Change SystemEDGE from Managed Mode to Unmanaged Mode • 267
- How to Change SystemEDGE from Unmanaged Mode to Managed Mode • 270
- How to Change the Configuration Mode for SystemEDGE • 263

- How to Configure AIX NIM Imaging • 484
- How to Configure and Provision Amazon EC2 Resources • 722
- How to Configure Hyper-V Management • 426
- How to Configure Microsoft Cluster Service Management Components • 524
- How to Configure SNMP and Access Control Lists • 88
- How to Configure SNMPv1/v2 Settings and Access Control Lists • 91
- How to Configure SNMPv3 • 106
- How to Configure Software Delivery • 511
- How to Configure Storage • 510
- How to Configure SystemEDGE and Service Response Monitor Through Policies and Templates • 157
- How to Configure the Amazon Virtual Private Cloud • 725
- How to Configure the Amazon Web Service (AWS) • 723
- How to Configure the Cisco UCS Management Components • 450
- How to Configure the PowerVM Management Components • 355
- How to Configure the Solaris Zones Management Components • 407
- How to Configure the vCenter Server Management Components • 278
- How to Configure the vCloud Director Management Components • 337
- How to Configure XenServer Management Components • 390
- How To Create a Solaris 8 Image • 495
- How to Create and Apply an Autowatcher to a System • 195
- How to Create SRM Policy • 207
- How to Create SystemEDGE Policy • 208
- How to Deploy a Wiki Web Page • 690
- How to Deploy Oracle WebLogic Server • 704
- How to Deploy Rapid Server Imaging Using CA ITCM • 751
- How to Deploy SystemEDGE and AIMs • 112
- How to Manage Port Profiles • 472
- How to Manage Server-level SNMP Settings • 102
- How To Modify Configuration Files • 500
- How to Monitor a Specific Windows Performance Registry Metric • 204
- How to Monitor User-specific Metrics (MIB Extensions) • 202
- How to Perform Disaster Recovery with RSI • 762

- How to Prepare Windows Templates for XenServer Provisioning • 399
- How to Provision an IBM AIX Image Using MKSYSB • 734
- How to Provision EC2 Instances • 727
- How to Provision Services • 668
- How to Set Up Rapid Server Imaging for AppLogic • 759
- How To Use Centralized Service Profiles • 462
- How to Use Policy Actions to Identify Performance Issues • 305
- Hyper-V • 82
- Hyper-V Management • 438
- Hyper-V Management Actions • 446
- Hyper-V Server Connection Failed • 431
- Hyper-V Windows Provisioning Password Configuration • 800

I

- IBM AIX Provisioning with NIM • 729
- IBM High Availability Cluster Multiprocessing • 82, 537
- IBM High Availability Cluster Multiprocessing (HACMP) • 916
- IBM PowerVM • 83
- IBM PowerVM (LPAR) • 354
- IBM PowerVM Configuration Use Cases • 358
- IBM PowerVM Management • 380
- IBM PowerVM Server Administration Overview • 374
- Imaging Services • 667
- Import a Monitoring Template to SystemEDGE Policy • 221
- Import a SystemEDGE Configuration to a Policy • 210
- Import a SystemEDGE Configuration to a Template • 225
- Import a Test Definition Template into SRM Policy • 251
- Import a Threshold Definition Template into SRM Policy • 254
- Import an Existing SRM Configuration • 257
- Import External Directories • 43
- Import Templates from CA AppLogic • 689
- Increase the Size of /tmp and /opt File Systems • 737
- Increase the size of the /tmp and /opt filesystems • 486
- Infrastructure Deployment Process • 150
- Install and Run the Sysprep Tool on Windows 2003 R2 • 401
- Install CA provisioning helper • 401
- Install Imaging on a JumpStart Server Using the Text Terminal Console • 492
- Install NIM Adapter on AIX NIM Server • 485, 736
- Install the JumpStart Adapter from the Command Line • 491
- Install the JumpStart Adapter using a Response File • 491
- Install the Sysprep Tool • 402
- Installation Target Cannot be Resolved • 901
- Installation With Rapid Server Imaging Fails • 899
- Instances • 370
- Integrated Virtualization Manager (IVM, LPAR) • 916
- Integration • 547
- Interaction Between AIX LPAR Management Components • 357
- Interaction Between Cisco UCS Management Components • 447, 453
- Interaction Between IBM HACMP Management Components • 538
- Interaction Between Remote Monitoring Components • 544
- Interaction Between Solaris Zones Management Components • 410
- Interactions Between Citrix XenServer Management Components • 392
- Interactions Between Hyper-V Server Management Components • 429
- Interactions Between MSCS Management Components • 527
- Interactions Between vCloud Management Components • 339
- internet Small Computer Systems Interface, iSCSI • 916
- Introduction • 17
- Invalid X Configuration When Provisioning Linux Images • 895

J

- Job Status Filter • 116
- JumpStart Adapter Installation • 491
- JumpStart for Solaris • 493
- JumpStart Prerequisites • 490
- JumpStart Provisioning Password Configuration • 800

K

- Key Performance Indicator Metrics • 546

Key Points About Metrics Collection • 656

L

Large Environments • 860
Layered Templates • 219
Layered Templates Concept • 162
Let Users Manage the Power Status of an IBM
PowerVM Logical Partition • 785
Let Users Perform Some Administrative Tasks • 799
lightweight process, LWP (Solaris) • 916
Linux and UNIX Provisioning Fails When Using
Resizing Option • 895
List of Predefined Action Types • 571
Listing OS Types, Depots, or Networks Returns No
Results • 895
Local and Remote Monitors Do Not Show the Same
Values • 878
Log In with Different User Credentials on the Same
Computer Using IE8 • 888
logical memory block, LMB (LPAR) • 916
logical partition, LPAR • 916
Logical Partitions • 783
Logical Volumes in Virtual Machines • 301
LPAR AIM Monitoring Recommendations • 852
LPAR Monitoring • 376
LPAR Provisioning for IBM AIX • 728

M

Make Amazon Machine Images Available from EC2 •
777
Make Hyper-V Virtual Machines Available to Users •
777
Make Operating System Images Available to Users •
791
Make Physical Systems Available to Users • 787
Make Services Available to Users • 792
Make Virtual Machines Available to Users • 769
Manage AppLogic Appliances • 484
Manage AppLogic Applications • 483
Manage AppLogic Resources • 482
Manage Central Service Profiles • 463
Manage Cluster Services • 329
Manage Distributed Switch
VMware vCenter • 611
Manage Fault Tolerance • 299
VMware vCenter • 613
Manage RSI Images and Drivers • 750
Manage Snapshots in VMWare Resource Pools • 771

Manage Unmanaged Resources • 59
Manage Virtual Switch
VMware vCenter • 618
Manage VM Snapshots
VMware vCenter • 615
Manage VM Status (Hyper-V) • 441
Manage VM Status (VMware) • 324
Manage VM Status (XenServer) • 404
Manage Windows Service • 620
Managed and Unmanaged Resources • 58
Managed Mode and Unmanaged Mode • 66
Management DB • 25
Management Information Base (MIB) • 916
Managing Configuration Entries • 557
Managing Credential Settings • 556
Managing SystemEDGE and Application Insight
Modules (AIMs) • 79
Managing Systems Performance • 47
Managing Systems Using Remote Monitoring • 554
Managing Users and User Groups • 29
Managing Virtual Environments • 275
Manual Installation of the Infrastructure
Deployment Primer Software • 154
MIB objects, MIB attributes • 917
Microsoft Cluster Server • 84
Microsoft Cluster Service • 523
Microsoft Cluster Service Management • 536
Microsoft Hyper-V Server • 424
Migrate a Virtual Machine • 329
Migrate Machine
VMware vCenter • 622
Migrate Servers with RSI • 748
Modify a Monitor Within SystemEDGE Policy • 239
Modify a Package Wrapper • 134
Modify Cluster Properties • 536
Modify CPU
VMware vCenter • 623
Modify Existing Template in SystemEDGE Policy •
240
Modify Memory
VMware vCenter • 624
Modify SRM Test • 245
Modify SRM Test Definition Template • 252
Modify SRM Threshold Definition • 248
Modify SRM Threshold Definition Template • 255
Modify System Attribute Values • 789
Modify SystemEDGE Monitoring Template • 222
Modify the Physical System Allocation Policy • 842
Monitor a Virtual Machine • 330

- Monitor an ESX Server • 331
- Monitor Distributed Virtual Switches Through Events • 319
- Monitor Fault Tolerance • 297
- Monitor MS Cluster Services • 537
- Monitor vApps Through Events • 312
- Monitored vSphere and vCenter Server Resources • 276
- Monitoring Clusters • 523
- Monitoring Software Settings • 70
- More vApp Operations • 312
- Multiple Data Centers • 858
- Multiple Distribution Servers • 118
- Multiple Shared-Processor Pools (MSPPs) • 917
- Multiple Virtual I/O Servers • 917
- Multi-Tenancy Environment • 805

N

- Native Security • 30
- Network Considerations • 847
- network installation manager, NIM (LPAR) • 917
- Network Properties • 317
- New System Name is not Displayed • 890
- NIM Provisioning Password Configuration • 801
- No Cisco UCS Manager in Explore Pane • 889
- No Resources Available Message When Requesting a VM • 901
- NodeCfgUtil Overview • 865
- non-global zone (Solaris) • 917
- Notes on Infrastructure Deployment Using IPv6 Addresses • 153

O

- Open HelpDesk Ticket • 626
- OpenSSL Software Compatibility Issues • 890
- Operations on vApps in vCloud • 352
- Oracle Solaris Zones • 84
- organizational unit • 917
- Organizational Units • 803
- OSIM Provisioning Password Configuration • 801
- Override Automatic Selection • 839
- Overview • 21, 112, 490

P

- Package Filter • 130
- Partitions • 371
- Password Change Causes Error Message • 902

- Password Changes May Cause Authentication Errors • 890
- Password Management • 33
- Perform a Point Agent Configuration • 69
- Performance Considerations during Initial Discovery • 852
- Performance DB • 26
- Persistent Data • 369
- platform management module, PMM • 917
- Policies • 316
- policy based configuration • 917
- Poll Groups • 372
- poll interval • 917
- Port Group Properties • 317
- Port Properties • 317
- POWER processors (LPAR) • 917
- Predefined Content and Configuration for Service Provisioning • 769
- Prepare an Image • 401
- Prepare CA Server Automation for Reservation Manager • 768
- Prepare Your Directories • 496
- Prepare Your Environment for Reservation Manager • 766
- Prerequisite Knowledge • 734
- Prerequisites • 721, 729, 777
- Prerequisites for Automatically Deploying CA Server Automation Infrastructure • 151
- Prerequisites for Supporting Reservations of Physical Systems • 787
- Prerequisites for Supporting Reservations of Virtual Machines • 769
- Prerequisites for XenServer Environments • 401
- Process and Service Autowatchers • 198
- processor pools (LPAR) • 918
- processor set, pset (Solaris) • 918
- project (Solaris) • 918
- Properties • 315
- Protocols for Transferring Packages Employed by IDManager • 154
- Provide the Deployment Management Certificate to a Primer Installation • 155
- Provision a Citrix XenServer Virtual Machine • 405
- Provision AppLogic Applications • 481
- Provision Jobs for VMware Resource Pools • 774
- Provision Machine
 - IBM LPAR • 627
 - Microsoft Hyper-V • 630
 - Solaris Zones • 633

- VMware vCenter • 636
- Provision Services • 690
- Provision Storage for a VMware Resource Pool • 773
- Provision the IBM AIX Image • 739
- Provision vApp from Template • 352
- Provision VMware vApp • 309
- provisioning • 918
- Provisioning Resources • 667
- Public Templates for End Users • 793

R

- Rapid Server Imaging • 740
- Rapid Server Imaging (RSI) Troubleshooting • 893
- Rapid Server Imaging Support for UCS • 764
- Reconfigure the SystemEDGE Agent Port • 144
- Rediscover a Network • 54
- Register a Cluster • 535
- Register a UCS AIM Server • 449, 456
- Register the RSI Agent Packages • 758
- Register the RSI Server • 743
- Register the RSI Server Package • 754
- Registered Server Fails • 896
- regular expressions • 918
- Related Publications • 17
- Remote and Multi-instance vCloud Director Support • 350
- remote deployment • 918
- Remote Deployment and Policy Configuration Overview • 847
- Remote Deployment and Policy Configuration Recommendations • 854
- Remote Deployment Architecture • 113
- Remote Deployment to Solaris Lists SPARC and x86 Systems • 879
- Remote Deployment to UNIX/Linux Using Non Privileged User Account • 146
- Remote Monitoring • 86, 543
- Remove a Cluster • 535
- Remove Disk
 - VMware vCenter • 639
- Remove Managed Mode Information from the Manager • 268
- Remove Managed Mode Information from the SystemEDGE Configuration • 267
- Remove Network Interface
 - VMware vCenter • 640
- Remove Server From Service • 641
- Remove Server from Services • 57

- Remove Unmanaged Mode Information from the Manager • 271
- Remove Unmanaged Mode Information from the SystemEDGE Configuration • 270
- Remove Users or User Groups from a User Group • 44
- Removing a vCenter Server Lets Objects of Another Managed vCenter Server Disappear • 879
- Rename a Package Wrapper • 135
- Rename a UCS Pool • 470
- Rename a Virtual Machine • 442
- Rename SRM Policy • 242
- Rename SRM Test Definition Template • 252
- Rename SRM Threshold Definition Template • 256
- Rename SystemEDGE Monitoring Template • 223
- Rename SystemEDGE Policy • 209
- Requirements for Solaris Zones Management • 409
- Reservation Manager Prerequisites • 765
- Reservation Manager Troubleshooting • 899
- Reserve a System • 792
- Reserve a Virtual Machine • 776
- Resetting the vCenter Server Password Causes Data Collection to Fail • 879
- Resilience • 547
- Resource Allocation • 301
- Resource Allocation and Forecast Charts • 812
- Resource Allocation Best Practices • 303
- Resource Allocation Limit • 303
- Resource Allocation Reservation • 302
- Resource Allocation Shares • 302
- resource control (Solaris) • 918
- resource pool (Solaris) • 918
- resource pool (VMware) • 919
- Restart Logical Partition • 387
- Resubmit a Deployment Job • 140
- Revert a Policy Back To an Earlier Version • 261
- Revert to a Snapshot • 331
- Review Common Requirements (SNMPv3) • 107
- Review Hyper-V Requirements • 427
- Review Interactions Between vCenter Server Management Components • 280
- Review Managed Mode and Unmanaged Mode Details • 264
- Review Monitoring Template Application Progress • 224
- Review Policy Application Progress • 259
- Review Requirements • 196, 264, 279, 356, 391, 408, 451, 525
- Review Requirements (Server-level) • 102

- Review Requirements (SNMPv1/2) • 92
- Review SNMP Configuration and Policy Relationships • 92
- Review SNMPv3 Configuration Details • 107
- Review vCloud Requirements • 338
- RSI
 - Image is not Deployed Because of Timeout • 894
 - Remote Server Discovery Fails • 896
- RSI Agent Error During Image Capture or Deploy • 896
- RSI and ITCM Servers • 894
- RSI Imaging Fails with SSL Errors • 897
- RSI Prerequisites • 741
- RSI Prerequisites for AppLogic • 742
- RSI Prerequisites for Hyper-V • 743
- RSI Server Errors • 893
- Rule Planning • 564
- Rules and Actions • 561
- Run Action • 643
- Run Action Sequence • 645
- Run Command Script • 647
- Run Frequently Used Reports • 813
- Run the Sysprep Tool on Windows 2003 R2 • 402
- Run the Sysprep Tool on Windows 2008 R2 • 402

S

- Scalability • 117, 547
- Scalability Best Practices • 845
- Scalability Limitation in Terms of Monitored Servers • 850
- Scalability Limitations in Terms of Monitored Objects • 850
- Scalability Overview • 845
- Scalability Recommendation for LPAR AIM Monitoring • 853
- Scalability Recommendation for Solaris Zones AIM Monitoring • 854
- Scalability Recommendations • 849
- Scalability Use Cases • 856
- Scheduled Jobs do not Run • 905
- SCVMM Server Connection Failed • 436
- Search for Users or User Groups • 37
- Security and Maintenance • 70
- Security Considerations for Active Directory • 30
- security group (Amazon EC2) • 919
- Select a Chargeback Tier for IBM PowerVM Logical Partitions • 822
- Server Connection to the Manager Failed • 362, 411, 455, 528
- Server Connection to the Manager Failed (Citrix XenServer) • 393
- Service Provisioning • 668
- Service Response Monitoring • 72
- service-oriented architecture (SOA) • 919
- Services • 54
- Set Automatic Cancellation of Unapproved Reservations • 843
- Set Health State • 648
- Set Limits on Virtual Machine Resources • 835
- Set Over Commitment of Memory on ESX Server or Cluster • 835
- Set Run Command Script Privileges • 42
- Set User Group Permissions • 41
- Set User Group Permissions for Services • 42
- Set User Group Privileges • 41
- Setting Up Reservation Manager • 765
- Setup and Configuration • 768
- Share Application Definitions and Service Templates • 688
- shared memory (Solaris) • 919
- Shut Down Logical Partition • 388
- Simple Network Management Protocol (SNMP) • 919
- Slots • 371
- snapshot • 919
- SNMP Consistency • 88
- SNMP V3 Engine ID • 521
- SNMPv3 • 919
- SOA • 919
- Software Delivery Adapter Errors • 907
- Software Delivery Configuration File • 714
- Solaris JumpStart Provisioning • 489
- Solaris SPARC Provisioning Fails on Discovery of DVD • 897
- Solaris Zones • 406
- Solaris Zones AIM Monitoring Recommendations • 853
- Solaris Zones AIM Reset if a Monitored System is Down • 880
- Solaris Zones Management • 418
- Specific Remote Deployment Use Cases • 142
- Specify a Folder for VMware Virtual Machines • 836
- Specify a Prefix VMWare for Virtual Machine Names • 776
- Specify a Timeout Value • 834
- Specify Default Policy for New Instances • 262

Specify Global SNMP Settings and Access Control Lists • 94

Specify Memory and CPU Selections • 839

Specify Read-Write Community Post-Install • 138

Specify Read-Write Community Prior To Deployment • 137

Specify the Maximum Number of NICs per Virtual Machine • 836

Specify the Time Period for User Expiration Notification • 797

Specify When to Send a Stalled Task Alert • 798

Specify When to Send Pending Approval Request Notification • 843

SQL User Password Change Causes Blank UI • 892

SRM Tests • 73

SSP - The Home content does not display in Internet Explorer 9 • 908

Start or Stop the NIM Adapter Daemon • 487, 738

State Management Model • 64

Stateless Monitoring • 65

Static IP Addresses for IBM PowerVM Logical Partitions • 786

Status Icon of Component Shows Not Configured • 880

Stop VMs from being Provisioned from Resource Pools • 772

Storage Area Network, SAN • 919

Storage Prerequisites • 509

Storage Provisioning Features • 508

storage tiers • 920

stringification • 920

Substitution Variables • 828

Substitution Variables for Templates • 774

Super Administrator • 806

Supported Features • 720

Suspend and Restart the Scheduling of Tasks • 814

Suspension and Restart of Individual Tasks • 815

Synchronize NIM Master Servers • 488

System User Password Change Causes Blank UI • 893

SystemEDGE and Advanced Encryption • 87

SystemEDGE Features • 60

Systems • 370

Systems Management • 47

Systems Management MIB • 62

T

Target Server Stops Responding During Reboot After Image Capture or Deployment • 898

task (Solaris) • 920

Templates and Email Types • 828

Tenant Administrator • 807

Tenant End User • 809

The AIM Instance Status Icon Shows Disabled • 368, 399, 417, 460, 534

The AIM Instance Status Icon Shows Discovery in Progress • 366, 397, 415, 458, 532

The AIM Instance Status Icon Shows Error • 366, 397, 416, 458, 532

The AIM Instance Status Icon Shows No Polling • 366, 397, 416, 458, 532

The Sysprep Tool • 401

Tier Label Changes on VMware Datastore • 902

tiers • 920

time-sharing scheduler, TS (Solaris) • 920

Track Deployment Job Status • 139

trap • 920

Troubleshoot Adding Amazon Machine Images • 781

Troubleshoot the AIM Instance Connection • 365, 396, 415, 457, 531

Troubleshoot the vCenter AIM Instance Connection • 287

Troubleshoot the vCenter Server Connection • 283

Troubleshoot the vCloud AIM Instance Connection • 345

Troubleshoot the vCloud Server Connection • 341

Troubleshooting • 873

U

UCS • 920

UCS Action Types • 474

UCS Manager • 920

UCS Organizations • 467

UCS Pools • 468

Unable to Connect to Microsoft SQL Server • 880

Unable to Find Package Entries for Personality AutoDeploy • 902

Unable to Retrieve Information from vCenter • 903

Understanding Packaging • 712

Uninstall a JumpStart Adapter • 492

Unmanage Managed Resources • 59

Unregister a Virtual Machine • 332

Update the Hashed Password Variable • 486, 737

Upgrading SystemEDGE • 881

Use a Predefined Action Type • 567

Use Case

- Adding a New Rule to a Service • 654

- Adding a Server to a Service • 653
- Defining an Action • 654
- Use Case Scenario • 549
- Use Cases for Policies • 653
- Use Help Desk for Reservation Approvals • 770
- Use Static IP Addresses • 794
- Use the Reservation Manager Mobile App • 824
- User Access Control • 29
- User Access to Reserved Systems • 800
- User Group Management • 37
- User Interface • 26
- User Interface is Unresponsive on Provisioning and Policy Screens • 881
- User Management • 803
- User Permissions and Access Requirements Reference • 79
- Using Generic Groups and Templates • 713
- Using Remote Deployment • 132
- Using Rules and Actions • 561
- Using the MKSYSB Utility • 733

V

- vApp Support • 307
- vApp Support in vCloud • 350
- vCenter AIM Instance Status Icon Shows Disabled • 291
- vCenter AIM Instance Status Icon Shows Discovery in Progress • 289
- vCenter AIM Instance Status Icon Shows Error • 289
- vCenter AIM Instance Status Icon Shows Multiple Instances • 291
- vCenter AIM Instance Status Icon Shows No Polling • 290
- vCenter AIM Monitoring Recommendations • 849
- vCenter Automation and Policy Actions • 333
- vCenter Management Limitations in Terms of Virtual Machines • 851
- vCenter Server (VMware) • 920
- vCenter Server Agent (VMware) • 920
- vCenter Server AIM Attributes Show Zero • 881
- vCenter Server Connection Failed • 284
- vCenter Server Database (VMware) • 921
- vCenter Server Folder Does Not Display in UI • 908
- vCenter Server in a Cluster • 313
- vCloud AIM Instance Status Icon Shows Disabled • 349
- vCloud AIM Instance Status Icon Shows Discovery in Progress • 347
- vCloud AIM Instance Status Icon Shows Error • 347
- vCloud AIM Instance Status Icon Shows No Polling • 348
- vCloud Folder Structure • 350
- vCloud Organizations • 352
- vCloud Server Connection Failed • 342
- Verify the Autowatcher • 201
- Verify the Cisco UCS in the Resources Tree • 460
- Verify the Citrix XenServer Group in the Resources Tree • 399
- Verify the Current Configuration Mode of SystemEDGE • 265
- Verify the Group in the Resources Tree • 368
- Verify the Hyper-V Server Folder in the Resources Tree • 438
- Verify the Microsoft Cluster Service in the Resources Tree • 534
- Verify the SNMPv3 Settings in the System Summary • 111
- Verify the Solaris Zones Group in the Resources Tree • 418
- Verify the SystemEDGE Configuration Mode • 272
- Verify the vCenter Server Folder Appearance in the Resources Tree • 292
- Verify the VMware vCloud Folder in the Resources Tree • 349
- View a UCS Pool • 468
- View Cisco UCS Resources • 464
- View Custom Specifications • 334
- View Deployed Packages • 141
- View Deployment History • 142
- View General Information • 334
- View Managed Object States • 76
- View Monitors Within a SystemEDGE Policy • 237
- View Resource Summary and Events • 380
- View SRM Tests • 77
- View SystemEDGE Monitors • 75
- View Systems in the Reservation Manager Inventory • 789
- Viewing Query Results • 555
- virtual disk (VMware) • 921
- Virtual I/O Server, VIOS (LPAR) • 921
- virtual LAN or VLAN • 921
- virtual local area network • 921
- Virtual Machine Counts • 318
- virtual machine, VM (VMware) • 921
- virtual NIC (VMware) • 921
- Virtual Standard Switches and Virtual Distributed Switches in the vNetwork Panel • 313

- virtual switch (VMware) • 922
- Visualization • 546
- VLAN Scoping • 809
- VM Reservation Fails
 - Could Not Find Computer UID for Software Delivery • 909
- VM Reservation Fails Because of CPU Limitation • 904
- VM Reservation Fails in a Clustered Environment • 905
- VM Resources are not Available for Dates Requested • 904
- VM Usage Values Do Not Update Immediately After Power Down • 882
- VMs Not Being Discovered • 909
- VMware Linux Provisioning Password Configuration • 802
- VMware vCenter • 87
- VMware vCenter Provisioning and Common Use Cases • 320
- VMware vCloud • 88, 335
- VMware vSphere and vCenter Server • 275
- VMware Windows Provisioning Password Configuration • 801
- vNetwork Distributed Switch, vDS (VMware) • 922
- vNetwork Standard Switch, vSwitch (VMware) • 922
- vNetwork Standard Switches (vSwitch) • 314
- vNIC Templates • 467
- vSwitch Properties • 317

W

- Windows Driver Collection Issue with IBM Servers • 898
- Windows Provisioning Fails • 899

X

- XML-RPC • 922