CA Configuration Automation®

User Guide

r12.8 SP02

CA Technologies Product References

This document references the following CA Technologies products:

- CA® Embedded Entitlements Manager (CA EEM)
- CA Spectrum®® Automation Manager
- CA® SiteMinder® Web Access Manager (CA SiteMinder)

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: Introducing CA Configuration Automation

CA Configuration Automation is a standards-based software product that lets you manage your enterprise's distributed hardware and software components from a centralized browser-based window. You can use CA Configuration Automation, to do the following:

- Discover the servers in your enterprise
- Find out what operating systems, databases, and software application components are installed on those servers
- Access complex data, information, and configuration settings from within those components
- Determine the relationships and dependencies between the servers in your enterprise
- Detect server and service configuration changes and differences
- Take and retain snapshots (point-in-time copies) of your services
- Ensure software component and configuration policy compliance to corporate standards and best practices
- Enact change on a collection of software component attributes within a service
- Troubleshoot and improve the mean time to repair your servers and services

The following sections describe the CA Configuration Automation software components and provide a high-level overview of key concepts.

CA Configuration Automation Concepts

This section describes CA Configuration Automation terminology and concepts that you may not be familiar with. This document also contains a glossary that briefly defines these, and other, terms.

CA Configuration Automation supports two different approaches for managing your enterprise's distributed software components:

- *Server-centric* management fulfills the need to manage the infrastructure of your enterprise
- *Service-centric* management fulfills the need to manage complex tiered or multi-component applications across your enterprise

CA Configuration Automation provides discovery, snapshot, refresh, change detection, comparison, and rule compliance operations for either approach.

Discovery

You can identify network segments on which you want to perform discovery operations to locate servers and software. You assign each network segment a unique name that can then be scanned for servers. In addition, you can define a Management Profile that specifies the type of scan you want to perform and how frequently you want to perform it. This Management Profile can then be assigned to one or more network segments, letting you automate the discovery operations across your enterprise.

CA Configuration Automation begins managing your enterprise applications by establishing a comprehensive, up-to-date inventory of servers and software components across your organization's networks. You can discover components to get a complete, cross-platform inventory of applications at a granular level, including directories, files, registries, database tables, and configuration parameters. The basis for application-based discovery are Blueprints, which outline the basic structure of an application to enable the CA Configuration Automation Agent to find that application on a server. Blueprints are described in detail later in this section

Chapter 2: Blueprint Overview

Blueprints are the abstract definitions or *metadata* for a software component. This metadata defines the directives and mechanisms to:

- Detect a software component on a given computer
- Capture file system and database elements of the component
- Express and show inter- and intra-component relationships and dependencies
- Locate, analyze, and manage the configuration information
- Define, execute, and interpret diagnostic macros
- Define recommended, best-practice values for all these elements

CA Configuration Automation presents each Blueprint in a standardized format that simplifies configuration and administration tasks. CA provides a library of predefined Blueprints for commonly used software components. You can also edit existing Blueprints and can create custom Blueprints.

Services

Within CA Configuration Automation, a *service* is defined as a collection of software components running on one or more managed servers. You can define a service by specifying servers, server groups and related Component Blueprints that need to be discovered. A service generally fulfills a unique business function in the enterprise, however multiple instances of a service can run within an enterprise.

CA Configuration Automation presents a standardized and annotated view of all services, including configuration details, dependencies and constraint rules, file system elements, runtime logs, diagnostics, utilities, and a component inventory.

Profiles

Profiles let you automate CA Configuration Automation discovery and management operations using the following profiles:

- Access Profiles are associated with servers and provide the rules for server access and Agent installation.
- Network Profiles provide the operational rules for discovering servers.
- Management Profiles can be created and assigned at the network and server level to manage discovery, blueprint, and software management tasks.
- Notification Profiles store notification details for creating email messages that are sent when certain operations are performed.

Snapshots

CA Configuration Automation can detect change in a component or server by monitoring your enterprise using snapshots. A snapshot is a point-in-time copy of a server's or service's software configuration. You can automatically recapture application inventories to archive configuration data into fully detailed snapshots that can be used for troubleshooting, record keeping, or release management and migration planning.

You can also designate a snapshot as the Gold Standard to use an application's states in the snapshot as a baseline for auditing and Change Detection.

Rule Compliance

You can ensure that complex applications meet internal and regulatory compliance by using the detailed information that CA Configuration Automation collects and running a Rule Compliance operation. CA Configuration Automation helps control applications and establishes best practices with flexible, in-depth policy definition and automated enforcement of the rules you define. Auditing your enterprise's performance configurations, security settings, and dependent variables hardens the application infrastructure, freeing organizations from manual, error prone reviews.

CA Configuration Automation Components

The CA Configuration Automation software includes the following components:

- CA Configuration Automation Server
- CA Configuration Automation Database
- CA Configuration Automation Agents
- CCA Grid Node
- CA Network Discovery Gateway
- CA EEM
- BusinessObjects reports server

These components are described in the sections that follow.

CA Configuration Automation Server

CA Configuration Automation Server provides a browser-based user interface that acts as a central registry through which you manage persistent storage, control data access, and manage communication with the CA Configuration Automation Agents. CA Configuration Automation Server controls all aspects of the product's operation, including discovery, configuration, reconciliation, and analysis functions. CA Configuration Automation Server is accessible from any Windows server with a supported browser.

CA Configuration Automation Database

The CA Configuration Automation Database stores all of the collected CA Configuration Automation data and configuration information, including the following:

- Server configurations (hardware, software, system information)
- Service configurations and components
- Server and Service snapshots
- Job Scheduler information
- Custom Reports definitions
- Custom Blueprints

Each instance of a CA Configuration Automation Server needs a corresponding database instance. Multiple CA Configuration Automation Servers can share the same database on the same database server, but each server must have its own set of tablespaces and tables within the database instance to store data.

CCA Agents

A CA Configuration Automation Agent is a light-weight executable that inspects and implements server-directed operations on Service Blueprint-based components running on CCA-managed servers in your enterprise. It can perform deep configuration management of both server and software configurations.

CA Configuration Automation Agents are installed as daemons on UNIX-based servers or as services on Windows-based servers.

You need to install a CA Configuration Automation Agent on every server in your enterprise on which you want CA Configuration Automation to manage servers and services in depth. In addition, we recommend that you install CA Configuration Automation Agent on each CA Configuration Automation Server machine to discover and manage the CA Configuration Automation Server components.

Note: CA Configuration Automation can also provide secure agentless interrogation and monitoring of subject systems using SSH. This option may be a viable alternative when installing an agent is not feasible or when a CA Configuration Automation Agent is not supported on a platform.

CCA Grid Nodes

Grid processing is used to increase performance by distributing operational workloads to multiple Grid Nodes. A server is capable of supporting multiple CCA Grid Nodes each with multiple threads. CA Configuration Automation operations are *Grid-enabled* so they can be divided into independent executable entities. These executable entities are distributed to available Grid servers, Grid nodes, and threads for execution.

CCA Grid Nodes are supported on Linux, UNIX, and Windows platforms and have their own installation programs. After installing a CCA Grid Node and registering it with the CA Configuration Automation Server, Grid processing is invisible to CA Configuration Automation users.

CA Network Discovery Gateway

The NDG Server is responsible for the CA Configuration Automation Discovery operations that locate and monitor servers and services in your enterprise. You must install the NDG Server on a supported Windows platform before installing the CA Configuration Automation Server. The CA Configuration Automation installation program prompts you for the name of the NDG Server and the port it uses for discovery operations.

CA EEM

CA Embedded Entitlements Manager (CA EEM) provides user and group management and role-based authentication services for the CA Configuration Automation user interfaces.

Business Objects

Business Objects is a third-party business intelligence platform shipped with CA Configuration Automation that provides interactive reporting. Predefined CA Configuration Automation reports are hosted on the Business Objects server.

Chapter 3: Using the CA Configuration Automation User Interface

This chapter introduces the CA Configuration Automation browser-based user interface. For information about the command-line interface (CLI), see Using the Command-line Interface (see page 421).

This section contains the following topics:

Log In to CA Configuration Automation

Log in to CA Configuration Automation to access the user interface. When you log in for the first time, enter the correct URL and log in as the default or user-defined CCA Administrator user. You can change your password after you access the UI.

Follow these steps:

1. Open a supported web browser and enter the appropriate following URL in the Address field.

   `http://<server>:port/CCAUI.html`

   `http://<server>:port/CCAUI.jsp`

   <server>

   Defines the CA Configuration Automation server name that you entered during the installation process.

   port

   Defines the port number that you entered during the installation process.

   Default: 8080

   The CA Configuration Automation Log In page opens.

2. (Optional) Click Favorites on the browser toolbar and select Add to Favorites from the menu to add the Log In page to your list of favorite Web pages.

3. On the Log In page, complete one of the following actions and click Log In:

   - If you accepted the default CCA Administrator during the CA Configuration Automation Server installation process, enter **ccaadmin** in the User Name and Password fields.
   - If you did not accept the default CCA Administrator during the CA Configuration Automation Server installation process, enter the name and password you specified for the CCA Administrator.

   The Tasks panel appears and displays the administrator user that you are logged in as. The panel also contains a link where you can change the associated password.

CA Configuration Automation UI Overview

When you log in to CA Configuration Automation, the Tasks panel opens by default.

You can access the following main UI panels from links in the top right corner:

- Management
- Dashboard
- Administration
- Tasks

The link to the online help system accompanies each panel link. The sections that follow introduce each panel.

Management Panel

To complete most day-to-day configuration management operations, use the Management panel.

To create, view, and manage objects of the relevant type, use the following tabs on the Management panel:

- Services
- Servers
- Software
- Network
- Storage
- Blueprints
- Compliance
- Remediation
- Jobs
- Log
- Reports

Each management tab page contains a table that lists the objects that are defined for that page. You can add objects to the tables manually or as the result of a discovery operation. You can import the objects from another application or the CA Configuration Automation installation program can install them as predefined data.

Except for the Reports tab, all management tabs contain a Filter pane with which to filter so the table displays only the selected objects. For information about creating filters, see Filter Table Views (see page 33).

Most of the management tabs also contain the following drop-down lists from which you can select management actions:

Select Actions

Contains the options for running, managing, exporting, and deleting objects and operations.

Table View

Contains the options for displaying the default table view or the custom views you create.

Table Actions

Contains the options for creating or importing objects (servers, services, and so on), and the following common tasks:

- Export to Excel
- Print
- Configure Table View

The Table Action drop-down list on every tab page lists the common tasks. For information about the common tasks, see Common Table Actions.

For more information about the Management panel, see the section that corresponds to the relevant tab.

Dashboard Panel

The Dashboard panel contains two tabs: Charts and Visualization.

Charts Tab

The Charts tab contains a Dashboard pane that includes two folders: Dashboards and Charts. The Dashboard folder contains the following predefined Dashboards that display graphical summaries of the objects you are managing with CA Configuration Automation:

- VM Hosting Servers
- VM Guest Software Components
- VM Guest Servers
- Virtualization
- Software Components

- Services
- Servers (Unmanaged)
- Servers (Managed)
- Servers
- Compliance
- Communication Relationships
- Change History

The Charts folder contains the following subfolders which contain related charts:

- All Charts
- Servers
- Relationships
- Virtual Environment
- Applications
- Software Components
- Services
- Rule Compliance
- Change Detection
- Grid Information

Dashboards, and their corresponding charts, contain options for displaying them, configuring them, removing them, refreshing them, and changing how they display information. Additionally, you can create new and custom Dashboards, and import and export Dashboards from and to other CA Configuration Automation implementations.

For detailed information about the Dashboard panel, see Dashboards (see page 365).

Visualization Tab

The Visualization tab contains a Visualization pane that includes two folders: Graphs and Templates. The Graphs and Templates folders both contain the following subfolders:

- Applications
- Servers
- Clusters and Service profiler
- Services
- Software Components

You can view any of the predefined views in the Graphs subfolders, or display, modify, and save any view in the Templates subfolder to create a custom graph.

For detailed information about the Visualization panel, see Visualization (see page 373).

Administration Panel

Define and manage CA Configuration Automation users, view and configure the CA Configuration Automation server settings, and manage port-based communication mappings on the following Administration panel tabs:

Configuration

Contains the following pages:

Properties

Contains the settings for viewing and editing how CA Configuration Automation looks and operates.

Security Certificates

Contains the settings for creating and managing CA Configuration Automation Server and CA Configuration Automation agents.

Communication Mappings

Lists the port numbers and communication types that the product commonly uses through the port. You can edit the communication type setting on this page.

Application Mappings

Lists the applications and regular expressions that identify the typical installation directory of the associated application. You can add, edit, and manage the mappings on this page.

For more information, see Configuration Settings.

Access Management

Links to the following access management pages that provide CA EEM integration functionalities:

Users

Provides the functionality for creating and managing users.

Policies

Provides the functionality for managing user access to specific CA Configuration Automation features.

Configure

Specifies where the product stores user and user group information and from where the product accesses it.

For more information, see Configuring Access Management.

Network

Contains the Network Discovery Gateways table, which displays the servers where Network Discovery Gateways are installed. You can create, manage, and delete NDG Servers from this page.

Catalyst Integration

Contains the Catalyst Attributes Profiles table which displays predefined and custom Catalyst Attributes Profiles. You can create, import, export, edit, delete, and copy Catalyst Attributes Profiles from this page.

Profiles

Defines the predefined and custom Catalyst Attributes Profiles.You can determine what CIs to export to the CA Catalyst server.

Jobs

Specifies the selected information (servers, or services, or storage systems, or blueprints) that is published from CA Configuration Automation to the catalyst server.

Log

Logs the operations that are related to profile and jobs.

Diagnostics

Links to the following diagnostics pages:

CCA Information

Displays the configuration details about the UI, and CA Configuration Automation integrations.

Database Information

Displays the configuration details about the CA Configuration Automation Database and the database schemas.

Grid Information

Displays the details for all Grid Nodes. This tab also displays the CA Configuration Automation Server grid jobs in a table view or a tree view.

Distributed Lock Information

Specifies the lock that is shared among the grids to orchestrate the allocation of services among the grids. The lock provides the server details where a scheduled job is executed. The services are assigned and reassigned in the event of failure.

Collect Diagnostics

Collects the information that CA Technologies Support requests to troubleshoot the CA Configuration Automation Server or CA Configuration Automation Server Grid Node issues.

Log Archives

Specifies the logs that are archived when the logs exceed the maximum storage size limit.

Data Migration

Provides the following options for migrating data from CA Cohesion CCA:

- From the Cohesion Database to the CA Configuration Automation r12.8 SP02
- From the Cohesion Database to a JAR file
- From a JAR file to the CA Configuration Automation Database
- Import Security Certificates from CA Cohesion ACM to CA Configuration Automation r12.8 SP02

For more information about how to migrate data from CA Cohesion ACM to CA Configuration Automation, see Migrating Data from CA Cohesion ACM.

Scripts

Links to the following script pages:

Scripts

Lists all the available scripts for the user and provides options to manage the scripts. See Scripts for more information.

Global Variables

Lists all the available global variables for the user and provides options to manage the global variables. See Manage Global Variables for more information.

Tasks Panel

Use the Tasks panel to complete the following common tasks:

- Discover Network
- Access Profile and Agent Deployment
- Discover Service
- Run Compliance Job
- Locate and Upgrade Agents

Click a task to open a wizard that contains a detailed description of the task and navigation buttons that link to the subtasks that are required to complete the task.

Filter Table Views

Each of the CA Configuration Automation Management tab pages (Services, Servers, Network, Blueprints, Compliance, Remediation, Jobs, Log, and Reports) contains a table that displays details about the corresponding objects. Some of these tables can be very large. To make it easier to work with large amounts of table data, you can create a filter that only displays the objects that are important to you.

To filter table data

1. Open any of the nine tab pages.

   The page displays a corresponding table. For example, if you select the Servers tab, the page displays the Servers table.

2. Create the filter by selecting options from the drop-down lists or typing in the following fields:

   Column

   Specifies the column in the table on which you want to filter. The drop-down list contains an option for each column in the table on which you can filter.

   Value

   Specifies the value in the selected column on which you want to filter. Some drop-down lists contain options for the values in the column selected in the Column field. If there are no options available, you must enter a text string in the field.

   Note:

   - The Value field is not case-sensitive.
   - The string must match exactly—partial matches are not returned. For example if you want to the Blueprints table to display all Apache Blueprints, entering Apache will not return any Blueprints.
   - Wildcards are supported. You can use an asterisk (*) or a percent sign (%) as a wildcard, for example Apache* returns all Blueprints that begin with Apache (Apache Tomcat Servlet Engine, Apache HTTP Server, and so on).

3. (Optional) Add additional filter criteria to create a more complex filter:
   a. Select one of the following options:
      - And—Specifies that table displays objects that match the entries in both pairs of Column and Value fields.
      - Or—Specifies that table displays objects that match entries in either pair of Column and Value fields.
   b. Select an option from the second Column drop-down menu.
   c. Enter or select a value in the second Value column.

      For example, if you create a filter on the Blueprints page with the first Column field set to Blueprint Name and the first Value field set to Apache*, and the second pair of fields set to Blueprint Version and 1.0.0, selecting the And option would display all Apache Blueprints with a version of 1.0.0. If you select the Or option, the table would display all Blueprints that begin with Apache (regardless of what version they are), and all Blueprints that are version 1.0.0 (regardless of what their name is).
4. Click Refresh.

   The table displays the rows of objects that match your filter criteria.

To clear a filter and display all table data

1. Open any of the eight tab pages.

   The page displays a corresponding table.
2. Do one of the following:
   - Click Reset to clear the filter fields.
   - Select the blank entry from both of the Column drop-down lists (the blank option is the first entry on the menu, it appears above the first text option).
3. Click Refresh.

   The filter is cleared and the table displays the first 50 rows of table data (rows 51 and above are displayed on different pages, 50 rows to a page).

Common Table Actions

Each of the CA Configuration Automation Management tab pages (Services, Servers, Networks, Profiles, Jobs, Blueprints, Reports, and Remediation) contains a table that displays details about the corresponding objects. Each table contains a Table Actions drop-down menu that contains page-specific table actions and the following three options that are common to all tables:

- Export to Excel
- Print
- Configure Table View

These options are described in the sections that follow.

Export Table Data to Excel

You can export table data and column headings to a Microsoft Excel spreadsheet to share CA Configuration Automation data with people who are not configured as CA Configuration Automation users.

To export table data to Excel

1. Open the tab page the contains the table that you want to export.
2. Select Export to Excel from the Table Actions drop-down list.

   The File Download window appears and prompts to open or save the file.
3. Do one of the following:
   - Click Save, enter a name and location for the file, and then click Save.

     The file is saved in the specified location.
   - Click Open.

     The table data displays in Excel. Select Save As if you want to save the exported data as a file.

Print Table Data

You can print table data if you want to have a paper copy.

Follow these steps:

1. Open the tab page that contains the table that you want to print.
2. Select Print from the Table Actions drop-down list.

   The table data is sent to your printer.

Create Table View

Tab pages that contain a table (for example, the Servers table) also contain a Create Table View option on the Table Actions drop-down list. You can use this option to define custom table views that display the table contents according to your personal preferences.

To create table views

1. Click the Management link then any tab that displays a tab page that contains an element table (for example, the Servers tab).

   The tab page appears and, in this example, contains the Servers table.
2. Select Create Table View from the Table Actions drop-down list.

   The Details page of the Create Table View wizard appears.
3. Enter the following information in the corresponding field, then click Next:

   Name

   Specifies a name for the table view.

   Refresh Interval

   Specifies the rate (in seconds) at which the table is automatically refreshed.

   Page Size

   Specifies the maximum number of rows per page in the table.

   Sort Column

   Specifies which column is used to determine the sort order. For example, if you select the Server Name column, the server names are sorted alphabetically. If you select the Creation Date/Time column, the server names are sorted chronologically.

Sort Order

Specifies Ascending or Descending. For example if a column was sorted alphabetically, and the Sort Order was set to Ascending, the order would be A through Z.

Shared View

Specifies whether this view is available to all users, or only to the table view creator.

The Columns page appears with all available columns displayed in the Selected Columns field (that is, by default, tables display all available columns).

4. Double-click one or more columns in the Selected Columns field that you want to remove from this custom view.

   The selected columns are moved to the Available Columns field.

5. Click Next.

   The Filter page appears.

6. Create a filter by selecting options from the drop-down menus or typing in the following fields:

   Column

   Specifies the column in the table on which you want to filter. The drop-down list contains an option for each column in the table on which you can filter.

   Value

   Specifies the value in the selected column on which you want to filter. Some drop-down lists contain options for the values in the column selected in the Column field. If there are no options available, you must enter a text string in the field.

   Note:

   - The Value field is not case-sensitive.
   - The string must match exactly—partial matches are not returned. For example if you want the Blueprints table to display all Apache Blueprints, entering Apache will not return any Blueprints.
   - Wildcards are supported. You can use an asterisk (*) as a wildcard, for example Apache* returns all Blueprints that begin with Apache (Apache Tomcat Servlet Engine, Apache HTTP Server, and so on).

7. (Optional) Add additional filter criteria to create a more complex filter:

   a. Select one of the following options:

      - And—Specifies that table displays objects that match the entries in both pairs of Column and Value fields.
      - Or—Specifies that table displays objects that match entries in either pair of Column and Value fields.

   b. Select an option from the second Column drop-down list.

   c. Enter or select a value in the second Value column.

   For example, if you create a filter on the Blueprints page with the first Column field set to Blueprint Name and the first Value field set to Apache*, and the second pair of fields set to Blueprint Version and 1.0.0, selecting the And option would display all Apache Blueprints with a version of 1.0.0. If you select the Or option, the table would display all Blueprints that begin with Apache (regardless of what version they are), and all Blueprints that are version 1.0.0 (regardless of what their name is).

8. Click Finish.

   The custom Table View is created and appears in the Table Views table.

Chapter 4: Service Management

This section contains the following topics:

Services Overview

A service is a named instance of a collection of software components running on one or more managed CA Configuration Automation servers. Services typically are created to represent a function, for example, all the servers and software that are used in an organization's Online Banking system.

CA Configuration Automation displays services and their components in a highly structured and annotated format that includes:

- Server hardware, network, and storage details
- Key component parameters and configuration details

- Macros
- Constraint rules
- Relationships and dependencies
- File system, registry, and database elements
- Sub-components inventory

Create Services Using the Wizard

The Create Service wizard defines a service and groups discovered software and managed servers into a logical business entity. The business entity can be the software and server to run a real-time inventory control solution, or to run an FTP site.

Follow these steps:

1. Click the Management link, and then the Services tab.
2. On the Services tab, select Create Service from the Table Actions drop-down list.
3. On the Service page of the Create Service wizard, complete the following fields:

   Name

   Specifies a name for the new service.

   Description

   Describes or identifies the service and its purpose.

   Management Profile

   Assigns a management profile to the service.

   Note: To assign the profile that the CA Configuration Automation administrator designated as the default, select Use Default Profile. The product does not predefine default profiles; the CA Configuration Automation administrator must designate them.

   Discovery Enabled

   Specifies whether the components in the service are updated during discovery.

   Management Enabled

   Specifies whether the components in the service are updated during refresh.

   Business Owner

   Specifies the owner of the service.

   Business Process

   Specifies the business operation that is associated with the service.

IT Owner

Specifies the IT organization that is responsible for the service components.

Location

Specifies the geographic location of the service.

Notes

Displays any additional information that you want to associate with the service.

4. Click Next.

   The Servers page appears, listing the available servers and the selected servers:

   Available Servers

   Lists all the servers that CA Configuration Automation manages in your enterprise network.

   Selected Servers

   Lists all the servers that are part of the service and server group.

5. Add or remove servers from the service as follows:

   - To move one or more servers from the Available Servers column to the Selected Servers pane, select them and then click the down arrow.
   - To move all the servers to the opposite column, click the up arrow.

6. Click Next.

7. On the Server Groups page, double-click one or more server groups to add to the service from the Available Server Groups column.

   The selected server group moves to the Selected Server Groups column.

   You can also add or remove server groups from the service as follows:

   - To move a server group from either column to the opposite column, click the group and then click the left or right single arrow.
   - To move all the server groups from either column to the opposite column, click the double left or right arrows.

8. Click Finish.

   The product creates the service and adds it to the Services table.

Create Services and Management Profiles from the Service Profiler

A configuration manager can use the Service Profiler UI to create a service and an associated management profile. The Service Profiler graphically represents the elements (servers, clusters, applications, and so on) similarly to the Visualization functionality in the CA Configuration Automation Server UI.

The following illustration shows the tasks that are used to create a service and management profile with the Service Profiler.

Create Services and Management Profiles Using the Service Profiler

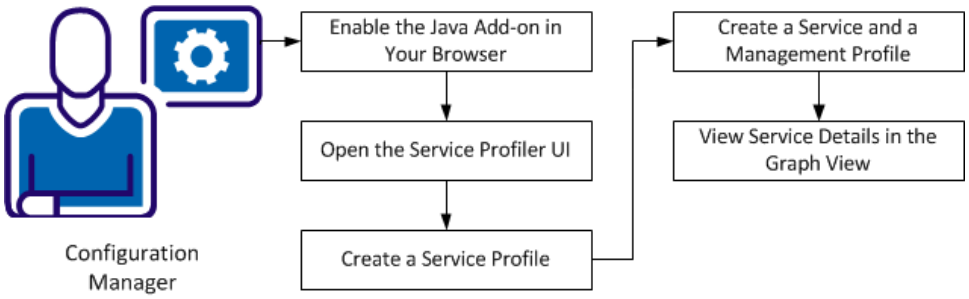

A configuration manager can perform the following tasks to create services and management profiles using the Service Profiler:

1. Enable the Java Add-on in your browser (see page 44)
2. Open the Service Profiler (see page 45)
3. Create a service profile (see page 45)
4. Create a service and a management profile (see page 51)
5. View service details in the Graph View (see page 55)

Enable the Java Add-on in Your Browser

Before you use Service Profiler or the visualization functionality, enable the Java Add-on in your browser.

Follow these steps: (Internet Explorer)

1. Open Internet Explorer and select Tools, Manage Add-ons.
2. In the Manage Add-ons dialog, locate Java Plug-In 2 SSV Helper, select it, and click Enable.
3. In the Enable Add-on dialog, click Enable to confirm the selection.

4. In the Manage Add-ons dialog, click Close.
5. Close and reopen the browser.

Follow these steps: (Firefox)

1. Open Firefox, click the Firefox drop-down menu.
2. Click Add-ons, then click Plugins.
3. Locate Java Platform *<version_number>* in the list of plugins, then click Enable.
4. Close and reopen the browser.

Open the Service Profiler

Use one of the following methods to open the Service Profiler:

- From the CA Configuration Automation Server UI, click the Management panel, click the Services tab, and then select the Services link. In the Services table, select Service Profiler from the Table Actions drop-down list.
- Enter the following URL in your browser address field:

  `http://`*`CA Configuration AutomationServerName`*`:`*`port`*`/TransformUI.jsp`

  For example, you could enter http://ccaServer1:8080/TransformUI.jsp.

The Service Profiler opens in a new browser window. The profiler contains the following panes:

- Service Profile (this pane contains the Details panel and the Graph Content panel)
- Graph View
- Service View

Create a Service Profile

You can create service profiles in the Service Profiler and can use them as templates to create services and management profiles. For example, assume that you commonly create services that are based on server operating systems that are combined with the following other service elements:

- Port numbers
- Discovered components
- Discovered applications

You can create a service profile for each operating system in your environment, then use the service profile to create a service and management profile.

Follow these steps:

1. In the Service Profiler, click + in the Service Profile pane to create a service profile.
2. Click + to expand the Graph Content panel, then click one or more of the following panels to select corresponding profile elements to add:

   Servers

   Contains the following subpanels:

   Discovered Servers

   Displays all the servers that are known to CA Configuration Automation. These servers were discovered in an NDG Discovery operation, imported from another CA Configuration Automation instance, or added manually. Select discovered servers on the following tabs:

   Inclusions tab: In the left pane, double-click a discovered server to *add* to the service.

   Exclusions tab: In the left pane, double-click a discovered server to *exclude* from the service.

   The selected servers move to the right pane.

   The Discovered Servers subpanel also includes the Limit Displayed Relationships to Target Servers/Applications Only check box. Select the check box to add only the selected servers. Servers with relationships to the selected servers are not added to the service.

   Operating Systems

   Displays the operating systems of servers that are known to CA Configuration Automation. Select operating systems on the following tabs:

   Inclusions tab: In the left pane, double-click an operating system to *add* to the service.

   Exclusions tab: In the left pane, double-click an operating system to *exclude* from the service.

   The selected operating systems move to the right pane.

   Clusters

   Displays the server clusters that are known to CA Configuration Automation. Select server clusters on the following tabs:

   Inclusions tab: In the left pane, double-click a server cluster to *add* to the service.

   Exclusions tab: In the left pane, double-click a server cluster to *exclude* from the service.

   The selected server clusters move to the right pane.

Virtualization

Contains the Nonvirtual Servers, VM Hosting Servers, and VM Guest Servers check boxes. Use these options to select the types of servers to include in the service.

Server State

Contains the following check boxes where you can select the servers with the following states: New (selected by default), Managed (selected by default), Unmanaged, and Imported.

Server Groups

Displays the server groups that are known to CA Configuration Automation. Select server groups on the following tabs:

Inclusions tab: In the left pane, double-click a server group to *add* to the service.

Exclusions tab: In the left pane, double-click a server group to *exclude* from the service.

The selected server groups move to the right pane.

Discovered Ports

Displays the discovered ports that are known to CA Configuration Automation. Select discovered ports on the following tabs:

Inclusions tab: In the left pane, double-click a discovered port to *add* to the service.

Exclusions tab: In the left pane, double-click a discovered port to *exclude* from the service.

The selected discovered ports move to the right pane.

The Discovered Ports panel also includes the Only Show Discovered Ports with Relationships check box. Select the check box to add only the selected discovered ports for which NDG Softagent NETSTAT processing found established server connections. When the check box is cleared, the panel displays all available open ports that NDG Softagent processing found.

Port Mapping

Displays the ports from the Communication Mappings table. The table shows which ports map to which applications in your environment. Select ports to map on the following tabs:

Inclusions tab: In the left pane, double-click a mapped port to *add* to the service.

Exclusions tab: In the left pane, double-click a mapped port to *exclude* from the service.

The selected mapped ports move to the right pane.

The Port Mapping panel also includes a field that lets you define filtering criteria for the list. For example, enter *8* to filter the list to include only ports that begin with the number 8 (80, 8080, 8081, 8031, and so on).

Discovered Applications

Displays the discovered applications that are known to CA Configuration Automation. Components included with the application are not included in this list. For example, CA Configuration Automation includes a version of Tomcat. CA Configuration Automation is listed in the Discovered Applications pane, and Tomcat is not. Tomcat is listed in the Discovered Components pane. Select discovered applications on the following tabs:

Inclusions tab: In the left pane, double-click a discovered application to *add* to the service.

Exclusions tab: In the left pane, double-click a discovered application to *exclude* from the service.

The selected discovered applications move to the right pane.

The Discovered Applications panel also includes the Only Show Discovered Applications with Relationships check box. Select the check box to add only the selected discovered applications. When the check box is cleared, the panel displays all available applications that CA Configuration Automation manages.

Application Mapping

Displays the predefined and user-defined application mappings that are known to CA Configuration Automation. Select application mappings on the following tabs.

Inclusions tab: In the left pane, double-click an application mapping to *add* to the service.

Exclusions tab: In the left pane, double-click an application mapping to *exclude* from the service.

The selected application mappings move to the right pane.

The Application Mapping panel also includes a field that lets you define filtering criteria for the list. For example, enter *mi* to filter the list to include only applications with names that begin with the letters mi (Microsoft Internet Explorer, Microsoft SQL Server Database, and so on).

Discovered Components

Displays the discovered components that are known to CA Configuration Automation. Select discovered components on the following tabs:

Inclusions tab: In the left pane, double-click a discovered component to *add* to the service.

Exclusions tab: in the left pane, double-click a discovered component to *exclude* from the service.

The selected discovered components move to the right pane.

The Discovered Components pane also includes a field that lets you filter the list. For example, enter *mi* to filter the list to include only components with names that begin with the letters mi (Microsoft Internet Explorer, Microsoft SQL Server Database, and so on).

Blueprints

Displays the following tabs:

Blueprints

Displays all predefined and custom blueprints. Select blueprints on the following tabs:

Inclusions tab: In the left pane, double-click a blueprint to *add* to the service.

Exclusions tab: In the left pane, double-click a blueprint to *exclude* from the service.

The selected blueprints move to the right pane.

Blueprint Groups

Displays the user-defined blueprint groups. Select blueprint groups on the following tabs:

Inclusions tab: In the left pane, double-click a blueprint group to *add* to the service.

Exclusions tab: In the left pane, double-click a blueprint group to *exclude* from the service.

The selected blueprint groups move to the right pane.

Categories

Displays the predefined blueprint categories. Select predefined blueprint categories on the following tabs:

Inclusions tab: In the left pane, double-click a predefined blueprint category to *add* to the service.

Exclusions tab: In the left pane, double-click a predefined blueprint category to *exclude* from the service.

The selected predefined blueprint categories move to the right pane.

Storage

Displays the following tabs:

Storage Systems

Displays all storage systems. Select storage systems on the following tabs:

Inclusions tab: In the left pane, double-click a storage system to *add* to the service.

Exclusions tab: In the left pane, double-click a storage system to *exclude* from the service.

The selected storage systems move to the right pane.

Managers

Displays the server running the storage manager software. Select storage managers on the following tabs:

Inclusions tab: In the left pane, double-click a manager to *add* to the service.

Exclusions tab: In the left pane, double-click a manager to *exclude* from the service.

The selected managers move to the right pane.

Vendors

Displays the manufacturer or vendor of the storage device. Select vendors on the following tabs:

Inclusions tab: In the left pane, double-click a vendor to *add* to the service.

Exclusions tab: In the left pane, double-click a vendor to *exclude* from the service.

The selected vendors move to the right pane.

3. Click Apply. The application displays the selected elements in the following places:

   Graph View pane

   Includes a legend to describe the graphical elements.

   Service View pane

   Displays the selected servers and blueprints in expandable folders.

4. In the Graph View pane, review the contents of the service profile. If necessary, repeat step 2 to add or remove elements.

5. In the Service Profile pane, click Save As (), enter a name and description in the Save As dialog, then click OK.

   The service profile is created. You can open the profile in the Service Profiler or on the Visualization tab of the Dashboard panel in the CA Configuration Automation Server UI.

Create a Service and a Management Profile with the Service Profiler

Use the Service Profiler to create a service and a management profile to use with a service. Using the Service Profiler has the following advantages over using the Create Service wizard:

- The Service Profiler creates a graphic representation of the service elements each time you save the service.
- When you add a server to the service with the Service Profiler, you can include any servers with a relationship to the added server.

Follow these steps:

1. In the Service Profiler, do one of the following actions:
   - Click the Open icon () in the Service Profile pane to create a service that is based on an existing service profile. In the Open Service Profile dialog, select a profile, then click OK.
   - Click (+) in the Service Profile pane to create a service.
2. Select the elements for the service from the panels in the Graph Content panel as described in step 2 of the preceding (see page 45) section.
3. In the Graph View pane, review the contents of the service. If necessary, repeat Step 2 to add or remove elements.

4. Click Apply. The application displays the selected elements in the following places:

   Graph View pane

   Includes a legend that describes the graphical elements in the service.

   Service View pane

   Displays the servers and blueprints in the service in expandable folders.

5. In the Service View pane, click Save Service.
6. On the Service tab in the Save Service dialog, enter the following information:

   Name

   Defines a name for the new service.

   Description

   Describes the service and its purpose.

   Management Profile

   Defines the management profile that is assigned to the service.

   By default, the application appends the contents of the Name field with Service Profiler (that is, *name*ServiceProfiler) in the Management Profile field. You can edit the name or you can select an existing profile (including the Use Default Profile option) from the drop-down list.

   Business Owner

   Defines the owner of the service.

   Business Process

   Defines the business operation that is associated with the service.

   IT Owner

   Defines the IT organization that is responsible for the service components.

   Location

   Defines the geographic location of the service.

   Note

   Defines additional information to associate with the service.

7. In the Save Services dialog, click the Discovery Options tab.

8. On the Discovery Options tab, enter the following information to define the management profile:

Collect Hardware Information

Specifies whether information about the physical host computer is discovered and managed in CA Configuration Automation.

Collect Networking Information

Specifies whether information about the network is discovered and managed in CA Configuration Automation.

Collect Storage Information

Specifies whether information about storage devices is discovered and managed in CA Configuration Automation.

Collect Server Properties

Specifies whether information about servers is discovered and managed in CA Configuration Automation. The Server Properties that NDG discovery operations returns are stored on the following pages that are linked from the Servers tab page:

- Server Details
- Virtualization
- Network Adapters
- Hardware
- Applications
- Services/Daemons
- Open Ports
- Relationships

The Server Properties component is created and updated under the following scenarios for servers in the Managed state:

- Servers that are discovered or updated during network discovery (NDG)
- Discovery or a refresh that a Management Profile initiates (scheduled, manual, or using the SDK)

Search the Registry

Specifies whether discovery operations associated with this profile search the Windows Registry.

Follow Symbolic Links

Specifies whether the discovery operations that are associated with the profile search networks and file systems that are connected with a symbolic link.

Include Network Drives

Specifies whether the discovery operations that are associated with the profile search network drives.

Refresh Previously Discovered Components

Specifies whether the discovery operation updates software components already in inventory.

Discovery Time Limit

Defines the interval after which to terminate a discovery operation. Most discovery operations require only a few minutes. The duration of the search can become long when searching large file systems, particularly if the Agent Priority is set to its lowest value. If a search exceeds the Discovery Time Limit value, the discovery operation ends without returning results.

Default Search Root

Defines the directory in which to begin the search for profiles that are used in the following situations:

- On some combination of computers that runs Windows, UNIX, and Linux
- Only on computers that run Linux or UNIX

Windows Search Root

Defines the folder in which to begin the search for profiles that are used only on Windows computers.

File Search Depth

Defines the number of directory levels that you want to search below the search root. If you leave this value blank, CA Configuration Automation searches all directories under the search root.

Perform Discovery on Servers

Specifies to run a discovery operation on the servers to which the profile is assigned.

Use Discovered Components from CCA Database

Specifies to search for service components that were previously discovered and are already stored in the CA Configuration Automation Database.

Perform Discovery Now

Specifies whether to run a discovery operation when you click OK to create the service.

9. Click OK.

   CA Configuration Automation creates a service and a management profile. The application stores the service in the Services table on the Services tab. The application stores the management profile in the Management Profiles table on both the Services tab and Servers tab. If you selected the Perform Discovery Now check box, a discovery operation begins.

View Service Details in the Graph View

The Service Profiler Graph View pane contains navigation tools that let you view specific graphic elements in the service. A legend is included to identify the elements. Depending on the element type, you can double-click or point to an element to display details about it.

Follow these steps:

1. In the Service Profiler, use the following buttons on the Service Profiler Graph View pane to isolate elements with details you want to view:

    **Select**

   Switches to the selection tool from the selected mode. The Select tool is active by default.

    **Pan**

   Lets you drag-and-drop the graph view to reposition the view in the main pane. The Pan tool does not change magnification.

    **Marquee Zoom**

   Lets you magnify a specific area of the graph by drawing a rectangle around the area to magnify.

Interactive Zoom

Lets you increase or decrease the magnification to show greater detail (fewer elements) or less detail (more elements).

- To zoom in (that is, to increase the magnification), click an outer edge of the graph and drag toward the center.
- To zoom out (that is, to decrease the magnification), click the center of the graph and drag toward the outer edge.

Fit

Increases or decreases the magnification to display the entire graph in the pane.

Tree Layout

Displays graph elements in a hierarchical tree layout. The layout attempts to show the flow of elements (for example, from top to bottom) and arranges similar elements in levels (that is, horizontal rows).

Symmetric Layout

Arranges the elements in a graph using quadrilateral symmetry (that is, it divides elements equally on both the horizontal and vertical axes).

Orthogonal Layout

Arranges the elements in a graph orthogonally (that is, at right angles to other elements).

Circular Layout

Determines the relationship of the elements in a graph, then arranges them in separate circles. The circles are then arranged in a radial tree layout.

Print

Prints the current graph view.

Print Preview

Displays a preview of the current graph as it will appear when printed.

Print Settings

Displays the Print Setup dialog so you can configure printer settings.

Save As Image

Displays the Save As Image dialog so you can save the current view to a graphic file.

Display Edge Labels

Displays text labels over the connecting lines that represent relationships.

Expand Nodes

Collapse Nodes

Expands or collapses all nodes in the graph elements depending on the current setting.

- If the nodes are currently expanded, the button displays the collapse node icon 
- If the nodes are currently collapsed, the button displays the expand node icon 

Note: To use the Select tool to expand or collapse individual nodes, click the appropriate button next to an element in the Graph View pane.

2. Do one or more of the following tasks:
   - Point to an arrow between elements to display relationship details about the elements that the arrow connects.
   - Point to an application icon to display details about the application.
   - Click an application icon to display a navigation icon. To display the child element of the application (for example, the server that hosts it), double-click the navigation icon. To return to the previous view, right-click the server icon and select Go To Parent.
   - To display the filter that created the current graph, right-click an application icon. You can edit the filter if you want to change the contents of the graph.
   - To expand all nodes, click the Expand Nodes button above the graph.
   - To expand a specific node, double-click an individual node icon in the graph.
   - To display server details (including the assigned Management Profile, status IP address, operating system, and so on), point to a server icon .
   - To display relationship details, right-click a server icon.

Delete Services

You can delete services that you are no longer using.

Note: When you delete services, all snapshots associated with the selected services are also deleted.

To delete one or more services

1. Click the Management link.

   The Management panel appears and displays the Services tab by default.

2. Click the check boxes next to the services that you want to delete, and then select Service Actions, Delete Services from the Select Actions drop-down list.

   You are prompted to delete the services.

3. Click OK.

   The specified services are deleted.

Export Services

You can export a service definition to use it in another instance of CA Configuration Automation. Only the service properties are exported, and the associated components, relationships and snapshots are not exported. The service properties include Service name, Service Description, Management Profile Name, Discovery Enabled, Management Enabled, Server Names and Server Group Names.

Follow these steps:

1. Click the Management link, then click the Services tab.
2. Select one or more services that you want to export from the Services page.
3. Click Select Actions, Service Actions, and Export Services.
4. Click Save in the File Download dialog.

   Default name for the export file is Services.csv.

5. Edit the file name if desired, select the location to save the file, and then click Save.

   The service definition is exported to the selected location.

Import Services

You can import a service definition as a CSV file exported from another instance of CA Configuration Automation.

Follow these steps:

1. Click the Management link, Services tab.
2. Click Table Actions, and select Import Services.
3. Click Browse in the Import Services dialog to select the CSV file that contains the service definition, and click OK.
4. Select the Overwrite Existing Services check box to overwrite the service with the same name.

   Select this option if you want to retain the changes that you made to the service on another instance of CA Configuration Automation.

   Note: The service is not updated if any current activity (for example, discovery or refresh) is running on the service.

   The file is imported and the services appear in the Services page. Default management profile is set to Imported Service if the corresponding Management Profile does not exist in the CA Configuration Automation database.

View Service Components

The Service View Components option displays a expandable tree view of a service and the associated components.

To display service components in a tree view

1. Click the Management link, then click the Services tab.

   The Services tab page appears.
2. Click the check box next to the services whose software components you want to view and then select View Components from the Select Actions drop-down list.

   The View Components and Configurations window opens in a new browser page. It displays the selected services in the Services pane, and details about the first service are displayed in the right pane. The Service pane displays a tree that uses the following icons to identify elements in the tree:

    Identifies the service being displayed.

    Identifies the servers in the selected service.

    Identifies the Blueprints for the server listed directly above this icon.

 Identifies the primary folders for the Blueprint directly above this icon.

 Identifies the subfolders and components folders.

 Identifies the parameters.

Note: The components listed for the service are defined by the service's Management Profile.

3. Do the following at any level of the tree:
   - Click an icon or the element name to view details of that element in the right pane.
   - Click any plus sign (+) next to an element to expand the tree to display the next level of elements.

Create a Service Snapshot

The Take Snapshot option creates a point-in-time copy of a service. You supply the snapshot name, and CA Configuration Automation time stamps the snapshot to identify the copy uniquely.

Snapshots can help you track and identify configuration changes made to your services when you run Change Detection and Compare operations to compare the point-in-time copy to the current server data to see what changed.

To create and save a service snapshot

1. Click the Management link, then click the Services tab.

   The Services tab page appears.

2. Click the check box next to the services you want to create a snapshot of, and then select Management Actions, Take Snapshot from the Select Actions drop-down list.

   The Snapshots Service dialog appears.

3. Enter a name for the snapshot in the Snapshot Name field.
4. (Optional) Enter a description for the snapshot in the Snapshot Description field.
5. Click OK.

   A snapshot is created for the selected service. You can view it in the Snapshots table on the Snapshots tab of the Service Details page.

Run Service Change Detection

Service Change Detection detects how a service has changed over a period of time by using *snapshots* (point-in-time copies) of service data to provide a detailed account of all detected configuration changes, as well as file system changes, including file ownership, file permission, and file modification times.

Service-based Change Detection provides options for finding differences between the current service data and a snapshot, or any two snapshots.

Component elements that have the Time Variant filter set in the corresponding Component Blueprint are the only items not checked for changes during time-based Change Detection. Log files and data modified at runtime are examples of Time Variant-filtered elements.

Note: Running Change Detection using the procedures in this section is considered running the operation manually.

To run Change Detection to identify changes to a service

1. Click the Management link, then click the Services tab.

   The Services tab page appears.

2. Click the check box next to the services you want to search for changes, and then select Management Actions, Run Change Detection from the Select Actions drop-down list.

   The Change Detection Across Time dialog appears.

3. Select a Source Snapshot from the following options:

   Current Data

   Specifies the current system data available for the service is used as the source.

   Most recent snapshot

   Specifies the snapshot with the most recent timestamp is used as the source.

   Second most recent snapshot

   Specifies the snapshot with the second most recent timestamp is used as the source.

   Most recent snapshot on a specific date

   Specifies the date of the snapshot to use as the source. If you select this option, the Source Snapshot Date field appears. Select the date of the snapshot you want to use. If there are multiple snapshots available on the specified date, the most recent snapshot on that date is used.

Selected snapshot

Specifies a user-selected snapshot to use as the source. If you select this option, the Source Snapshot field appears. Select the snapshot you want to use.

Baseline

Specifies the snapshot designated as the Baseline is used as the source.

Gold Standard

Specifies the snapshot designated as the Gold Standard is used as the source.

Silver Standard

Specifies the snapshot designated as the Silver Standard is used as the source.

Bronze Standard

Specifies the snapshot designated as the Bronze Standard is used as the source.

4. Select a Target Snapshot (the options are the same as those listed in step 3).
5. Select one of the following options:

Include All Component Blueprints

Specifies that all component blueprints are searched for changes to the service.

Select Component Blueprints by Name

Specifies that one or more component blueprint is searched for changes to the service. If you select this option you must select the blueprints in the Available Blueprints column, then click the single right-facing arrow to move it to the Select Blueprints column.

6. Click Next.

The Filters page appears.

7. Select the following options to specify what differences are included in the Change Detection results:

No children comparison if a hierarchical object exists on source or target only

Specifies whether the Change Detection operation is performed on the child components if the object only exists in one of the services. When this option is selected, the operation ignores an object if it is not part of both services.

All Differences

Specifies that all differences to the service are included.

Component Inventory Differences Only

Specifies that only the services that are in the component inventory are included.

Filters

Specifies that one or more of the following objects are included:

- Folders—Accept the default setting (All) or click the Select option, then select one or more folders to search for changes.
- Categories—Accept the default setting (All) or click the Select option, then select one or more categories to search for changes. Categories are assigned in the Component Blueprint and are the organizational groupings to which an element belongs.
- Weights—Accept the default setting (All) or click the Select option, then select the weights to search for changes. Weights are assigned in the Component Blueprint and represent the relative importance of an element. Unweighted elements (no weight assigned) are considered Medium.

You can press Ctrl+click to select multiple non-consecutive list entries, or Shift+click to select multiple consecutive list entries.

8. Click Finish.

   The Change Detection operation runs, and the results as described in Viewing the Results of Change Detection and Compare Operations (see page 63).

   .

Viewing the Results of Change Detection and Compare Operations

The results of Server Change Detection, Service Change Detection, Compare Servers, and Compare Service operations are all displayed in the same interface. The results are shown in a multi-tab window that contains the following tabs:

- **Summary**—Shows the server name and whether the Change Detection or Compare operation was performed on the server or service's current data or a snapshot, the target (snapshot or current data), the result (differences found or not found), and the total number of differences found.
- **Summary by Weight**—Shows the number of changes for the following color-coded weights: Low, Medium, and High.

  You can display the Summary by Weight in either a graphical chart or a table view.
- **Summary by Category**—Shows the number of changes by category.

  You can display the Summary by Category in either a graphical chart or a table view.

- **Tree**—Shows the results in a hierarchical tree view. If there are less than 20 components, the tree view is automatically expanded to the component level, otherwise you must expand the tree to see the components.

  Color-coded icons next to an item in the tree view indicate that a difference was found between the source and target. CA Configuration Automation displays a legend to the right of the tree that states what the following color-coded indicators mean:

  
  Element has different values in the source and target

  
  Element exists only in the source

  
  Element exists only in the target

  
  Element is not different, but a child of the element is different

  - Right-click the server or service name and select Expand All to expand all nodes of the tree.
  - Expand the tree until you find the element that is different, or click to expand all children of the element.
  - Click an element with one of the aforementioned color-coded icons.

    The details about the element appear in the details pane.

    If the element exists in both the source and target and the values are different, the details pane highlights the differences in yellow and displays the value for the source and target elements.

- **Flat Table**—Shows the results in a table that contains the following columns: Select, Server, Type of Change, Software, Parameters, Manage Files, Managed Data, Configuration Files, Configuration Executables, Configuration Data, and Registry.
  - Click the server name in the Server column to display details of the changes on that server.
  - Click the option button in the Select column and select Print or Export to Excel from the Table Action drop-down list.

Compare Services or Components

The Compare Services or Components feature finds the differences between the following:

- Current service data
- Most recent snapshot
- Second most recent snapshot
- Most recent snapshot on a specific date
- Selected snapshot
- Baseline snapshot
- Gold standard snapshot
- Silver standard snapshot
- Bronze standard snapshot

Component elements that have the Time Variant filter set in the corresponding Component Blueprint are the only items not checked for differences. Log files and data modified at runtime are examples of Time Variant-filtered elements.

To compare services or components

1. Click the Management link, then the Services tab.

   The Services tab page appears.

2. Click the check box next to the services you want to compare, and then select Management Actions, Run Compare from the Select Actions drop-down list.

   The Compare Services wizard appears with the first service you selected listed in the Source Service field and the second service listed in the Target Service field.

3. Select a Source Snapshot from the following options:

   Current Data

   Specifies the current system data available for the service is used as the source.

   Most recent snapshot

   Specifies the snapshot with the most recent timestamp is used as the source.

   Second most recent snapshot

   Specifies the snapshot with the second most recent timestamp is used as the source.

Most recent snapshot on a specific date

Specifies the date of the snapshot to use as the source. If you select this option, the Source Snapshot Date field appears. Select the date of the snapshot you want to use. If there are multiple snapshots available on the specified date, the most recent snapshot on that date is used.

Selected snapshot

Specifies the snapshot to use as the source. If you select this option, the Source Snapshot field appears. Select the snapshot you want to use.

Baseline

Specifies the snapshot designated as the Baseline is used as the source.

Gold Standard

Specifies the snapshot designated as the Gold Standard is used as the source.

Silver Standard

Specifies the snapshot designated as the Silver Standard is used as the source.

Bronze Standard

Specifies the snapshot designated as the Bronze Standard is used as the source.

4. Select a Target Snapshot (the options are the same as those listed in step 3).
5. Click Next.

   The Components page appears.
6. Select the following options to use for the comparison:

   Include All Component Blueprints

   Specifies that all component blueprints are included.

   Select Component Blueprints by Name

   Specifies that one or more component blueprints are included. If you select this option you must double-click a blueprint in the Available Blueprints column to move it to the Select Blueprints column. If you select this option you also activate the Select Components area.

   Enable Component Selection

   Specifies whether you want to include specific software components in the comparison.

   Source Component

   Specifies the source software component to use for the comparison.
7. Click Next.

   The Filters page appears.

8. Select the following options to specify what differences are included in the comparison results:

   No children comparison if a hierarchical object exists on source or target only

   Specifies whether the comparison operation is performed on the child components if the object only exists in one of the services. When this option is selected, the operation ignores an object if it is not part of both services.

   All Differences

   Specifies that all differences in the services are included.

   Component Inventory Differences Only

   Specifies that only the components that are in the component inventory are included.

   Filters

   Specifies that one or more of the following objects are included:

   - Folders—Accept the default setting (All) or click the Select option, then select one or more folders to compare.
   - Categories—Accept the default setting (All) or click the Select option, then select one or more categories to compare. Categories are assigned in the Component Blueprint and are the organizational groupings to which an element belongs.
   - Weights—Accept the default setting (All) or click the Select option, then select the weights to compare. Weights are assigned in the Component Blueprint and represent the relative importance of an element. Unweighted elements (no weight assigned) are considered Medium.

     You can press Ctrl+click to select multiple non-consecutive list entries, or Shift+click to select multiple consecutive list entries.

9. Click OK.

   The service comparison runs, and the results appear as described in Viewing the Results of Change Detection and Compare Operations (see page 63).

Run Service Rule Compliance

Service-based Rule Compliance enables you to check service and snapshot data against the following:

- Default Value rules
- Data Type rules
- Constraint rules defined in Component Blueprints
- Constraint rules defined in both services and Component Blueprints

Note: Running Rule Compliance using the procedure described in this section is considered running the operation manually. To schedule Rule Compliance operations, see Create Rule Compliance Jobs.

To run service-based Rule Compliance

1. Click the Management link, then the Services tab.

   The Services tab page appears.

2. Click the check box next to the services you want to check for compliance, and then select Management Actions, Run Rule Compliance from the Select Actions drop-down list.

   The Service page of the Run Rule Compliance wizard appears.

3. Select the lowest severity level of messages you want Rule Compliance to report from the Severity drop-down list.

   The available Rule Category options (described in step 5) are determined by the severity level you select in this step.

   Information

   Displays Information, Warning, Error, and Critical messages. Enables you to select either Default Value Rules or Data Type Rules in step 5.

   Warning

   Displays Warning, Error, and Critical messages. Disables Data Value Rules in step 5.

   Error

   Displays Error and Critical messages. Disables Data Value Rules in step 5.

   Critical

   Displays only Critical messages. Disables both Default Value Rules and Data Type Rules in step 5, and requires that you define Explicit Rules as described in step 7.

4. Select one of the following options from the Remediation drop-down list to specify whether you want to use remediation to reset non-compliant values:

   None

   Specifies that non-compliant values appear in the Rule Compliance results, but does not use remediation to reset these values.

   Rule Value Only

   Specifies that remediation is used to reset non-compliant values to the values that are defined in the rules.

   Rule Value or Blueprint Default Value

   Specifies that remediation is used to reset non-compliant values to the values that are defined in the rules. If an explicit rule is not defined for the component, the non-compliant value is reset to the default value defined in the Component Blueprint.

5. Create a set of rules against which to run Rule Compliance by selecting one of the following categories in the Rule Category area (if you selected Critical in the Severity drop-down list, these options are not available):

   Default Value Rules

   Verifies current service or snapshot values against specified default values.

   When default values are specified in a Component Blueprint, CA Configuration Automation automatically creates rules that check to see if the actual value deviates from the default value. Default rule deviations show as Information messages in the results.

   Data Type Rules

   Verifies current service or snapshot values against specified values for the corresponding data type.

6. Accept the default Explicit Rules settings or click the Select Rule Groups option and double-click the rule group you want to use in the Available Rule Groups column to move it to the Selected Rule Groups column.

   Blueprint Rules

   Verifies current service or snapshot values against constraint rules defined in Component Blueprints.

   Includes both user-defined rules and built-in rules, such as data type checking.

   Instance Rules

   Verifies current service or snapshot values against constraint rules defined in the service and Component Blueprints.

7. Select the snapshot you want to use to establish compliance.

   Note: If you selected multiple services in step 2, and want to use snapshot data, select each service's most recent snapshot on specified date option and enter a date in the Snapshot Date field. If there is no snapshot for the selected services on the specified date, CA Configuration Automation displays an error message on the results page. If there are multiple snapshots available on the specified date, the most recent snapshot on that date is used.

8. Click Next.

   The Components page appears.

9. Select one of the following Component Blueprint options:

   Include All Component Blueprints

   Specifies that all component blueprints are used on the selected services.

   Select Component Blueprints by Name

   Specifies that one or more component blueprints are used on the selected services. If you select this option you must select the blueprints in the Available Blueprints column, then click the single right-facing arrow to move it to the Select Blueprints column.

10. Click Next.

    The Filters page appears.

11. Accept the default setting (All) or create a filter to determine which of the following items are considered by the Rule Compliance operation:

    Folders

    Specifies whether all folders or only the selected folders are searched by the Rule Compliance operation.

    Categories

    Specifies whether all categories or only the selected categories are searched by the Rule Compliance operation. Categories are assigned in the Component Blueprint and are the organizational groupings to which an element belongs.

    Weights

    Specifies whether all weights or only the selected weights are searched by the Rule Compliance operation. Weights are assigned in the Component Blueprint and represent the relative importance of an element. Unweighted elements (no weight assigned) are considered Medium.

    You can press Ctrl+click to select multiple non-consecutive list entries, or Shift+click to select multiple consecutive list entries.

12. Click OK.

    The Rule Compliance results appear as described in Viewing the Results of Rule Compliance Operations (see page 71).

Viewing the Results of Rule Compliance Operations

The results of Service Rule Compliance and Server Rule Compliance operations are displayed in the same interface. The results are shown the following tabs:

- **Summary**—Displays the server name or service name. Also displays whether the Rule Compliance operation was performed on the server or service current data or a snapshot, and whether rule violations were found.
- **Summary by Weight**—Displays the number of violations and passed for the following color-coded weights: Low, Medium, and High.

  You can display the Summary by Weight in either a graphical chart or a table view.
- **Summary by Category**—Displays the number of violations and passed by category.

  You can display the Summary by Category in either a graphical chart or a table view.
- **Tree**—Displays the results appear in a hierarchical tree view.

  The results are grouped based on the severity level of the failed rules groups as per the severity. When you expand the tree, the following nodes appear:

  Rule Failure By Severity—Expand the node to view the server where Rule Compliance found the rule violation (displayed in brackets) and its element value.

  The Rule Compliance results tree lets you perform the following tasks:

  - View the description that was provided during the rule creation when you move the mouse pointer over a value.
  - Click any value in the tree to display the details in the right pane. You cannot modify elements from the results window.
  - View the severity level. The following cons indicate the severity level:

    
    Specifies the Information severity level and displays the rule violation text.

    
    Specifies the Warning severity level and displays the rule violation text.

    
    Specifies the Error severity level and displays the rule violation text.

    
    Specifies the Critical severity level and displays the rule violation text.

- **Flat Table**—Displays the results in a table that contains the following columns: Select, Server, Software, Total Failures, Total Passed, and Total Run.
  - Click the server name in the Server column to display details of the changes on that server.
  - Click the option button in the Select column and select Print or Export to Excel from the Table Action drop-down list.
- **Rule Exceptions**—Displays the rule exceptions that are excluded in the rule compliance operation for selected Servers and Services. You can filter the rule exceptions that are based on the software, server, and service name.

Refresh Services

The refresh services operation obtains the most current service component data.

Note: The basic state of a component does not change as a result of a refresh operation—Inventoried components remain Inventoried components and Managed components remain Managed components. The state of a component changes only as a result of the Management Profile directing components to be managed or as a result of service-related discovery of a component.

To refresh service software components

1. Click the Management link, then the Services tab.

   The Services tab page appears.

2. Click the check box next to the services whose components you want to refresh, and then click Refresh (above the Services table).

   The Current Activity column of the Services table displays Refresh during the operation. The column is blank when the service refresh is complete.

Run Service Discovery

The Service Discovery operation searches the servers that are defined in the service for updates to the service's software components.

Note: The procedure that is described in this section is a manual operation to run a Service Discovery operation immediately. Management Profiles are typically used to automate the Service Discovery operations using the scheduling information that is provided in the profile. See Management Profiles (see page 79) for more information.

Follow these steps:

1. Click the Management link, then the Services tab.

   The Services tab page appears.

2. Select the services whose components you want to discover, or update by discovery. Then select Management Actions, Run Discovery from the Select Actions drop-down list.

   Note: Ensure that you set the corresponding Management Profile to the Enabled state. During the service discovery, the previously discovered servers or components of a service are delete if the servers are not defined in a service.

   The Current Activity column of the Services table displays Discovery during the operation.

Configure Component Discovery Using Telnet

A configuration manager can create an Access Profile with the Telnet access mode to use for software component discovery operations.

Configure Component Discovery Using Telnet

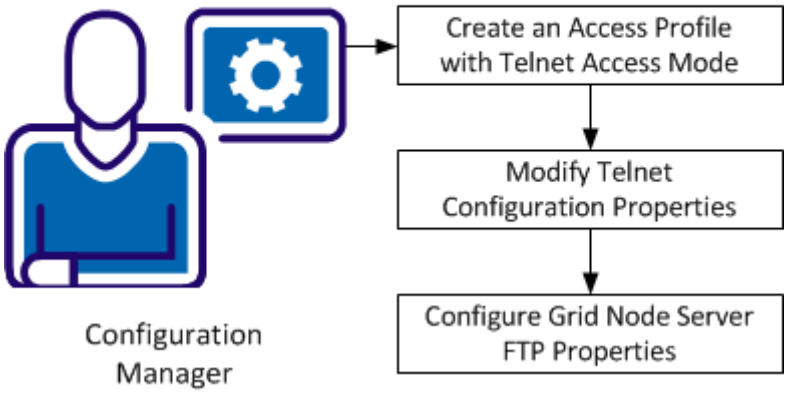

A configuration manager can perform the following tasks to configure CA Configuration Automation to perform component discovery using telnet access:

1. Create an Access Profile with Telnet access mode (see page 74)
2. Modify telnet configuration properties (see page 75)
3. Configure Grid Node Server FTP properties (see page 75)

Create an Access Profile with Telnet Access Mode

Follow these steps:

1. In the CA Configuration Automation Server UI, click the Management link, then click the Access Profile link on the Servers tab.
2. Complete *one* of the following tasks:
   - On the Access Profile tab, select Create Access Profile from the Table Actions drop-down list.
   - To modify an existing profile, select the name in the Profile Name field.
3. On the Access Mode page, select Telnet from the Access Mode drop-down list and complete all required fields as described in the *Online Help*.

   Notes:

   - The Connection Timeout value is important to the Telnet access mode and CA Configuration Automation sets it to a high value by default. CA Configuration Automation uses the Connection Timeout value to stop idle connections.
   - The Prompts field values locate and identify the prompts that CA Configuration Automation displays while establishing a Telnet session with the targeted server.

Modify Telnet Configuration Properties

1. In the CA Configuration Automation Server UI, click the Administration link (top right), the Configuration tab, and then the Properties link.
2. On the properties tab page, click the right-facing arrow (>) above the Default Value column to display the second page of properties (page one shows properties 1 through 50, page two shows properties 51 through 100).
3. To modify a property, click the entry in the Value column for the following values in the cca group:

   telnet.connection.retries

   Specifies the number of connection retries if the Telnet connection is lost for any reason in the middle of the discovery process.

   Default: 3. Increase this value to 6 if telnet drops connections during discovery.

   telnet.read.timeout_secs

   Specifies the maximum time interval to wait, to gather the results after issuing a command.

   Default: 2. Increase this value to 4 if telnet drops connections during discovery.

   telnet.read.byte_to_byte_delay_secs

   Specifies the maximum time interval to wait for the next byte while reading the results. This value is used only if the Look for Prompts option is *not* selected in the Access Profile.

   Default: 900. Increase this value to 1500 if telnet drops connections during discovery.

Configure Grid Node Server FTP Properties

1. In the CA Configuration Automation Server UI, click the Administration link (top right), the Configuration tab, and then the Properties link.
2. On the properties tab page, click the right-facing arrow (>) above the Default Value column to display the second page of properties (page one shows properties 1 through 50, page two shows properties 51 through 100).
3. To modify a property, click the entry in the Value column for the following values in the grid group:

   ftp.account

   Specifies the account name used to connect to the FTP Server.

   ftp.password

   Specifies the password used for connecting to the FTP Server.

ftp.port

Specifies the port number where the FTP Server is listening.

ftp.root

Specifies the root path of the FTP Server (FTP site directory).

Report Actions for Services

The Reports Actions operation lets you run reports from Services. You can perform the following actions to run reports:

Run Compliance Report

Run the compliance reports for the selected services. By default, the corresponding Compliance Rule groups filter criteria is applied when you run the report.

Run Report

Run a report from the available reports. Based on the selected report, provide the necessary inputs to generate the report.

To run Report Actions for services

1. Click the Management link.
2. Click the Services tab in the Management panel.
3. Select services in the Services page, and then select Reports Actions from the Select Actions drop-down list.
4. Select any one of the following actions:
   - Run Compliance Report—Run compliance report for the selected services.
   - Run Report—Run reports from the available reports.
5. If you select Run Compliance Report, select the Report Name, and click Ok.
6. If you select Run Report, select one of the following options in the Select Report panel:
   - Templates—Run a report using any of the predefined report templates.
   - Saved Reports—Run the reports that are saved earlier.

   Report is generated.

Stop Service Discovery

You can manually stop a service discovery operation if it is taking too long, or for any other reason.

To stop a service discovery operation

1. Click the Management link.

   The Management panel appears and displays the Services tab by default.

2. Click the check box next to the service whose discovery operation you want to stop, and then select Management Actions, Stop Discovery from the Select Actions drop-down list.

   The Current Activity column of the Services table displays Discovery while the operation is running, and is blank when the operation stops.

View Relationships and Management Operation Results in the Visualization UI

You can display a graphical representation of relationships and management operations (for example, Change Detection or Rule Compliance) in the Visualization UI.

Follow these steps:

1. In the CA Configuration Automation Server UI, click the Management panel, then click one of the following locations:
   - Services tab
   - Servers tab
   - Server Groups tab
   - Storage tab
2. On the tab page that appears, select one or more check boxes in the table.

   For example, if you selected the Services tab, select one or more services in the Services table.
3. From the Select Actions drop-down list, select Visualization, then select one of the following visualization templates:

   Services tab:
   - Service Communication Relationships
   - Service Change Detection
   - Service Rule Compliance

   Servers tab and the Server Groups tab:
   - All Server Relationships
   - Server Change Detection
   - Server Communication Relationships
   - Server Configuration Relationships
   - Server Rule Compliance
   - Server Storage Relationships
   - Server Virtualization

   Storage tab:
   - Server Storage Management Relationships
   - Server Storage Relationships

   For a description of each template, see Select a Visualization Template (see page 379).

Management Profiles

Management profiles provide the operational rules for updating services, discovering servers and software components, and for managing those components when they are located. These operation rules include:

- Which software components or categories of software components are discovered.
- Whether or not previously discovered, but no longer existing components, are removed from the component inventory.
- How frequently and when software component discovery is performed.
- What depth and level the discovery takes place and its importance relative to other operations.
- Which software components or categories of software components are managed.
- Which, when, and how frequently CA Configuration Automation management operations (Refresh, Snapshot creation and deletion, Change Detection, Compare, and Rule Compliance) are run automatically.

CA Configuration Automation includes the following predefined Management Profiles:

- Discover and Manage Application Server
- Discover and Manage CA Software
- Discover and Manage OS and Hardware (this profile is defined as the default and is assigned to a Network Profile)
- Discover and Manage Relational Databases
- Discover and Manage Server Properties
- Discover and Manage Virtualization
- Discover and Manage Web Server

The predefined profiles are created by system_user. The default profile is assigned to each newly discovered server or newly created service. You can also create your own profiles and assign them to servers and services as described later in this section.

Create Management Profiles

You can create a Management Profile to automate and configure how service and server discovery operations perform their search, and how they respond to the objects they locate.

To create a Management Profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Management Profiles link (below the main tabs).

   The existing profiles appear in the Management Profiles table.

3. Click Table Actions, then select Create Management Profiles.

   The Profile page of the Create Management Profile wizard appears.

4. Enter the following information in the corresponding field:

   Name

   Specifies the name of the Management Profile.

   Description

   Describes the Management Profile.

   Default

   Specifies whether this Management Profile is assigned to newly discovered servers or services when they are placed in the managed mode. When the check box is empty, the profile is *not* designated as the default.

   Enabled

   Specifies whether this Management Profile is available to be assigned, assigned to a Network Profile. or set as the default profile. When the check box is empty, the profile is *not* enabled.

   Enable Catalyst Integration

   Specifies whether CA Configuration Automation exports configuration items (CIs) to Catalyst and other consuming products using the CCA connector.

   Attributes Profile

   Specifies the Catalyst Attributes Profile to use when the Catalyst Integration is enabled. The options are Use Default, All Catalyst Attributes, and the user-defined profiles stored in the Catalyst Attributes Profiles table.

5. Click Next.

   The Blueprints page appears.

   Note: The top of the Blueprint page contains links to Blueprints Groups and Categories pages. You can specify any combination of Blueprints, Blueprint Groups, and Categories to use in the discovery operations managed by this profile.

6. Do one of the following:

   - Click the Move All Blueprints to Discover check box if you want all of the blueprints to be used for discovery and managed.
   - Ensure the Move All Discover to Manage check box is checked, then double-click one or more blueprints listed in the Available column to move them to the Discover and Manage columns.
   - Clear the Move All Discover to Manage check box, then double-click one or more blueprints in the Available column to move them to the Discover column.

   As an alternative to double-clicking blueprints to move them, click a blueprint and then click the single left- or right-facing arrow button to move the blueprint to an adjacent column. Click the double left- or right-facing arrow button to move all blueprints to an adjacent column.

   The selected blueprints are moved into the Discovery or Manage columns, and used by the profile for discovery and management operations.

   Note:

   - The components specified by these Blueprints (and those associated with the Blueprint Groups described in step 7) are the components that appear on the Components tab of the Server Detail and Service Details pages. The All Components tab on the Server Details page shows all of the components installed on a server, some of which may be from a Service Management Profile.
   - Operations including Snapshot, Change Detection, Compare, and Rule Compliance also reference the components in the Management Profile for the service or server.

7. (Optional) Click the Blueprint Groups above the check boxes and repeat step 6, but this time specify the blueprint groups you want to include in discovery and management operations.

   The selected blueprint groups are moved into the Discovery or Manage columns, and used by the profile for discovery and management operations.

8. (Optional) Click the Categories link above the check boxes and repeat step 6, but this time specify the categories you want to include in discovery and management operations.

   The selected categories are moved into the Discovery or Manage columns, and used by the profile for discovery and management operations.

9. Click Next.

   The Discovery Options page appears.

10. Select the following discovery options or enter the appropriate information in the corresponding field:

    Collect Hardware Information

    Specifies whether information about the physical host computer is discovered and managed in CA Configuration Automation.

    Collect Networking Information

    Specifies whether information about the network is discovered and managed in CA Configuration Automation.

    Collect Storage Information

    Specifies whether information about storage devices, storage managers, and their relationships are discovered and managed in CA Configuration Automation.

    Collect Server Information

    Specifies whether information about servers is discovered and managed in CA Configuration Automation. The Server Properties returned by NDG discovery operations are stored on the following pages linked from the Servers tab page:

    - Server Details
    - Virtualization
    - Network Adapters
    - Hardware
    - Applications
    - Services/Daemons
    - Open Ports
    - Relationships

    The Server Properties component is created and updated under the following scenarios for server's with a state of Managed:

    - Servers discovered or updated during network discovery (NDG)
    - Discovery or refresh initiated by a Management Profile (either scheduled, manual, or using the SDK)

    Search the Registry

    Specifies whether discovery operations associated with this profile search the Windows Registry.

Follow Symbolic Links

Specifies whether discovery operations associated with this profile search networks and file systems connected with a symbolic link.

Include Network Drives

Specifies whether discovery operations associated with this profile search network drives on Windows computers. If you select this option, you must also complete the following steps:

a. Log on to the CA Configuration Automation Agent host computer as a user with administrator privileges for the host.

b. Stop the agent service if it is running.

c. Edit the agent.conf file to disable the restart property as follows:

```
#restart every # calls
#restart=1000
```

d. Open the command prompt and change to the CCA Agent installation directory.

e. Run the agent using the following command:

```
CCagent.exe -p agent.conf
```

f. Run the Management Profile on the Agent host to discover components on the network drives.

Agent Priority

Specifies the priority given to the CA Configuration Automation Agent during discovery operations associated with this profile. As the CA Configuration Automation Agent searches the server file system for matches, it can impact the performance of other file system operations on the target server. To control the impact, you can set the priority to one of five following levels:

- Highest (fastest)—The CA Configuration Automation Agent is unrestricted. Generally, this priority has a brief, but typically acceptable impact on the target server activity and can be used in most circumstances. This is the default setting.
- High, Medium, Low—Each lower priority reduces the impact on file system activity by approximately 20%.
- Lowest (slowest)—The search yields to other processes. This results in a discovery with almost no impact on existing file system activity, but discoveries take substantially longer to complete.

Discovery Time Limit

Specifies the amount of time before the discovery is terminated. Most discovery operations take only a few minutes, however when searching large file systems, the duration of the search can become long, particularly if the Agent Priority is set to its lowest value. If a search exceeds the limit, the discovery ends—there is no such thing as a partial discovery.

Default Search Root

Specifies the directory to begin the search for profiles that are used on a combination of Windows and UNIX or Linux computers (or if the profile is only used on Linux and UNIX computers).

Windows Search Root

Specifies the folder to begin the search for profiles that are used only on Windows computers.

File Search Depth

Specifies the number of directory levels below the search root that you want to search. If you leave this blank, all directories under the search root are searched.

Enable Pruning

Specifies whether the Server Software Inventory Pruning Mode is enabled. If the check box is marked, pruning is enabled, and missing software components are removed from the inventory. If you disable this feature, all discovered software remains in the inventory regardless of whether future discovery operations find it or not. Clearing the check box effectively turns the inventory into a history of software on the server.

Pruning Mode

Specifies which of the following modes is used to process missing software components:

- Mark missing components, but keep in inventory—Discovery operations that fail to verify the existence of previously inventoried software components mark those components as missing, but retain them in the inventory.
- Delete missing components from inventory—Components that are no longer verified by discovery are removed from the server's inventory. Note that components that are part of a managed service are not removed even if this option is selected.

Perform discovery on servers

Specifies to run a discovery operation on the servers to which the profile is assigned.

Use discovered components from CCA database

Specifies to search for service components that have been previously discovered and are already stored in CA Configuration Automation.

11. Click Next.

The Management Options page appears and displays the Change Detection and Compare management options.

Note: The top of the Management Options page also contains links to Rule Compliance, Snapshots, and Filters pages.

12. Select the following Change Detection and Compare options:

Component Inventory Differences Only

Specifies whether change detection or compare operations only return results for components already being managed.

No children comparison if a hierarchical object exists on source or target only

Specifies whether the comparison operation is performed on the child components if the object only exists in one of the services or servers. When this option is selected, the operation ignores an object if it is not part of both services or servers.

Current data with

Specifies that the change detection operation searches for changes between the current data and one or more of the following snapshots:

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

Compare to Another Server

Specifies whether the comparison is made with another server. When checked, you must specify the server or snapshot with which to compare.

Compare to Another Service

Specifies whether the comparison is made with another service. When checked, you must specify the service or snapshot with which to compare.

Change Detection Adhoc

Specifies that when the management profile is run manually, an alert is sent over the CCA Catalyst Connector to a consuming CA product (for example, CA Spectrum Service Assurance).

This field is hidden unless the sdk.events.enabled property is set to true as described in View and Edit CA Configuration Automation Properties.

Change Detection Scheduled

Specifies that when the management profile is run as a scheduled job, an alert is sent over the CCA Catalyst Connector to a consuming CA product (for example, CA Spectrum Service Assurance).

This field is hidden unless the sdk.events.enabled property is set to true as described in View and Edit CA Configuration Automation Properties.

The selected options are included in the profile.

13. Click the Rule Compliance link above the check boxes.

    The Rule Compliance page appears.

14. Select the following rule compliance options:

Run Rule Compliance

Specifies whether the rule compliance operation is performed when this profile is run.

Rule Severity

Specifies the severity level of messages you want rule compliance to report. Rule compliance returns all messages for the level you specify and above (for example, if you want to see Error and Critical messages, select Error; if you want to see all messages, select Information).

Remediation

Specifies whether violations located by the rule compliance operation are reset on the target service or server. The Remediation options are as follows:

- None—Specifies that no remediation action is performed after the rule compliance operation.
- Rule Value—Specifies that the value that violated the rule is reset to the value defined in the rule.
- Default Value From Blueprint—Specifies that the value that violated the rule is reset to the default values from the software component's Blueprint.

Default Value Rules

Verifies current service or snapshot values against specified default values.

When default values are specified in a Blueprint, CA Configuration Automation automatically creates rules that check to see if the actual value deviates from the default value. Default rule deviations show as Information messages in the results.

Data Type Rules

Verifies current service or snapshot values against specified values for the corresponding data type.

Rule Category

Specifies that explicit rules from one or both of the following categories are used for the rule compliance operation:

Blueprint Rules

Verifies current service or snapshot values against constraint rules defined in Blueprints.

Includes both user-defined rules and built-in rules, such as data type checking.

Instance Rules

Verifies current service or snapshot values against constraint rules defined in the service and Blueprints.

Rule Groups

Specifies the rules defined in rule groups are used for the rule compliance operation. Click the Select Rule Groups option and double click the rule group you want to use in the Available Rule Groups column to move it to the Selected Rule Groups column.

Rule Compliance Adhoc

Specifies that when the management profile is run manually, an alert is sent over the CCA Catalyst Connector to a consuming CA product (for example, CA Spectrum Service Assurance).

This field is hidden unless the sdk.events.enabled property is set to true as described in View and Edit CA Configuration Automation Properties.

Rule Compliance Scheduled

Specifies that when the management profile is run as a scheduled job, an alert is sent over the CCA Catalyst Connector to a consuming CA product (for example, CA Spectrum Service Assurance).

This field is hidden unless the sdk.events.enabled property is set to true as described in View and Edit CA Configuration Automation Properties.

The selected options are included in the profile.

15. Click the Snapshots link above the check box.

    The Snapshots page appears.

16. Click the Create Snapshot check box if you want this profile to create a service or server snapshot when run.

    If you select this option, the following fields are activated:

    Maximum Count

    Specifies a limit to the number of snapshots that are stored in CA Configuration Automation. When selected, the counter can be set to the desired number.

    Maximum Age

    Specifies a limit to the age of snapshots that are stored in CA Configuration Automation. When selected, the counters can be set to the desired number of days, weeks, or months.

    Note: If you select the Create Snapshot option, and do not select either the Maximum Count or Maximum Age options, an unlimited number of snapshots are stored in CA Configuration Automation until they are manually deleted.

    The selected options are included in the profile.

17. Click the Filters link above the check box.

    The Filters page appears.

18. Accept the default setting (All) in the Folders, Categories, and Weights areas, or click Select in one or more area, and select the options you want included in the profile (use Ctrl+click or Shift+click to select multiple options).

    - Folders are how CA Configuration Automation presents discovered software components. They are contained in a hierarchical tree view. If you choose the Select (Folders) option, the rule compliance, change detection, and compare operations only search the specified folders for this profile.
    - Categories are assigned in the Component Blueprint and are the organizational groupings to which an element belongs. If you choose the Select (Categories) option, the rule compliance, change detection, and compare operations only search the specified categories for this profile.
    - Weights are assigned in the Component Blueprint and represent the relative importance of an element. Unweighted elements (no weight assigned) are considered Medium. If you choose the Select (Weights) option, the rule compliance, change detection, and compare operations only search the specified weights for this profile.

    The selected options are included in the profile.

19. Click Next.

    The Scheduling page appears and displays the Discovery tab by default.

20. Click the Run Management After Discovery check box if you want the management operation associated with this profile to run on the same schedule as the discovery operation, when the discovery operation completes.

    If you select this option you do not need to define a schedule for the management operation.

21. Define the schedule for automatically running discovery operations using this profile by selecting one of the following from the Frequency drop-down list:

    Not Scheduled

    Specifies that the profile does not run automatically. It can be run manually or scheduled in the future.

    Once

    Specifies that the profile is run automatically one time. If you select this option, you also need to specify when it is run in the Time field.

    Minutes

    Specifies that the profile is run on a recurring basis using at an interval defined in minutes. If you select this option, you also need to specify the following:

    - Start Time—Specify the time the profile starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
    - Begin Date—Specify the date the profile is first run.
    - End Date—Specify the date the profile is run for the last time.
    - Recur every # minutes—Specify the interval at which the profile runs.

    For example, if you want the profile to run every 10 minutes starting at 11:00 p.m., you would specify a Start Time of 11:00:00PM, and specify Recur every 10 minutes. The profile would run at 11:00 p.m., 11:10 p.m., 11:20 p.m., 11:30 p.m., and so on until the end of the hour (midnight in this example). If the current profile has not finished running by the time the next interval occurs, the next run waits until the previous one completes, and then starts.

    Hourly

    Specifies that the profile is run on a recurring basis using at an interval defined in hours. If you select this option, you also need to specify the following:

    - Start Time—Specify the time the profile starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
    - Begin Date—Specify the date the profile is first run.
    - End Date—Specify the date the profile is run for the last time.
    - Recur every # hours—Specify the interval at which the profile runs.

For example, if you want the profile to run every four hours throughout the day starting at 11:00 p.m., you would specify a Start Time of 11:00:00PM, and specify Recur every 4 hours. The profile would run at 11:00 p.m., 3:00 a.m., 7:00 a.m., 11:00 a.m., 3:00 p.m., and 7:00 p.m.. If the current profile has not finished running by the time the next interval occurs, the next run will wait until the previous one completes, and then start. Also note that if the Start Time has already passed in the current day, the profile runs immediately, then resumes the recurring schedule you specify.

Daily

Specifies that the profile is run on a recurring basis using at an interval defined in days. If you select this option, you also need to specify the following:

- Start Time—Specify the time the profile starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
- Begin Date—Specify the date the profile is first run.
- End Date—Specify the date the profile is run for the last time.
- Recur every # days—Specify the interval at which the profile runs.

Weekly

Specifies that the profile is run on a recurring basis using at an interval defined in weeks. If you select this option, you also need to specify the following:

- Start Time—Specify the time the profile starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
- Begin Date—Specify the date the profile is first run.
- End Date—Specify the date the profile is run for the last time.
- Recur every # weeks—Specify the interval at which the profile runs.

Monthly

Specifies that the profile is run on a recurring basis using at an interval defined in months. If you select this option, you also need to specify the following:

- Start Time—Specify the time the profile starts to run.
- Begin Date—Specify the date the profile is first run.
- End Date—Specify the date the profile is run for the last time.
- Recur every # months—Specify the interval at which the profile runs.

22. Define the notification that is sent when the profile is run in the following fields:

    Notification Profile

    Specifies the notification profile to use when discovery operations using this profile are run as scheduled. For information about creating notification profiles, see Create Notification Profiles (see page 98).

    Subject

    Specifies the subject line of the email message that is sent by the selected notification profile.

    The schedule for discovery operations associated with this profile is defined.

    - If you clicked the Run Management After Discovery check box, skip to step 25.
    - If you did not click the Run Management After Discovery check box, you must define a management operation. Continue with step 23.

23. Click the Management tab.

    The Management page appears.

24. Repeat step 21, but this time define the schedule for the management operations associated with this profile.

25. Click Finish.

    The management profile is created and appears in the Management Profiles table.

Run Management Profiles on Services

You can run a Management Profile manually to update the configuration settings of software components in a service instead of waiting for the scheduled time specified in the profile. You have the following options for manually running Management Profiles:

- Run Management Profile with Discovery
- Run Management Profile without Discovery

The procedures associated with these options are described in the sections that follow.

Note: You can also use the Run Discovery option to run a Service Discovery manually without using a Management Profile. If you use a Management Profile, you get the benefits of the profile including change detection, compare, rule compliance, snapshots, pruning mode, search mode, component refresh mode, and so on.

- For information about running a discovery without a Management Profile, see Run Service Discovery (see page 133).
- For information about defining Management Profiles, see Create Management Profiles (see page 80).

Run Management Profile with Discovery

You can manually run a management profile with discovery instead of waiting for the scheduled time specified in the profile. When you run the profile with the discovery operation, the configuration settings of existing software components in a service are updated. Additionally, new software components that are discovered are added to the software inventory.

To run the management profile with discovery manually

1. Click the Management link, then the Services tab.

   The Services tab page appears.

2. Click the check box next to one of more services you want to update as defined by the management profile, and then select Management Actions, Run Management Profile with Discovery from the Select Actions drop-down list.

   Note: The Management Profile that corresponds with the service must be enabled.

   The Current Activity column of the Services table displays Management Profile. The column is blank when the operation is complete.

Run Management Profile without Discovery

You can manually run a management profile without a discovery operation when you want to update existing services with information about their associated servers and software components without searching for new components on the server.

To manually run the management profile without discovery

1. Click the Management link, then the Services tab.
2. Click the check box next to one or more services you want to update as defined by the management profile, and then select Management Actions, Run Management Profile without Discovery from the Select Actions drop-down list.

   Note: The Management Profile that corresponds with the service must be enabled.

   The Current Activity column of the Services table displays Management Profile. The column is blank when the operation is complete.

Assign Management Profiles to Services

You can assign existing management profiles to services to control how software components and servers are discovered and managed.

To assign a management profile to a service

1. Click the Management link.

   The Management panel appears and displays the Services tab by default.

2. Click the check box next to the services to which you want to assign a management profile, and then select Service Actions, Assign Management Profile from the Select Actions drop-down list.

   The Assign Profiles dialog appears.

3. Click Change, select a management profile from the drop-down list, and then click OK.

   The Management Profile column of the Services table displays the name of the assigned profile.

Set a Management Profile as the Default

You can designate one Service Management Profile to be the *default.* The default profile is automatically assigned to newly discovered servers and newly created services. This saves you from having to create a new profile and make specific profile assignments for each new server or service.

To set a Management Profile as the default profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Management Profiles link (below the main tabs).

   The existing profiles appear in the Management Profiles table.

3. Click the check box next to the Management Profile that you want to be the default profile, then click Select Actions and select Set As Default.

   The Is Default column displays a check mark next to the profile you selected.

Enable Management Profiles

You can enable Management Profiles to make them available to assign to services or servers.

To enable a Management Profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Management Profiles link (below the main tabs).

   The existing profiles appear in the Management Profiles table.

3. Click the check box next to one or more Management Profiles that you want to enable, then click Select Actions and select Enable Profiles.

   The Is Enabled column displays a check mark next to the profiles you selected.

Disable Management Profiles

You can disable Management Profiles to prevent them from being assigned to services or servers.

To enable a Management Profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Management Profiles link (below the main tabs).

   The existing profiles appear in the Management Profiles table.

3. Click the check box next to one or more Management Profiles that you want to disable, then click Select Actions and select Disable Profiles.

   The Is Enabled column displays an X next to the profiles you selected.

Delete Management Profiles

You can delete Management Profiles when you no longer have a need for them. If you think you may reuse the profiles in the future, you can disable them instead of deleting them.

To delete a Management Profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Management Profiles link (below the main tabs).

   The existing profiles appear in the Management Profiles table.

3. Click the check box next to one or more Management Profiles that you want to delete, then click Select Actions and select Delete Profiles.

   The selected profiles are deleted and removed from the table.

Import Management Profiles

You can import a Management Profile as a Java Archive (JAR) file from another instance of CA Configuration Automation.

To import a Management Profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Management Profiles link (below the main tabs).

   The Management Profiles page appears.

3. Click Table Actions, then select Import Management Profiles.

   The Import Management Profiles dialog appears.

4. Enter or select the following information in the corresponding field:

   JAR File to Import

   Specifies the name of the JAR file that contains the Management Profile you want to import. You can click Browse to navigate to the file.

   Overwrite Existing Management Profiles

   Specifies whether the file being imported overwrites a file with the same name. Select this option if you want the changes made to the profile on another instance of CA Configuration Automation to be retained.

5. Click one of the following buttons:

   Import All

   Imports all of the Management Profiles in the JAR file.

   Import On Selected

   Displays a dialog where you can select the Management Profiles in the JAR file to import.

   The file is imported and the profiles appear in the Management Profiles table.

Export Management Profiles

You can export a Management Profile as a JAR file to use in another instance of CA Configuration Automation.

To export a Management Profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Management Profiles link (below the main tabs).

   The Management Profiles page appears.

3. Click the check box next to the profile you want to export, then click Select Actions and select Export Management Profiles.

   The File Download dialog appears.

4. Click Save.

   The Save As dialog appears and the export JAR file is assigned a default name using the following format:

   ExportDatabaseObject_*<year>*_*<month>*_*<date>*_*<hour>*_*<minutes>*_*<seconds>*.jar

   For example: ExportDatabaseObject_2009_12_29_04_20_00.jar

5. Edit the file name if desired, select the location to save the file, and then click Save.

   The profile is exported to the selected location.

Export Profiles and Objects to Tenants

CA Configuration Automation tenant administrators can export profiles (Management Profiles, Network Profiles, Notification Profiles, and so on) and other objects (including Blueprints and Structure Classes) from any CA Configuration Automation instance to a tenant instance on their master instance. The export can be performed from the Select actions drop-down list in the following places in the CA Configuration Automation Server UI:

- Services tab
  - Management Profiles page
  - Notification Profiles page
- Servers tab
  - Management Profiles page
  - Notification Profiles page
- Networks tab
  - Network Scan Policies page
  - Notification Profiles page
- Blueprints tab
  - Blueprints page
  - Blueprint Groups page
  - Structure Classes page
- Compliance tab
  - Rule Groups page
- Remediation tab
  - Remediation Profiles page

Additionally, the Export to Tenants functionality is available for Dashboards and Visualization objects as described in Export Dashboards and Visualization Objects to Tenants (see page 394).

Follow these steps:

1. Log in to the CA Configuration Automation Server UI as a tenant administrator user, then navigate to the page that contains the profile or object you want to export to a tenant.

   The selected page appears with the objects listed in a table.

2. Click the check box next to one or more objects that you want to export, the select Export to Tenants from the Select Actions drop-down list.

   The Import dialog appears.

3. Click the Overwrite Existing Objects check box if you want to overwrite the existing object with the new version on the tenant instance.
4. Double-click one or more tenants in the Available Tenants column.

   The selected tenant appears in the Selected Tenants column.
5. Click OK to the import the selected profile or object into the selected tenants.

   The Results pane confirms the import was successful, or displays an error description.

Notification Profiles

Notification Profiles can be created and assigned to services and servers. They automate how, and to whom notifications are sent when components in the services or servers to which they are assigned change or are updated.

Create Notification Profiles

You can define Notification Profiles that send email notifications when server, service, or network operations complete, or if the operations fail.

Follow these steps:

1. Click the Management link, then click one of the following tabs:
   - Services
   - Servers
   - Network
2. On the selected tab, click the Notification Profiles link.

   The Notification Profiles table displays the existing profiles.
3. Click Table Actions, then select Create Notification Profiles.
4. On the Profile page of the Create Notification Profile wizard, complete the following fields:

   Name

   Defines the name of the Notification Profile.

   Description

   Describes the Notification Profile.

Default

Specifies whether to assign this Notification Profile to newly discovered servers or services when they are put in the managed mode. Select the Default check box to designate this Notification Profile as the default.

Mode

Specifies one of the following notification modes:

Notify On Completion

Sends a notification when the scheduled job finishes.

Notify On Error

Sends a notification if the scheduled job cannot run because an error condition occurred, changes were detected, or rules failed compliance.

Do Not Notify

Sends no notification.

Send Email To

Defines the email address of the person who receives the email notification.

Notification Subject

Defines the subject line of the email notification.

5. Click OK.

   The Notification Profiles table displays the new profile.

Set a Notification Profile as the Default

You can designate one Notification Profile to be the *default.* The default profile is automatically assigned to newly discovered or created services or servers. This saves you from having to create a new profile and make specific profile assignments for each new service.

To set a Notification Profile as the default profile

1. Click the Management link, then the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Notification Profiles link (below the main tabs).

   The existing profiles appear in the Notification Profiles table.

3. Click the check box next to the Notification Profile that you want to be the default profile, then click Select Actions and select Set As Default.

   The Is Default column displays a check mark next to the profile you selected.

Import Notification Profiles

You can import a Notification Profile as a Java Archive (JAR) file from another CA Configuration Automation instance.

Follow these steps:

1. Click the Management link, then click one of the following tabs:
   - Services
   - Servers
   - Network
2. On the selected tab, click the Notification Profiles link.

   The Notification Profiles table displays the existing profiles.
3. Click Table Actions, then select Import Notification Profiles.
4. On the Import Notification Profiles dialog, complete the following fields:

   JAR File to Import

   Defines the name of the JAR file that contains the Notification Profile to import. Click Browse to navigate to the file.

   Overwrite Existing Notification Profiles

   Specifies whether to overwrite a file with the same name. Select this option to retain the profile from another CA Configuration Automation instance.
5. Click one of the following buttons:

   Import All

   Imports all of the Notification Profiles in the JAR file.

   Import On Selected

   Displays a dialog on which to select the Notification Profiles to import from the JAR file.

   The application imports the file and the Notification Profiles table displays the profiles.

Delete Notification Profiles

You can delete Notification Profiles when you no longer need them. If you think you may reuse the profiles in the future, you can disable them instead of deleting them.

To delete a Notification Profile

1. Click the Management link, then the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Notification Profiles link (below the main tabs).

   The existing profiles appear in the Notification Profiles table.

3. Click the check box next to one or more Notification Profiles that you want to delete, then click Select Actions and select Delete Profiles.

   The selected profiles are deleted and removed from the table.

View Scheduled Jobs

The Jobs table displays details about all of the scheduled jobs that are associated with Management Profiles and Network Profiles.

To view scheduled jobs, click the Management Link then do one of the following actions:

- Click the Services tab, and then the Jobs link (below the main tabs)
- Click the Servers tab, and then the Jobs link (below the main tabs)
- Click the Compliance tab, and then the Jobs link (below the main tabs)
- Click the Remediation tab, and then the Jobs link (below the main tabs)
- Click the Reports tab, and then the Jobs link (below the main tabs)
- Click the Jobs tab

The scheduled jobs appear in the Jobs table. The Jobs tab maintains a history of completed and failed scheduled jobs. Click the completed or error jobs links to view the logs for the selected job. The completed and the error jobs are archived. The archiving of the jobs is based on the following configuration information available in the Administration tab:

job.archive.threshold

Specifies the number of records to exceed before a completed job history archive is created.

Default: 500

job.archive.skip.records

Specifies the number of records to skip before job history archiving remainder.

Default: 200

job.archive.minimum.records

Specifies the number of records before a completed job history archive is created.

Default: 200

View the Service Log

CA Configuration Automation maintains a service log that documents each service-based transaction. It can be accessed using the UI for each service being managed by CA Configuration Automation.

To view a service's activity log

1. Click the Management link, the Services tab, and then the Log tab.

   The Log table appears and displays all activity for each service. The default view is to show the most recent service activity in the first row, and the oldest activity in the last row. You can click the sort icons at the top of sortable columns to present the table data differently.

2. (Optional) Create a filter to view specific service events captured in the log as described in Filter Table Views.

   Note: The Filter functionality on this page contains the following additional fields only found on Log pages:

   Start Date/Time

   Specifies the time to begin searching for service activity. For example, if you want to see all service activity during the week you were away on vacation, you would specify the day you left as the start time by clicking the the Start Date/Time check box, clicking the calendar icon, and selecting the day and time you left.

   End Date/Time

   Specifies the time to end searching for service activity. Continuing the previous example, if you want to see all service activity during the week you were away on vacation, you would specify the day you returned as the end time by clicking the the End Date/Time check box, clicking the calendar icon, and selecting day and time you returned.

View and Edit Service Details

You can view and edit details about any service being managed by CA Configuration Automation.

To view and edit service details

1. Click the Management link.

   The Management panel appears.

2. Click the Services tab.

   The Services tab page appears.

3. Click the name of the service whose details you want to view or edit in the Services table.

   The service's Service Details page appears.

4. Edit the fields as appropriate, then click Save.

   The service is updated and a confirmation message appears.

Add Servers to Services

You can add a server to an existing service from the following locations:

- Server Details page
- Service Details page

From the Server Details Page

Follow these steps:

1. Click the Management link.
2. Click the Servers tab.
3. Click the name of the server to which to add a service in the Server table.
4. On the Server Details page, click the Services tab.

   The Available Services and Selected Services columns appear.

   Note: If you cannot select the listed services in the Available Services column, the server is not in the Managed state. Set the server state to Managed before you try to add it to a service.

5. Double-click one or more services in the Available Services column to which to add the selected server.

   The product moves the selected server to the Selected Services column.

   You can also add or remove services from the server as follows:

   - To move a service from either column to the opposite column, click the service and then click the left or right single arrow.
   - To move all the services from either column to the opposite column, click the double left or right arrows.

6. Click Save.

   The product adds the selected servers to the service.

From the Service Details Page

Follow these steps:

1. Click the Management link.
2. Click the Services tab.
3. Click the name of the service to which to add a server in the Services table.
4. On the Service Details page, click the Servers tab.

   The Available Servers and Selected Servers columns appear.

   Note: The Available Servers column only displays servers in the Managed state. Set the server state to Managed before you try to add a service.

5. Add or remove servers from the service as follows:

   - To move one or more servers to the Selected Servers pane, select them from the Available Servers column and then click the down arrow.
   - To move one or more servers from the Selected Servers pane to the Available Servers pane, click the up arrow.

6. Click Save.

   The product adds the selected servers to the service.

Add Server Groups to a Service

You can add a server group to an existing service from a service's Service Details page. Server groups are managed servers that are logically grouped into a single entity to simplify managing them.

To add a server group to a service

1. Click the Management link, then click the Services tab.

   The Services tab page appears.

2. Click the name of the service to which you want to add to a server group in the Services table.

   The selected service's Service Details page appears and displays the Service tab.

3. Click the Server Groups tab.

   The Available Services and Selected Services columns appear.

4. Double-click one or more server group that you want to add to the selected service in the Available Server Groups column.

   The selected service is moved to the Selected Service column.

   Alternatively, you can add or remove a server group from the service as follows:

   - Click a server group in either column and click the left- or right-facing single arrow to move it to the opposite column.
   - Click the double left- or right-facing arrows to move all the server groups to the opposite column.

5. Click Save.

   The server group is added to the service.

Managing Service Components

CA Configuration Automation enables you to manage software components from a Components table that is available on the each service's Service Details page. The table shows only the components that are included in the selected service. A similar view is available for server-centric management of software components from the Server Details page.

You can perform the following component management operations from the Service Details, Components page:

- View Components by Service (see page 106)
- Refresh Services (see page 72)
- Delete Components from Service (see page 107)

View Components by Service

The View Components and Configurations page displays a hierarchical tree presentation of the selected service and the software components included in it. You can navigate the tree to display configuration attributes and settings for the components.

To display service components and configurations in a tree view

1. Click the Management link.

   The Management panel appears.

2. Click the Services tab.

   The Services tab page appears.

3. Click the check box next to one or more services whose components you want to view, then select View Components from the Select Actions drop-down list.

   The Services page displays the selected services in the Services pane, and the first service is displayed in the right pane. The service name appears at the top of the pane.

4. Click a plus sign (+) next to a service.

   The node is expanded and displays the components contained in the service.

5. Click any plus sign (+) to expand the node to display the folders and configuration elements under each service or component, or click the component name to display details about the component.

Refresh Services

The refresh services operation obtains the most current service component data.

Note: The basic state of a component does not change as a result of a refresh operation—Inventoried components remain Inventoried components and Managed components remain Managed components. The state of a component changes only as a result of the Management Profile directing components to be managed or as a result of service-related discovery of a component.

To refresh service software components

1. Click the Management link, then the Services tab.

   The Services tab page appears.

2. Click the check box next to the services whose components you want to refresh, and then click Refresh (above the Services table).

   The Current Activity column of the Services table displays Refresh during the operation. The column is blank when the service refresh is complete.

Delete Components from Services

You can delete software components that you no longer want to manage as part of a service.

To delete one or more components from a service

1. Click the Management link.

   The Management panel appears.

2. Click the Services tab.

   The Services tab page appears.

3. Click the name of the service from which you want delete a component in the Service table.

   The selected service's Service Details page appears and displays the Service tab.

4. Click the Components tab.

   The Components table appears.

5. Click the check box next to the components you want to delete, select Delete Components from the Select Actions drop-down list.

   The Delete Components dialog appears.

6. Do one of the following:

   - Click OK to delete the component from the selected service.

     The component is deleted from the service. It is still managed in the software inventory and may be added to, or already be part of, other services.

   - Click the Delete the Selected Components from the Server check box, then click OK to delete the component from the selected service and from the server.

     - If the selected component is included in other services, a message appears stating the component cannot be deleted from the server.

     - If the selected component is not included in other services, the component is deleted from the selected service, and from the server. It is no longer managed in the software inventory.

View Relationship Details

You can view virtual, static, and dynamic relationships between the servers in a service. When you display details about a service, the Relationships table lists all the types of communications for the servers in the selected service. The Server Name *n* column can use the types listed in the Communication Type column to list the servers in the selected service.

The communication type is configured on the Communication Mapping page as described in View and Edit Communication Mappings.

Follow these steps:

1. Click the Management link, then click the Services tab.
2. Complete one of the following actions:
   - In the Service table, click the name of the service to edit or for which to view details.
   - Ctrl+Click the Service Name field of the service to edit or for which to view details.

   The Details for Service page opens and displays the Services tab.
3. Click the Relationships tab, then click one of the following links:

   Virtualization

   The Virtual Environment Relationships table appears.

   Communication

   The Communication Relationships table appears.

   Configuration

   The Configuration Relationships table appears.

   Note: The Configuration Relationships table shows only the relationships corresponding to the components discovered within the service.

Configuration relationships resolve both the server name and IP address.

- If the CCA managed servers list includes the target server, the target server information includes IPv4 and IPV6 addresses.
- If the CCA managed servers list does not include the target server, the product takes the following actions:
  - The product uses "reverse lookup" to resolve the target server information.
  - The product includes IPv4 and IPV6 addresses.
- If the CCA managed servers list does not include the target server and is not resolvable using "reverse lookup", the IP information is empty.

4. Click a link in one of the following columns to display details about the specific relationship:
   - Virtual Environment (Virtual Environment Relationships table)
   - Communication Type (Communication Relationships table)
   - Parameter Name (Configuration Relationships table)

   You can also take the following actions:
   - Click the Relationships link to return to the Relationships table.
   - Click the Services link to return to the Services table.

Managing Service Snapshots

A service snapshot is a point-in-time copy of a service being managed by CA Configuration Automation. The Snapshot tab of the Service Details page displays a list of all existing service snapshots for the selected service. By default the list sorts the snapshots chronologically with the oldest snapshot at the bottom and the newest snapshot at the top.

You can perform the following snapshot management operations from the Service Details, Snapshots page:

- View Service Snapshots (see page 110)
- Set a Service Snapshot as the Baseline (see page 111)
- Set Service Snapshots as the Gold, Silver, or Bronze Standard (see page 112)
- Remove the Baseline Designation from a Service Snapshot (see page 113)
- Remove a Gold, Silver, or Bronze Standard Designation from a Service Snapshot (see page 114)
- Delete Service Snapshots

View Service Snapshots

You can view any service snapshot you have created.

To view a service snapshot

1. Click the Management link.

   The Management panel appears.

2. Click the Service tab.

   The Service tab page appears.

3. Click the name of the service you want to view snapshot details about in the Service table.

   The selected service's Service Details page appears and displays the Service tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the service snapshots you want to view.

   The View Components and Configurations page is displayed. For information about this page, see Viewing Components and Configurations.

You can also create a filter to search service snapshots.

1. Select a filter option from the Column drop-down-list to search service snapshots in the Filters section. The filter options are as follows:
   - Blueprint Name
   - Created By
   - Designation
   - Snapshot Description
   - Snapshot Name
   - Snapshot Origin
2. Select an option from the Value drop-down list or enter a value.
3. Click And or Or to further refine the filter by selecting another option in the second pair of Column and Value fields.
4. Click Go.

   The service snapshots are listed based on the filter criteria.

Set a Service Snapshot as the Baseline

CA Configuration Automation enables you to designate one service snapshot as the *Baseline*—the snapshot that is used as the reference for service Change Detection Across Time operations.

Notes:

- A similar snapshot feature enables you to designate multiple service snapshots as the Gold, Silver, or Bronze Standard. These snapshots are used as the reference for service comparison operations.
- The same service snapshot can be designated as the Baseline and as the Gold, Silver, or Bronze Standard.
- Baseline designations can also be made for service snapshots.

To designate a service snapshot as the Baseline

1. Click the Management link.

   The Management panel appears.

2. Click the Service tab.

   The Service tab page appears.

3. Click the name of the service from which you want delete a snapshot in the Service table.

   The selected service's Service Details page appears and displays the Service tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the service snapshot you want to designate as the Baseline, then select Set as Baseline from the Select Actions drop-down list.

   The Designation column displays Baseline for the corresponding snapshot. This snapshot is used when the Baseline snapshot option is selected as the source or target snapshot as described in Run Service Change Detection (see page 61).

Set Service Snapshots as the Gold, Silver, or Bronze Standard

CA Configuration Automation enables you to designate service snapshots as the Gold, Silver, or Bronze Standard. These snapshots are used as the reference for service comparison operations.

The gold, silver, and bronze designations do *not* necessarily imply a hierarchy where the Gold Standard is a higher standard than the Silver Standard, and the Silver Standard is a higher standard than the Bronze Standard. While they can be implemented this way, they are designed to give you three different standards to apply to the service snapshots used as the reference for comparison operations.

To increase the flexibility, you can designate more than one service snapshot as the Gold, Silver, or Bronze Standard in the case where you want a specific component of the snapshot to be used for the comparison. For example, you could have a Gold Standard for the operating system, another Gold Standard for the database, a Bronze Standard for Linux mail servers, and another Bronze Standard for Windows mail servers.

Notes:

- Because snapshots are differentiated by timestamp and not by name, you can create multiple snapshots with same name and designate them as the Gold, Silver, or Bronze Standard.
- If you designate multiple snapshots as the Gold, Silver, or Bronze Standard, be careful when naming and creating the description of the snapshots to reflect how they are used.
- A similar snapshot feature enables you to designate one service snapshot as the *Baseline*—the snapshot that is used as the reference for Change Detection operations. There can be only one Baseline snapshot per service.
- The same service snapshot can be designated as the Gold, Silver, or Bronze Standard and as the Baseline.
- Gold, Silver, or Bronze Standard designations can also be made for service snapshots.

To designate service snapshots as the Gold, Silver, or Bronze Standard

1. Click the Management link.

   The Management panel appears.

2. Click the Service tab.

   The Service tab page appears.

3. Click the name of the service you want to view or edit snapshot details about in the Service table.

   The selected service's Service Details page appears and displays the Service tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the service snapshots you want to designate as the Gold, Silver, or Bronze Standard, then select one of the following options from the Select Actions drop-down list:

   - Set as Gold Standard
   - Set as Silver Standard
   - Set as Bronze Standard

   The Designation column in the Snapshots table displays Gold Standard, Silver Standard, or Bronze Standard for the corresponding snapshots. These snapshots are used when the Gold, Silver, or Bronze Standard snapshot option is selected as the source or target snapshot as described in Compare Services or Components.

Remove the Baseline Designation from a Service Snapshot

You can remove a Baseline designation from a service snapshot when you no longer want to use it as the reference for service Change Detection Across Time operations.

To remove a Baseline designation from a service snapshot

1. Click the Management link.

   The Management panel appears.

2. Click the Service tab.

   The Service tab page appears.

3. Click the name of the service whose snapshot is designated as the Baseline.

   The selected service's Service Details page appears and displays the Service tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the service snapshot whose Designation column displays Baseline, then select Remove Baseline from the Select Actions drop-down menu.

   You are prompted to confirm the removal.

6. Click OK to confirm the removal.

   The Baseline designation is removed from the server snapshot.

Remove the Gold, Silver, or Bronze Standard Designation from a Service Snapshot

You can remove a Gold, Silver, or Bronze Standard designation from a service snapshot if you no longer want it to be the reference for service comparison operations.

To remove a Gold, Silver, or Bronze Standard designation

1. Click the Management link.

   The Management panel appears.

2. Click the Service tab.

   The Service tab page appears.

3. Click the name of the service whose snapshot is designated as the Gold, Silver, or Bronze Standard.

   The selected service's Service Details page appears and displays the Service tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the service snapshot whose Designation column displays Gold, Silver, or Bronze Standard, then select one of the following options from the Select Actions drop-down menu:

   - Remove Gold Standard
   - Remove Silver Standard
   - Remove Bronze Standard

   You are prompted to confirm the removal.

6. Click OK to confirm the removal.

   The designation is removed from the service snapshot.

Delete Service Snapshots

You can delete server snapshots that you no longer need.

To delete a server snapshot

1. Click the Management link.

   The Management panel appears.

2. Click the Service tab.

   The Service tab page appears.

3. Click the name of the service whose snapshot you want delete in the Service table.

   The selected service's Service Details page appears and displays the Service tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the service snapshots you want to delete, select Delete Snapshots from the Select Actions drop-down list, and then click OK to confirm the deletion.

   Note: You cannot delete a Gold, Silver, or Bronze Standard Standard or Baseline snapshot. You must remove the designation before deleting these snapshots.

   The selected service snapshots are deleted.

Chapter 5: Server Management

This section contains the following topics:

Add Servers Manually

Automatically discovering or importing server data files are the more common methods of adding servers. You can also manually add servers for CA Configuration Automation to manage.

Follow these steps:

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

3. Select Create Server from the Table Actions drop-down list on the Servers tab.

4. Complete the following fields on the Create Server page:

   Name

   Specifies the name or IP address of the server.

   Network Discovery Gateway

   Specifies the NDG Server that is used for discovery.

   Perform Network Discovery

   Specifies whether the server performs a network discovery.

   Default: Yes (selected)

   Note: Enabling the network discovery displays the alias names that are associated with an IP address in the Hostname Aliases field of the Server Details page. Network discovery reconciles the servers that have multiple DNS aliases and displays the accurate count of the discovered servers. The product does not discover the aliases that are associated with the server as separate servers.

5. Click Next.

   The Details page appears and displays the server name or IP address you entered on the previous page.

6. Complete the following fields:

   IPv4 Address

   Defines the server IPv4 address.

   IPv6 Address

   Defines the server IPv6 address.

   IP Locked

   Defines a user-assigned IP address for the server that is different from the one that the Reconcile IP operation discovers or resolves.

OS Name

Defines the server operating system.

State

Specifies the server state:

- **New:** The product inventories the server, but assigns no management operations to it.
- **Managed:** The product runs Discovery and Management Profiles at scheduled times for servers in this state.

Access Profile

Defines the access profile that accesses the server.

Management Profile

Defines the management profile that manages software components on the server.

Note: To assign the profile that the CA Configuration Automation administrator designated as the default, select Use Default Profile. The product does not predefine default profiles; the CA Configuration Automation administrator must designate them.

Network Realm

Defines the realm in environments that have multiple, private networks. These private networks are independent of each other, which can cause conflicts when you try to discover and manage servers with duplicate IP addresses. To identify private networks uniquely, assign each a Network Realm string.

Note: You can customize the Server table on the Servers tab to display the Network Realm column (as described in Filter Table Views (see page 33)). The column displays the name of the Network Realm that is associated with each server. To modify the realm, click the link in the Server Name column to display the Server Details page, and then select a Network Realm or enter a new name.

Business Owner

Defines a user or organizational entity that is responsible for using the server.

Business Process

Defines the tasks or jobs the server is responsible for performing.

IT Owner

Defines the IT entity. Depending on the organization of the IT department, the owner can be a person or group responsible for managing and maintaining the server.

Location

Defines the physical location of the server (any combination of country, state, city, building, floor, room number, and so on).

Notes

Lists any other information about the server that is useful for managing servers.

7. Click Next.
8. Double-click one or more services to add to the server from the Available Services column.

   The selected service is moved to the Selected Services column.

   You can also add or remove services as follows:

   - To move a service from either column to the opposite column, click the service and then click the left or right single arrow.
   - To move all the services from either column to the opposite column, click the double left or right arrows.
9. Click Next.
10. On the Server Groups page, double-click one or more server groups to add to the service from the Available Server Groups column.

    The selected server group moves to the Selected Server Groups column.

    You can also add or remove server groups from the service as follows:

    - To move a server group from either column to the opposite column, click the group and then click the left or right single arrow.
    - To move all the server groups from either column to the opposite column, click the double left or right arrows.
11. Click Finish.

    The product creates the service and adds it to the Services table.
12. (Optional) Select Test Servers from the Select Actions drop-down list to test server communications. Ensure that the product can resolve or find the specified IP address or server name. For more information, see Test Servers (see page 120).

Test Servers

You can test communications between CA Configuration Automation and one or more servers using the Test Servers operation. The operation attempts to contact the selected servers and the CA Configuration Automation Agent if installed on the target server.

To test server communication

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more servers you want to test, then select Server Actions, Test Servers from the Select Actions drop-down list.

   The Test Servers dialog appears and displays the following results in the Server Test Results table:

   Servers

   Specifies the names of the server that is selected for the test.

   Access Mode

   Specifies the mode that is used to access the server (WMI, SSH, and so on).

   Ping Test

   Displays a check mark if the server responded to being pinged or an X if there was no response.

   Server Access Test

   Displays a check mark if the server is accessed by the test or an X if it could not be accessed.

   If the test server is successful, the Is Accessible column value in Server table is Yes. Else, the value is No.

Group Servers

You can organize your managed servers into logical groupings to simplify server-based management operations by selecting multoiple servers as a group instead of selecting them individually.

Note: This functionality is similar to the Create Server Groups (see page 150) option on the Server Groups page.

Follow these steps:

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to the servers you want to group, then select Server Actions, Group Servers from the Select Actions drop-down list.

   The Add Servers to Server Group dialog appears.

3. Select one of the following options, then click OK:

   New Server Group

   Creates a new server group using the name you enter.

   Existing Server Group

   Enables you to select an existing server group.

   The selected servers are added to the new or existing server group. The group appears in the table on the Server Groups page.

Create Server Snapshots

The Take Snapshots option creates a point-in-time copy of a server. You supply the snapshot name, and CA Configuration Automation time stamps the snapshot to identify the copy uniquely.

Snapshots can help you track and identify configuration changes made to your servers when you run Change Detection and Compare operations to compare the point-in-time copy to the current server data to see what changed.

To create a server snapshot

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more servers whose state is Managed for which you want to create a snapshot, then select Management Actions, Take Snapshot from the Select Actions drop-down list.

   The Take Snapshot dialog appears.

3. Enter a name and description for the snapshot, then click OK.

   A snapshot is created for each server you selected. You can click the server name, then the Snapshot tab in the Server Details page to view and manage the new snapshot.

Run Server Change Detection

The Change Detection operation determines how a server has changed over a period of time. Change Detection uses snapshots (point-in-time copies) of server data to provide a detailed account of all detected configuration changes, as well as file system changes, including file ownership, file permission, and file modification times. Server-based Change Detection provides options for finding differences between any two of the following options:

- Current server data
- Most recent snapshot
- Second most recent snapshot
- Most recent snapshot on a specific date
- Selected snapshot
- Baseline snapshot
- Gold standard snapshot
- Silver standard snapshot
- Bronze standard snapshot

Component elements that have the Time Variant filter set in the corresponding Component Blueprint are the only items not checked for changes by the Change Detection operation. Log files and data modified at runtime are examples of Time Variant-filtered elements.

Note: Running Change Detection using the procedure in this section is considered running the operation manually.

To run Change Detection to identify changes to servers

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the check box next one or more servers on which you want to search for changes, and then select Management Actions, Run Change Detection from the Select Actions drop-down list.

   The Differences Between page of the Run Change Detection wizard appears.

4. Select a Source Snapshot from the following options:

   Current Data

   Specifies the current system data available for the server is used as the source.

   Most recent snapshot

   Specifies the snapshot with the most recent timestamp is used as the source.

   Second most recent snapshot

   Specifies the snapshot with the second most recent timestamp is used as the source.

   Most recent snapshot on a specific date

   Specifies the date of the snapshot to use as the source. If you select this option, the Source Snapshot Date field appears. Select the date of the snapshot you want to use. If there are multiple snapshots available on the specified date, the most recent snapshot on that date is used.

   Selected snapshot

   Specifies a user-selected snapshot to use as the source. If you select this option, the Source Snapshot field appears. Select the snapshot you want to use.

   Baseline

   Specifies the snapshot designated as the Baseline is used as the source.

   Gold Standard

   Specifies the snapshot designated as the Gold Standard is used as the source.

   Silver Standard

   Specifies the snapshot designated as the Silver Standard is used as the source.

   Bronze Standard

   Specifies the snapshot designated as the Bronze Standard is used as the source.

5. Select a Target Snapshot (the options are the same as those listed in step 4), then click Next.

   The Component Blueprints page appears.

6. Select one of the following options:

   Include All Component Blueprints

   Specifies that software components of all component blueprints are searched for changes on the server.

   Select Component Blueprints by Name

   Specifies that software components of one or more component blueprint is searched for changes on the server. If you select this option you must select the blueprints in the Available Blueprints column, then click the single right-facing arrow to move it to the Select Blueprints column.

7. Click Next.

   The Filters page appears.

8. Select one of the following options to specify what differences are included in the Change Detection results:

   No children comparison if a hierarchical object exists on source or target only

   Specifies whether or not the child objects are compared for changes if the object only exists on either the source or target server.

   All Differences

   Specifies that all differences on the servers are included.

Component Inventory Differences Only

Specifies that only the differences in the the component inventory are included.

Filters

Specifies that one or more of the following objects are included:

- Folders—Accept the default setting (All) or click the Select option, then select one or more folders to search for changes.
- Categories—Accept the default setting (All) or click the Select option, then select one or more categories to search for changes. Categories are assigned in the Component Blueprint and are the organizational groupings to which an element belongs.
- Weights—Accept the default setting (All) or click the Select option, then select the weights to search for changes. Weights are assigned in the Component Blueprint and represent the relative importance of an element. Unweighted elements (no weight assigned) are considered Medium.

You can press Ctrl+click to select multiple non-consecutive list entries, or Shift+click to select multiple consecutive list entries.

9. Click Finish.

   The Change Detection operation runs, and the results appear as described in Viewing the Results of Change Detection and Compare Operations (see page 63).

Compare Servers or Components

The Compare Servers operation determines differences between any two of the following options on two or more different servers:

- Current server data
- Most recent snapshot
- Second most recent snapshot
- Most recent snapshot on a specific date
- Selected snapshot
- Baseline snapshot
- Gold standard snapshot

- Silver standard snapshot
- Bronze standard snapshot

Note: Component elements that have the Time Variant filter set in the corresponding Component Blueprint are the only items not checked for differences. Log files and data modified at runtime are examples of Time Variant-filtered elements.

To compare servers or components

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to the servers you want to compare, and then select Management Actions, Run Compare from the Select Actions drop-down list.

   The Run Compare wizard appears with the first server you selected listed in the Source Server field and the other servers listed in the Target Server fields.

3. Select a Source Snapshot from the following options:

   Current Data

   Specifies the current system data available for the server is used as the source.

   Most recent snapshot

   Specifies the snapshot with the most recent timestamp is used as the source.

   Second most recent snapshot

   Specifies the snapshot with the second most recent timestamp is used as the source.

   Most recent snapshot on a specific date

   Specifies the date of the snapshot to use as the source. If you select this option, the Source Snapshot Date field appears. Select the date of the snapshot you want to use. If there are multiple snapshots available on the specified date, the most recent snapshot on that date is used.

   Selected snapshot

   Specifies the snapshot to use as the source. If you select this option, the Source Snapshot field appears. Select the snapshot you want to use.

   Baseline

   Specifies the snapshot designated as the Baseline is used as the source.

   Gold standard

   Specifies the snapshot designated as the Gold Standard is used as the source.

Silver standard

Specifies the snapshot designated as the Silver Standard is used as the source.

Bronze standard

Specifies the snapshot designated as the Bronze Standard is used as the source.

4. Select a Target Snapshot (the options are the same as those listed in step 4).
5. Click Next.

   The Components page appears.
6. Select one of the following Component Blueprint options:

   Include All Component Blueprints

   Specifies that all software components of all component blueprints are compared on the selected servers. If you select this option, skip to step 10.

   Select Component Blueprints by Name

   Specifies that software components of one or more component blueprints are compared on the selected servers. If you select this option you must select the blueprints in the Available Blueprints column, then click the single right-facing arrow to move it to the Select Blueprints column.

   After you select one or more component blueprints, the ability to select software components is activated. Continue with step 8.
7. Click the Enable Component Selection check box if you want to select the specific software components that are compared on the selected servers.

   The Source Component field is activated.
8. Select the source and target software components that you want to compare from the Source Comparison drop-down list and by moving target components to the Selected Components column.
9. Click Next.

   The Filters page appears.
10. Select one of the following options to specify what differences are included in the comparison results:

    No children comparison if a hierarchical object exists on source or target only

    Specifies whether or not the child objects are compared for changes if the component only exists on either the source or target server.

    All Differences

    Specifies that all differences on the servers are compared.

Component Inventory Differences Only

Specifies that only the differences in the component inventory are compared.

Filters

Specifies that one or more of the following objects are compared:

- Folders—Accept the default setting (All) or click the Select option, then select one or more folders to compare.
- Categories—Accept the default setting (All) or click the Select option, then select one or more categories to search for changes. Categories are assigned in the Component Blueprint and are the organizational groupings to which an element belongs.
- Weights—Accept the default setting (All) or click the Select option, then select the weights to search for changes. Weights are assigned in the Component Blueprint and represent the relative importance of an element. Unweighted elements (no weight assigned) are considered Medium.

You can press Ctrl+click to select multiple non-consecutive list entries, or Shift+click to select multiple consecutive list entries.

11. Click OK.

The comparison results appear as described in Viewing the Results of Change Detection and Compare Operations (see page 63).

Run Server Rule Compliance

Server-based Rule Compliance lets you check server and snapshot data against the following:

- Default value rules
- Data type rules
- Constraint rules defined in Component Blueprints

Note: Running Rule Compliance using the procedures in this section is considered running the operation manually. To schedule Rule Compliance operations, see Create Rule Compliance Jobs.

To run server-based Rule Compliance

1. Click the Management link, then Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more managed servers on which you want run Rule Compliance, then select Management Actions, Run Rule Compliance from the Select Actions drop-down list.

   The Server Criteria page of the Run Rule Compliance wizard appears.

3. Select the lowest severity level of messages you want Rule Compliance to report from the Severity drop-down list.

   The available Rule Category options (described in step 5) are determined by the severity level you select in this step.

   Information

   Displays Information, Warning, Error, and Critical messages. Enables you to select either Default Value Rules or Data Type Rules in step 5.

   Warning

   Displays Warning, Error, and Critical messages. Disables Data Value Rules in step 5.

   Error

   Displays Error and Critical messages. Disables Data Value Rules in step 5.

   Critical

   Displays only Critical messages. Disables both Default Value Rules and Data Type Rules in step 5, and requires that you define Explicit Rules as described in step 7.

4. Select one of the following options from the Remediation drop-down list to specify whether you want to use remediation to reset non-compliant values:

   None

   Specifies that non-compliant values appear in the Rule Compliance results, but does not use remediation to reset these values.

   Rule Value Only

   Specifies that remediation is used to reset non-compliant values to the values that are defined in the rules.

   Rule Value or Blueprint Default Value

   Specifies that remediation is used to reset non-compliant values to the values that are defined in the rules. If an explicit rule is not defined for the component, the non-compliant value is reset to the default value defined in the Component Blueprint.

5. Create a set of rules against which to run Rule Compliance by selecting one of the following categories in the Rule Category area (if you selected Critical in the Severity drop-down list, these options are not available):

   Default Value Rules

   Verifies current server or snapshot values against specified default values.

   When default values are specified in a Component Blueprint, CA Configuration Automation automatically creates rules that check to see if the actual value deviates from the default value. Default rule deviations show as Information messages in the results.

   Data Type Rules

   Verifies current server or snapshot values against specified values for the corresponding data type.

6. Accept the default Explicit Rules settings or click the Select Rule Groups option and double-click the rule group you want to use in the Available Rule Groups column to move it to the Selected Rule Groups column.

   Blueprint Rules

   Verifies current service or snapshot values against constraint rules defined in Component Blueprints.

   Includes both user-defined rules and built-in rules, such as data type checking.

   Instance Rules

   Verifies current service or snapshot values against constraint rules defined in the service and Component Blueprints.

7. Select the snapshot you want to use to establish compliance.

   Note: If you selected multiple servers in step 2, and want to use snapshot data, select each server's most recent snapshot on specified date option and enter a date in the Snapshot Date field. If there is no snapshot for the selected servers on the specified date, CA Configuration Automation displays an error message on the results page. If there are multiple snapshots available on the specified date, the most recent snapshot on that date is used.

8. Click Next.

   The Components page appear.

9. Select one of the following Component Blueprint options:

   Include All Component Blueprints

   Specifies that all component blueprints are used on the selected servers.

   Select Component Blueprints by Name

   Specifies that one or more component blueprints are used on the selected servers. If you select this option you must select the blueprints in the Available Blueprints column, then click the single right-facing arrow to move it to the Select Blueprints column.

10. Click Next.

    The Filters page appears.

11. Accept the default setting (All) or create a filter to determine which of the following items are considered by the Rule Compliance operation:

    Folders

    Specifies whether all folders or only the selected folders are searched by the Rule Compliance operation.

    Categories

    Specifies whether all categories or only the selected categories are searched by the Rule Compliance operation. Categories are assigned in the Component Blueprint and are the organizational groupings to which an element belongs.

    Weights

    Specifies whether all weights or only the selected weights are searched by the Rule Compliance operation. Weights are assigned in the Component Blueprint and represent the relative importance of an element. Unweighted elements (no weight assigned) are considered Medium.

    You can press Ctrl+click to select multiple non-consecutive list entries, or Shift+click to select multiple consecutive list entries.

12. Click OK.

    The Rule Compliance results appear as described in Viewing the Results of Rule Compliance Operations (see page 71).

Refresh Servers

The refresh servers operation obtains the most current server component data.

Note: The basic state of a component does not change as a result of a refresh operation—Inventoried components remain Inventoried components and Managed components remain Managed components. The state of a component changes only as a result of the Management Profile directing components to be managed or as a result of service-related discovery of a component.

To refresh server software components

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to the servers whose components you want to refresh, and then click the Refresh button (above the Servers table).

   The Current Activity column of the Servers table displays Refresh during the operation. The column is blank when the server refresh operation is complete.

Run Server Discovery

The Server Discovery operation searches your networks for new software components and updates the existing components on the server.

Note: The procedure described in this section is considered a manual operation for times that you want to run a discovery immediately. Management Profiles are typically used to automate Server Discovery using the scheduling information provided in the profile. See Management Profiles (see page 79) for more information.

To manually run a Server Discovery

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the check box next to the servers you want to discover or update, and then select Management Actions, Run Discovery from the Select Actions drop-down list.

   Note: The corresponding Management Profile must be in the Enabled state.

   The Current Activity column of the Servers table displays Discovery during the operation.

Report Actions for Servers and Server Groups

The Reports Actions operation lets you run reports from Servers and Server Groups. You can perform the following actions to run reports:

Run Compliance Report

Run the compliance reports for the selected servers and server groups. By default, the corresponding Compliance Rule groups filter criteria is applied when you run the report.

Run Report

Run a report from the available reports. Based on the selected report, provide the necessary inputs to generate the report.

Follow these steps:

1. Click the Management link.
2. Click the Servers tab in the Management panel.
3. Select servers in the Servers page, and then select Reports Actions from the Select Actions drop-down list.
4. Select any one of the following actions:
   - Run Compliance Report—Run compliance report for the selected server and server groups.
   - Run Report—Run reports from the available reports.
5. If you select Run Compliance Report, select the Report Name, and click Ok.
6. If you select Run Report, select one of the following options in the Select Report panel:
   - Templates—Run a report using any of the predefined report templates.
   - Saved Reports—Run the reports that are saved earlier.

   Report is generated.

Stop Server Discovery

You can manually stop a server discovery operation if it is taking too long, or for any other reason.

To stop a server discovery operation

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the check box next to the server whose discovery operation you want to stop, and then select Management Actions, Stop Discovery from the Select Actions drop-down list.

   The Current Activity column of the Servers table displays Discovery while the operation is running, and is blank when the operation stops.

Run Management Profiles on Servers

You can run a Management Profile manually to update the configuration settings of software components on a server instead of waiting for the scheduled time specified in the profile. You have the following options for manually running management profiles:

- Run Management Profile with Discovery
- Run Management Profile without Discovery

The procedures associated with these options are described in the sections that follow.

Note: You can also use the Run Discovery option to run a server discovery without using the management options in a management profile. If you run a discovery using Run Discovery option, the server is searched for new software components based on the blueprints defined in the management profile. If you use management options, you get the benefits of managing snapshots, Run Change detection, Run Rule compliance in addition after the successful completion of discovery.

- For information about running a server discovery without using the management options in a Management Profile, see Run Server Discovery (see page 133).
- For information about defining Management Profiles, see Create Management Profiles (see page 80).

Run Management Profiles with Discovery

You can manually run a management profile with discovery instead of waiting for the scheduled time specified in the profile. When you run the profile with the discovery operation, the configuration settings of software components on a server are updated based on the discovery. If any new components are discovered in the process, they are added to the server.

To run the management profile with discovery manually

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one of more managed servers you want to update as defined by the management profile, and then select Management Actions, Run Management Profile with Discovery from the Select Actions drop-down list.

   Note: The Management Profile that corresponds with the server must be enabled.

   The Current Activity column of the Servers table displays Management Profile. The column is blank when the operation is complete.

Run Management Profile without Discovery

You can manually run a management profile without a discovery operation when you want to update the previously discovered servers and associated software components without searching for the new components on the servers.

To manually run the management profile without discovery

1. Click the Management link, then the Servers tab.
2. Click the check box next to one or more servers you want to update as defined by the management profile, and then select Management Actions, Run Management Profile without Discovery from the Select Actions drop-down list.

   Note: The Management Profile that corresponds with the server must be enabled.

   The Current Activity column of the Servers table displays Management Profile. The column is blank when the operation is complete.

Assign Profiles to Servers

You can assign existing management and access profiles to servers to control how software components and servers are discovered, managed, and accessed.

To assign a profile to a server

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to the servers to which you want to assign a profile, and then select Server Actions, Assign Profile from the Select Actions drop-down list.

   The Assign Profiles dialog appears and displays a Management Profile and Access Profile selection areas.

3. Click Change, select a profile from the drop-down list, and then click OK for one or both profile types.

   The corresponding profile column of the Servers table displays the name of the assigned profile.

Reconcile Server IPs

The Reconcile Servers functionality reconciles a server whose IP address conflicts with records in the CA Configuration Automation Database. The Reconcile Servers operation updates the CA Configuration Automation Database with all servers' correct IPv4 or IPv6 address and ensures that all IP addresses are unique.

To reconcile servers' IP addresses in the CA Configuration Automation Database

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more servers whose IP address you want to reconcile, then select Server Actions, Reconcile Servers from the Select Actions drop-down list.

   A progress indicator displays as the operation retrieves all server entries in the CA Configuration Automation Database, resolves the IPv4 and IPv6 addresses based on host name, and then updates the database.

Manage Conflicting IP Addresses

NDG discovers all the configuration items available in a network and there is a possibility of IP addresses conflict. For example, a server is disabled, and its IP address is assigned to another server in the network. The disabled server's IP address instance remains in the discovered servers list. A new discovery lists old *and* new servers with same IP addresses in the discovered servers.

CA Configuration Automation lets you manage the conflicting IP addresses and reconcile the duplicates. Reconciliation lets you eliminate duplicates and maintain consistency of data. The Reconcile Servers IPs operation updates the CCA Database with all servers' correct IPv4 or IPv6 address.

Follow these steps:

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click Table Actions tab, and select Manage Conflicting IP Addresses from the drop-down list.

   The Manage Conflicting IP Addresses page appears. The page displays the conflicting ipv4 and ipv6 addresses with the corresponding server names and MAC Addresses. By default, the conflicting servers IP address are sort by IPV4 Address.

3. Click Server Name to view the server details.

   Use Filter tab to filter the IP addresses based on specific IP4 address, IPV6 address, or subnet mask for IPV4 and IPV6 addresses.

4. Select one or more servers whose IP address you want to reconcile, and then select the Reconcile Servers IPs from the Select Actions drop-down list.

   A new NDG discovery id done to resolve the old conflict IP address.

5. Select Delete Servers from the Select Actions drop-down list if you want to Delete redundant servers.

   A message confirms the reconciliation or deletion of the selected servers. If error occurs, view the logs in the server log tables to know more about the error.

Secure Agents

CA Configuration Automation provides the capability for servers to communicate with CA Configuration Automation Agents using an SSL secured connection. The Secure Agents option lets you create the SSL certificates required for CA Configuration Automation Agent security.

Secured communications requires a certificate for identification on both the server side and the agent side. You must create a certificate authority before you can secure Agent communications. See Creating and Managing Security Certificates for more information.

Note: You cannot use Agent Security with an SSH proxy, and you must use the Manually Configured Agent selection for Agent Mode in the Access Profile.

When you switch from Manual Agent to Secured Agent, the access profile associated with the selected servers is modified as follows:

- If you are using a predefined access profile (that is, one of the following profiles installed by the CA Configuration Automation installation program: Manual Agent, Port Probe, Secured Agent,Self Registered, SSH, WMI, or WMI - SSH), a new access profile is created. This new profile is:
  - Automatically assigned a name using the Manual Agent<*timestamp*> convention (for example, Manual Agent<*timestamp1*>, Manual Agent<*timestamp2*>, and so on)
  - Defined as a Secure Agent (that is, the Secure Agent checkbox is checked in the Access Mode tab of the Edit Server Access Profile dialog box)
  - Listed in the table on the Access Profiles tab
  - Assigned to the selected servers
- If you are using a user-defined access profile that is used by multiple servers, a new access profile is created. This new profile is:
  - Renamed using the <*user-assigned_name*><*timestamp*> convention (for example, if you are using a profile you created called TestProfile, the new profile would use that name appended with a timestamp: TestProfile<*timestamp1*>, TestProfile<*timestamp2*>, and so on)
  - Defined as a Secure Agent (that is, the Secure Agent check box is checked in the Access Mode tab of the Edit Server Access Profile dialog box)
  - Listed in the table on the Access Profiles tab
  - Assigned to the selected servers
- If you are using a user-defined access profile that is used by *only* the selected server, the access mode of the current access profile is modified to Secure Agent (that is, the Secure Agent check box is checked). A new access profile is *not* created.

To secure CA Configuration Automation Agents on one or more servers

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to the servers with agents that you want to secure (the Agent Installed column of the Servers table contains a check mark for servers that have agents installed), then select Agent Actions, Secure Agents from the Select Actions drop-down list.

   The Secure Agents dialog appears.

3. Enter the appropriate information in the following fields, then click OK:

   Agent Certificate Password

   Specifies the password for the certificate.

   Enter the password text again. .

   Confirm Password

   Confirms the password entered in the Agent Certificate Password field. The two passwords must be identical.

   Certificate Authority Password

   Specifies the certificate authority password required to create the agent certificate.

   CA Configuration Automation creates the certificate for the agent, installs this new certificate in the agent installation directory, configures the agent to only accept secure connections, and restarts the agent with the new configuration. After CA Configuration Automation has successfully completed these steps, the Access Mode column of the Server table displays Secure Agent.

To revert to an unsecure agent

1. Do one of the following to edit the Access Profile for the server whose agent you want to set as unsecure:

   - Clear the Secure Agent check box on the Access Mode tab, then click OK.
   - Assign a new profile that does not use Secure Agent mode if do not want to change the existing access profile (in the case where other servers may be using it).

2. Go to CA Configuration Automation Agent installation directory and edit agent.conf file manually by changing the secure option to 0 (zero), and then restart the agent.

   The agent runs in unsecure mode.

Install CCA Agents Remotely

CA Configuration Automation uses the access and configuration details from the assigned Access Profile to install CCA Agents remotely on managed servers.

To install a CCA Agent on a server, the product requires information about the target server operating system, the IP address, and the listening port. You can obtain these details from the Server Details page. The remote agent installation program automatically detects whether the target server uses a 32- or 64-bit operating system. The installation program then installs the corresponding agent version.

Note: The access profile user account must be valid on both the Grid Node host and the remote server where you install the agent.

The user types that an access profile defines must also meet the following requirements:

Local user

A local user ID in the access profile must exist on the Grid Node host and the remote server on which you install the agent. Both user IDs must have the same password, which must also exist in the access profile.

Domain user

When the access profile defines a domain user account and the Grid Node host or the remote server on which you install the agent belongs to a workgroup, then the computer with the password specified in the access profile requires a local user account. Therefore, when the computer belongs to a workgroup, the product uses a local account even though the access profile defines a domain.

You cannot install CA Configuration Automation Agents remotely in the following circumstances:

- When the target server resides behind a firewall.
- When the root directory on the target server does not have enough space to copy the installation program.
- When the administrative shares are removed from a target Windows server.
- When a discovery operation assigns an operating system of "Linux or UNIX" to a target Linux server. Manually change the operating system to Linux before you install the CA Configuration Automation Agent.

- When you want to install from a Linux or UNIX CA Configuration Automation Server to a Windows 2003 or 2008 server.
- When the target UNIX server is not running Secure Shell. To determine whether the server is running Secure Shell, run the following command:

  ```
  # ps -ef | grep sshd
  ```

  The target server allows root access through Secure Shell (SSH/FTP) and enable SFTP in the SSH daemon configuration file. To ensure that the target server allows the root user Secure Shell access, run the following command from the root account of a different computer:

  ```
  # ssh <target_installation_server_name>
  ```

  Note: The root user must have Secure Shell login permissions.

Follow these steps:

1. Click the Management link, and then click the Servers tab.
2. On the Servers tab, select the check box for each server on which to install an agent.
3. From the Select Actions drop-down list, select Agent Actions, Install Agents.

   The Current Activity column displays "Agent Installation in Progress." When the installation finishes, the product clears the Current Activity column.

   Notes:

   - Ensure that the install directory specified in the access profile is a valid Windows or UNIX path. You can also use a System Environment variable as discussed in the Use Environment Variables (see page 143) section to define the installation directory.
   - If you install on a former CA Configuration Automation managed agentless server, rediscover the services that reference it to get the available components and component folders. For more information, see Refresh Services (see page 72).

Use Environment Variables in Install Directory

Consider the following standards for using the environment variables in Install Directory:

- Use the following pattern:

  $<VARIABLENAME>$

  Where the <VARIABLENAME> is an environment variable that is defined on a target system.

Examples:

  - Refer the %ProgramFiles% environment variable on a target machine as $ProgramFiles$ in the access profile.
  - Specify the C:\Program Files\CA\CCAAgent installation path as $ProgramFiles$\CA\CCAAgent in the access profile.

- Ensure that a specific environment variable is defined on the target machine for the CCA agent installation
- The $ character is not supported in the agent installation path or the environment variable. If $ character is used, follow the $<VARIABLENAME>$ pattern.
- Restart the CCA Grid node, or the CA Configuration Automation Server if the environment variables do not get resolved during the remote agent installation.

Uninstall CCA Agents

CA Configuration Automation provides support for uninstalling CA Configuration Automation Agents from the servers with existing Agents that you select in the Server Management display list. CA Configuration Automation uses the access and configuration details contained in the assigned Access Profile to remotely uninstall CA Configuration Automation Agents.

To remotely uninstall CA Configuration Automation Agents from managed servers

1. Click the Management link, then the Servers tab.

   The Servers tab page appears. Servers that have CA Configuration Automation Agents installed display Yes in the Agent Installed column.

2. Select Server Actions, Test Servers from the Select Actions drop-down list to test server communications.

   Test server communications to ensure the servers you selected can be resolved or found.

3. Click the check box next to one or more servers from which you want to uninstall an agent, then select Agent Actions, Uninstall Agent from the Select Actions drop-down list.

   The Current Activity column displays Agent Uninstallation in Progress. When the uninstall is complete, the Agent Installed column displays No.

Locate Agents and SSH Access

CA Configuration Automation enables you to identify the servers that have the CA Configuration Automation Agents that are installed or are configured to allow SSH access. After you identify the servers, set their state to Managed or Rejected, and assign a Management and Access Profiles to them.

Note: Enable the TCP Echo port (7) of the servers to locate the Agent or SSH.

To identify servers with CA Configuration Automation Agents or SSH access

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more servers on which you want to search for an agent,search for SSH access, or obtain agent details, then select Agent Actions, Locate Agents and SSH from the Select Actions drop-down list..

   The Auto-Locate dialog appears with the Locate CA Configuration Automation Agents option and Port 8063 selected by default.

3. (Optional) Do one or more of the following:
   - Clear the Locate CA Configuration Automation Agents check box if you do not want to locate servers with CA Configuration Automation Agents installed.
   - Change the CCA Ports (from 8063) to locate CA Configuration Automation Agents listening on a non-default port. You can specify multiple ports by entering comma-separated port numbers, or a range of ports separated by a hyphen (-). For example: 8000-8050, 8063.
   - Click the Locate SSH Remote Access check box if you want to locate CCA-managed servers with SSH access configured. Selecting this option activates the Verify SSH Remote Access check box.
   - Change the SSH Ports (from 22) to locate CCA-managed servers with SSH configured, but listening on a non-default port.
   - Click the Verify SSH Remote Access check box, and filter the managed servers with SSH configured to only those using a specific user ID, password, or both.
4. Click Auto-Locate.

   The selected servers are searched and the results appear on the Auto-Locate Agents/SSH Results tab.

Manage Servers Found by Auto-Locate

The servers identified by the Locate Agents or SSH operation (described in the previous (see page 144) section) are displayed in the Server table that includes Agent Found and SSH Found columns.

- A green check is displayed if an Agent or SSH port was found.
- A red X is displayed if an Agent or SSH port not was found.

To manage the servers identified by Auto-Locate

1. Click the check box next to the servers you want to manage or reject, then click the corresponding button (Manage Selected Servers or Reject Selected Servers).
2. (Optional) Assign an existing Management Profile or Access Profile to the servers located:
   a. Click one or more servers to which you want to assign a profile, then click the Assign Profiles button .

      The Assign Profile dialog box is displayed with the Do Not Change option selected for each profile type.

   b. Click Change next to the profile types you want to assign to the selected servers.

c. Specify the profile you want to assign for profile types with the Change option selected.

d. Click OK.

The profile is assigned to the selected servers.

Set Network Discovery Gateway

When you install CA Configuration Automation, you must specify the Network Discovery Gateway (NDG) server that CA Configuration Automation uses to perform discovery operations. If you install another NDG Server, and prefer to use that NDG server for discovery operations, you can use the Set Network Discovery Gateway option to assign the new NDG Server to your CA Configuration Automation Server.

NDG Servers are also associated with the servers that they discover. The Server table contains a Network Discovery Gateway column that shows the NDG Server that discovered it.

To assign an NDG server to a CA Configuration Automation Server

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more servers that you want to associate with an NDG server, then select NDG Actions, Set Network Discovery Gateway from the Select Actions drop-down list.

   The Set Network Discovery Gateway dialog appears.

3. Select the NDG server from the Network Discovery Gateway drop-down list, then click OK.

   The Network Discovery Gateway column of the Server table is update with the new NDG Server.

View Relationships and Management Operation Results in the Visualization UI

You can display a graphical representation of relationships and management operations (for example, Change Detection or Rule Compliance) in the Visualization UI.

Follow these steps:

1. In the CA Configuration Automation Server UI, click the Management panel, then click one of the following locations:
   - Services tab
   - Servers tab
   - Server Groups tab
   - Storage tab
2. On the tab page that appears, select one or more check boxes in the table.

   For example, if you selected the Services tab, select one or more services in the Services table.
3. From the Select Actions drop-down list, select Visualization, then select one of the following visualization templates:

   Services tab:
   - Service Communication Relationships
   - Service Change Detection
   - Service Rule Compliance

   Servers tab and the Server Groups tab:
   - All Server Relationships
   - Server Change Detection
   - Server Communication Relationships
   - Server Configuration Relationships
   - Server Rule Compliance
   - Server Storage Relationships
   - Server Virtualization

   Storage tab:
   - Server Storage Management Relationships
   - Server Storage Relationships

   For a description of each template, see Select a Visualization Template (see page 379).

Delete Servers

You can permanently delete servers if you no longer want to manage them in CA Configuration Automation.

To delete one or more servers

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more servers you want to delete, then select Server Actions, Delete Servers from the Select Actions drop-down list.

   A message prompts you to confirm the deletion.

3. Click OK.

   The specified servers are removed from the Server table.

Export Servers

You can export a server definition to use in another instance of CA Configuration Automation.

Note: Only the server properties are exported; the associated components and snapshots are not exported.

To export a server

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click the check box next to one or more servers you want to export, then select Server Actions, Export Servers from the Select Actions drop-down list.

   The File Download dialog appears.

3. Click Save.

   The Save As dialog appears and the export file is assigned the default name Servers.csv.

4. Edit the file name if desired, select the location to save the file, and then click Save.

   The server definition is exported to the selected location.

Import Servers

You can import a server definition as a CSV file exported from another instance of CA Configuration Automation.

To import a server definition

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click Table Actions, then select Import Servers.

   The Import Servers dialog appears.

3. Enter or select the following information in the corresponding field, then click OK:

   CSV File

   Specifies the name of the CSV file that contains the server definitions you want to import. You can click Browse to navigate to the file.

   Network Discovery Gateway

   Specifies the NDG Server that is assigned to all the servers being imported. This NDG Server is used to reconcile the IP addresses of the imported servers.

   Network Realm

   Defines the realm in environments that have multiple, private networks. These private networks are independent of each other, which can cause conflicts when you try to discover and manage servers with duplicate IP addresses. To identify private networks uniquely, assign each a Network Realm string.

   Note: You can customize the Server table on the Servers tab to display the Network Realm column (as described in Filter Table Views (see page 33)). The column displays the name of the Network Realm that is associated with each server. To modify the realm, click the link in the Server Name column to display the Server Details page, and then select a Network Realm or enter a new name.

   Perform Network Discovery

   Specifies whether a server import starts a Network Discovery to discover or confirm details about the server including the server name, reconcile the IP address, and classify the device (OS Family, OS Name).

   Allow Unknown Hosts

   Specifies whether servers defined in the JAR file that do not respond to a ping command can be imported and managed by CA Configuration Automation.

   The file is imported and the servers appear in the Servers table.

Create Server Groups

You can organize your managed servers into logical groupings either manually or automatically using the Server Groups option. Server groups can simplify server-based management operations by letting you select servers as a group instead of selecting them individually.

To create a server group

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the Server Groups link.

   The Server Groups page appears.

3. Select Create Group from the Table Actions drop-down list.

   The Create Server Group page appears.

4. Enter the following information:

   Name

   Specifies the name for the server group.

   Description

   Describes the intended purpose or function of the server group.

5. Click Next.

   The Servers page appears.

   - Continue with step 6 to create a dynamic server group.
   - Skip to step 7 to manually add servers to the server group.

6. Create a filter to automatically add managed servers of a certain type to this server group when they are discovered or imported:

   a. Click the Dynamic Server Group check box.

      The Filter fields appear.

   b. Select an option from the Column drop-down list. For example, select OS Name.

   c. Select an option from the Value drop-down list or enter a value. For example, type Windows 2008 Server.

   d. (Optional) Click And or Or to further refine the filter by selecting another option in the second pair of Column and Value fields. For example, select the Column option of Server State and assign the value of Imported to create a filter the adds all imported Windows 2008 Servers to this Server Group.

   Skip to step 8.

7. Double-click one or more servers that you want to add to the server group in the Available Servers column.

   The servers are moved to the Selected Servers column.

8. Click Save.

   The server group is created and appears in the Server Groups table.

Edit Server Groups

You can edit details about an existing server group.

To edit server groups

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the Server Groups link.

   The Server Group page appears.

3. Click the name of the server group you want to edit in the Server Groups table.

   The server group's Details page appears.

4. Edit the following information on the Server Group tab as appropriate:

   Name

   Specifies the name for the server group. If you change the name, you are actually creating a new group and deleting the existing one.

   Description

   Describes the intended purpose or function of the server group.

5. Double-click one or more server groups that you want to add or remove from the server group.

   The selected server group is moved to the opposite column.

   Alternatively, you can add or remove server groups from the service as follows:

   - Click a server group in either column and click the left- or right-facing single arrow to move it to the opposite column.
   - Click the double left- or right-facing arrows to move all the server groups to the opposite column.

6. (Optional) Click the Log tab to view the log of all server group activity.

7. Click Save.

   The server group is updated with your changes.

Delete Server Groups

You can delete one or more existing Server Groups if you no longer need them.

To delete Server Groups

1. Click the Management link, then the Server tab.

   The Server tab page appears.

2. Click the Server Groups link.

   The Server Group page appears.

3. Click the check box next to the Server Groups you want to delete, then select Delete Server Groups from the Select Actions drop-down list.

   The selected Server Groups are deleted.

Export Server Groups

You can export a server group definition to use in another instance of CA Configuration Automation.

Note: Only the server group properties are exported; the associated components and snapshots are not exported.

To export a server group

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click the Server Groups link.

   The Server Groups page appears.

3. Click the check box next to one or more server groups you want to export, then click Select Actions, and select Export Server Groups.

   The File Download dialog appears.

4. Click Save.

   The Save As dialog appears and the export file is assigned the default name ServerGroup_Export_*<timestamp>*.jar

5. Edit the file name if desired, select the location to save the file, and then click Save.

   The server group definition is exported to the selected location.

Import Server Groups

You can import a server group definition as a JAR file exported from another instance of CA Configuration Automation.

To import a server group definition

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click the Server Groups link.

   The Server Groups page appears.

3. Click Table Actions, then select Import Server Groups.

   The Import Server Groups dialog appears.

4. Enter or select the following information in the corresponding field, then click OK:

   JAR File to Import

   Specifies the name of the file that contains the server group definitions you want to import. You can click Browse to navigate to the file.

   Overwrite Existing Access Profiles

   Specifies whether the file being imported overwrites a file with the same name. Select this option if you want the changes made to the profile on another instance of CA Configuration Automation to be retained.

5. Click one of the following buttons:

   Import All

   Imports all of the Server Groups in the JAR file.

   Import On Selected

   Displays a dialog where you can select the Server Groups in the JAR file to import.

   The file is imported and the profiles appear in the Server Groups table.

View Cluster Details

All Softagent-based Network Discovery operations discover server clusters without you having to specify any cluster-specific discovery options. If Softagent access is successful for any of the servers that participate in a cluster, NDG successfully discovers all servers in the cluster and all cluster details. You can view these details and relationships on the Clusters tab or in the Visualization tool.

The following cluster technologies can be discovered and managed:

- Microsoft Cluster Server (MSCS)—Failover and load balancing
- IBM PowerHA—High availability
- Red Hat Cluster Suite (RHCS)—A component of the high availability and load balancing add-ons for Red Hat Enterprise Linux (RHEL)

Note:

- Cluster relationships do not appear on the Relationships tab.
- Blueprint Discovery of server clusters do not appear on the Clusters tab.

Follow these steps:

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the Clusters link.

   The clusters discovered by CA Configuration Automation appear in the Clusters pane, and the details of the first cluster appear in the right pane.

3. Do one of the following tasks in the Clusters pane:

   - Click the name of the cluster whose details you want to view.

     The details of the selected cluster appear in the right pane.

   - Click the plus sign (+) next to the cluster whose servers you want to view.

     The cluster node expands and the servers in the cluster appear.

4. (Optional) Click a server name.

   The details of the selected server appear in the right pane.

Access Profiles

Access Profiles provide the rules for server access and CA Configuration Automation Agent installation, including:

- The methodology the CA Configuration Automation Server uses to gain access to and collect information about the servers and components in your enterprise
- The security and proxy information required to successfully install CA Configuration Automation Agents on selected servers, and to securely communicate with the agents

Note: Access Profiles are not required for basic Agent and agentless discovery operations. Their primary purpose is to allow you to install Agents and set up SSH and WMI access on multiple machines. Secondarily, they provide an easy way for you to make changes to or to switch between access modes.

You create new profiles or edit existing profiles using the Access Profiles link on the Servers tab of the Management page. You make the actual profile assignments on the Server tab page using the Assign Profiles option on the Select Actions drop-down list.

CA Configuration Automation includes the following predefined Access Profiles to get you started with server access operations:

- Manual Agent
- Port Probe
- SSH
- Secured Agent
- Self Registered
- WMI
- WMI - SSH

These profiles have sensible defaults for gaining communication access to servers within your enterprise using the various modes available. However, because Access Profiles are unique to your enterprise, you can create custom access profiles, or copy and modify the predefined profiles.

The Access Profiles page displays the table of known profiles in alphabetic order by name, and also displays the description, the creator and creation timestamp, modification details, access mode (Agent, SSH, Telnet, WMI, WMI and SSH), proxy type, and whether or not the profile is assigned to a server.

Create Access Profiles

In addition to using the predefined profiles, you can create your own server-specific Access Profiles.

Follow these steps:

1. Click the Management link, and then the Servers tab.
2. On the Servers tab, click the Access Profiles link.
3. From the Table Actions drop-down list, select Create Access Profile.
4. Enter a profile name and description and then click Next to open the Access Mode page.

5. From the Access Mode drop-down list, select one of the following modes, and then click Next:

Port Probe

Accesses other servers in your enterprise in Port Probe Access Mode. Select this option if you do not intend to install a CA Configuration Automation Agent on the servers that are associated with this profile.

Agent

Accesses other servers in your enterprise in Agent Access Mode. Select this option if you have installed or intend to install a CA Configuration Automation Agent on the servers that are associated with this profile. If you select this option, complete the following Agent and Agent Installation fields:

Agent Mode

Specifies how the server where the agent is installed is configured with the CA Configuration Automation Server:

Self Registered Agent: Enables an CA Configuration Automation Agent to register itself when it tests communications with the CA Configuration Automation Server. After the agent self-registers, the product considers the host a managed server. When you select this option, the agent obtains the agent listening port.

Manually Configured Agent: Enables you to specify the CA Configuration Automation Agent listening port number. Select this option on a server that uses a Pass-through Agent for communication with the CA Configuration Automation Server. You can also select this option if your CA Configuration Automation Server is configured to communicate securely using SSL. Define the Pass-through Agent proxy type in Step 6.

Agent Port

Defines the CA Configuration Automation Agent listening port number.

Default: 8063

Secure Agent

Specifies whether the CA Configuration Automation Server communicates with the agent using an SSL-secured connection. Create a certificate authority before you secure a CA Configuration Automation Agent.

Install Directory

Identifies the location where the CA Configuration Automation Agent software is installed.

Default (Windows):

```
\Program Files\CA\CA Configuration Automation Agent
```

Default (UNIX):

```
/opt/CA/CCAAgent
```

Install JVM

Specifies whether to install the Java Virtual Machine (JVM) distributed with CA Configuration Automation. The CA Configuration Automation Agent installation requires a JVM on the target server:

Yes: Install the CA Technologies-supplied JVM.

No: Use a previously installed JVM.

Note: To locate the existing JVM, provide a commonly known installation location. If the CA Configuration Automation Agent installation fails, select Yes in this field and try installing the agent again.

System Account

Defines the user ID of the administrative user with authorization and privileges to connect to and install the CA Configuration Automation agent.

System Password

Defines the password that is associated with the specified System Account.

Retype Password

Verifies that the password matches the string that you entered in the System Password field.

Enable Use of sudo

Specifies whether you can use the sudo command to access and gather information about the remote UNIX and Linux servers. The sudo command enables the users that are defined in the /etc/sudoers configuration file to run commands. The sudo command lets users run commands as if they were users with different (in the case of the root user, unlimited) permissions.

If you enable sudo, comment out the Default requiretty entry in the /etc/sudoers file as follows:

```
# Default requiretty
```

For more information, see Configuring sudo for UNIX and Linux Softagent Discovery (see page 405).

Agent Logging

Specifies whether to enable agent logging.

Note: To conserve space and enhance security, some environments discourage writing log files to servers. You can also enable or disable the agent logging in the CA Configuration Automation Agent configuration file (agent.conf).

Server Ping

Specifies whether to enable the server ping. Clear the Server Ping check box to disable the server ping in the following instances:

• You encounter IP address and name resolution conflicts.

• The server has multiple Network Interface Cards (NICs). The product pings the server to ensure that it identifies with the intended NIC.

• Agents are installed on servers that have a firewall between that server and the CA Configuration Automation Server.

You can also enable or disable the server ping in the CA Configuration Automation Agent configuration file (agent.conf).

Note: The product requires you to enable the server ping to populate the CA Configuration Automation Agent-related details accurately on the attribute sheets and manage lists.

SSH

Accesses and collects data from associated servers using the Secure Shell (SSH). SSH provides authentication and secure encrypted communications over insecure networks. If you select this option, complete the following SSH fields:

SSH Mode

Specifies whether SSH with Credentials or SSH with Key File is used to access and retrieve data from discovered servers.

Port

Defines the SSH communications port.

Default: 22

Account

Defines the SSH login account.

Enable Use of sudo

Specifies whether you can use the sudo command to access and gather information about the remote UNIX and Linux servers. The sudo command enables the users that are defined in the /etc/sudoers configuration file to run commands. The sudo command lets users run commands as if they were users with different (in the case of the root user, unlimited) permissions.

If you enable sudo, comment out the Default requiretty entry in the /etc/sudoers file as follows:

```
# Default requiretty
```

For more information, see Configuring sudo for UNIX and Linux Softagent Discovery (see page 405).

Connection Timeout

Defines the interval (in milliseconds) before the product considers an SSH connection request to a remote server to have failed.

Default: 900000 (15 minutes)

Trust

Specifies whether the product performs file-based server verification and automatically verifies the remote servers. To increase security, clear the check box and provide the known hosts file name in the SSH Host File field.

Hosts File

Defines the file that the product uses to validate remote servers.

Default: <*home-directory*>/.ssh/known_hosts

Secure File Transfer Client

Specifies whether to use the Secure File Transfer Clients (SFTP) or Secure Copy (SCP) to perform SSH Discovery.

The product requires the SFTP during discovery using SSH and WMISSH access modes. The SFTP handles the following functions:

- Transfer a file to a remote server.
- Get a file from a remote server.
- Remove a file from a remote server.
- Run a script on a remote server.

If you select SFTP and the SFTP service is not running on the remote server, the CA Configuration Automation Server logs the following message:

`<message_number>: Discovery failed on Server "<server_name>"`

You can either start the SFTP service on the remote server or use SCP as the secure file transfer client.

Default: SFTP

Account Password

Defines the SSH password. The product displays this field only when you select the SSH with Credentials mode.

Retype Account Password

Confirms that the password matches the text string that is entered in the Account Password field. This field appears only when the SSH with Credentials mode is selected.

Private Key File

Defines the private key file. To create the public and private key files, use puttygen.exe or a similar utility. After you create the files, copy the private key to the CA Configuration Automation Server home directory and the public key to the SSH server. For example, on copSSH, copy the public key into \copSSH\home\Administrator\.ssh\authorized-keys.

The product displays this field only when you select the SSH with KeyFile mode.

Public Key File

Specifies the key format, either ssh-dss or ssh-rsa. The product displays this field only when you select the SSH with KeyFile mode.

Passphrase

(Optional) Defines a key file protection passphrase. Associate the passphrase with the key files when they are created. The product displays this field only when you select the SSH with KeyFile mode.

Telnet

Specifies whether the CA Configuration Automation Server uses the Telnet Access Mode to access other servers in your enterprise. If you select this option, complete the following Telnet fields:

Port

Defines the Telnet listening port.

Default: 23

Account

Defines a valid user account on the remote server.

Account Password

Defines the password that is associated with the specified account.

Retype Account Password

Verifies that the password matches the string that you entered in the Account Password field.

Enable Use of sudo

Specifies whether you can use the sudo command to access and gather information about the remote UNIX and Linux servers. The sudo command enables the users that are defined in the /etc/sudoers configuration file to run commands. The sudo command lets users run commands as if they were users with different (in the case of the root user, unlimited) permissions.

If you enable sudo, comment out the Default requiretty entry in the /etc/sudoers file as follows:

```
# Default requiretty
```

For more information, see Configuring sudo for UNIX and Linux Softagent Discovery (see page 405).

Connection Timeout

Defines the interval (in milliseconds) before the product considers a Telnet connection request to a remote server to have failed.

Default: 900000 (15 minutes)

Look For Prompts

Specifies whether the discovery looks for the Login Prompt, the Password Prompt, and the Shell Prompt values while it attempts to access the remote server. These fields contain the standard Telnet prompts by default, but they can vary on some systems.

Login Prompt

Defines the login prompt for which the access profile gains access. When the profile locates the specified prompt, it enters the information in the Account field of an access profile. You can edit this field if the target server uses a prompt other than the login prompt.

Password Prompt

Defines the password prompt for which the access profile gains access. When the profile locates the specified prompt, it enters the information in the Account Password field of an access profile. You can edit this field if the target server users a prompt other than password.

Shell Prompt

Defines the shell prompt for which the access profile looks to issue commands after gaining access. You can edit this field if the target server users a shell prompt other than #.

WMI

Specifies whether the CA Configuration Automation Server uses Microsoft Windows Management Instrumentation (WMI) Access Mode to access other servers in your enterprise to discover software components.

If your CA Configuration Automation Server is installed on a UNIX or Linux host, you cannot use a WMI Access Profile to access target Windows servers. To access target the Windows servers using a WMI Access Profile, at least one CA Configuration Automation Grid Server must be installed on a Windows server.

If you select this option, complete the following fields:

User

Defines a valid user account on the remote server. To use a WMI access profile to access the target servers, run the CCA Server service, and the CCA Grid service with domain credentials. You must have administrator privileges on the host and target servers to run the services.

Password

Defines the password that is associated with the specified User.

Confirm Password

Verifies that the password matches the string that you entered in the Password field.

WMI and SSH

Specifies whether CA Configuration Automation Server accesses other servers in your enterprise using a combination of WMI and SSH Access Modes. This combination enables the discovery to use methodologies that are optimized for discovering and accessing both Windows and Linux/UNIX servers.

If you select this option:

a. Specify whether to use SSH with Credentials or SSH with KeyFile

b. Complete the appropriate SSH and WMI fields.

The Proxy page appears.

6. Select a proxy type from the Proxy Type drop-down list. The proxy type specifies how the CA Configuration Automation Server communicates with the CA Configuration Automation Agent.

No Proxy

Specifies that the CA Configuration Automation Server communicates with the CA Configuration Automation Agent directly. No Proxy is the default setting.

Port Forwarding Proxy

Specifies that the CA Configuration Automation Server communicates with the CA Configuration Automation Agent through a firewall gateway. If you select this proxy type, complete the following fields:

Proxy Server

Defines the name or IP address of the proxy server.

Proxy Port

Defines the port on which the proxy server listens.

SSH

Specifies that the communication between the CA Configuration Automation Server and the CA Configuration Automation Agent is secure and encrypted. If you select this proxy type, complete the following fields:

SSH Server Host

Specifies the name or IP address of the SSH host computer. The drop-down list is populated with all the servers listed in the Server table on the Server tab.

SSH Server Port

Defines the SSH communications port.

Default: 22

Account

Defines the SSH login account.

Account Password

Defines the SSH account password.

Retype Account Password

Verifies that the password matches the string that you entered in the SSH Account Password field.

Pass-through Agent

Specifies that the CA Configuration Automation Server communicates with the CA Configuration Automation Agent through an intermediary agent. The pass-through agent consolidates communications from multiple CA Configuration Automation Agents and is a single point of communication with the CA Configuration Automation Server. The drop-down list is populated with all the servers listed in the Server table on the Server tab.

If you select Pass-through Agent, set the Agent Mode to Manually Configure Agent and then set the Access Mode page Agent Port field. Step 5 defines how to set the Agent Mode to Manually Configure Agent.

You can configure cascading pass-through agents, but only the first proxy-hop supports SSH.

The following illustration is an example of the cascading pass-through agents:

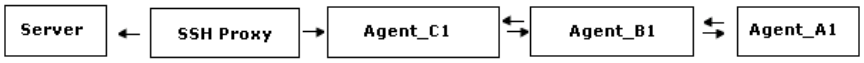

WMI

Specifies that the communication between the CA Configuration Automation Server and the target server is through a proxy server. If you select this proxy type, complete the following fields:

WMI Server Host Name

Defines the name of the proxy server.

Account

Defines a valid user account with administrator credentials to log in to the proxy server using WMI.

Password

Defines the administrator password that is associated with the Account value.

Retype Password

Verifies that the password matches the string that you entered in the Password field.

FTP Server

Defines the name of the proxy server. The product updates this value with the proxy server name that is specified in the WMI Server Host Name field.

Note: If you change the FTP Server value, the product updates the WMI Server Host Name value. The FTP connection must communicate from the CA Configuration Automation server to the FTP server, and from the target server to the FTP server. The FTP connection must use the short name or host name.

FTP Server Port

Defines the FTP server listening port.

FTP Root Directory

Defines the FTP root directory.

Account

Defines the name of the FTP server user account.

Password

Defines the password that is associated with the Account value.

7. Click Finish.

The product creates the profile and displays it in the Access Profile table.

Import Access Profiles

You can import an Access Profile as a Java Archive (JAR) file from another instance of CA Configuration Automation.

To import an Access Profile

1. Click the Management link, then either the Services or Servers tab.

   The Services or Servers tab page appears.

2. Click the Access Profiles link (below the main tabs).

   The Access Profiles page appears.

3. Click Table Actions, then select Import Access Profiles.

   The Import Access Profiles dialog appears.

4. Enter or select the following information in the corresponding field:

   JAR File to Import

   Specifies the name of the JAR file that contains the Access Profile you want to import. You can click Browse to navigate to the file.

   Overwrite Existing Access Profiles

   Specifies whether the file being imported overwrites a file with the same name. Select this option if you want the changes made to the profile on another instance of CA Configuration Automation to be retained.

5. Click one of the following buttons:

   Import All

   Imports all of the Access Profiles in the JAR file.

   Import On Selected

   Displays a dialog where you can select the Access Profiles in the JAR file to import.

   The file is imported and the profiles appear in the Access Profiles table.

Edit Access Profiles

You can edit any existing Access Profiles including the predefined profiles and profiles that are assigned to servers.

To edit an Access Profile

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the Access Profiles link (below the main tabs).

   The Access Profiles page appears and displays the existing profiles.

3. Click the name of the profile you want to edit.

   The Details page appears for the selected profile.

4. Edit the entries in any of the fields on any of the tabs, then click Save. For details about the fields, see Create Access Profiles (see page 155).

   The profile is updated.

Copy Access Profiles

You can copy any existing Access Profile.

To copy an Access Profile

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the Access Profiles link (below the main tabs).

   The Access Profiles page appears and displays the existing profiles.

3. Click the name of the profile you want to copy.

   The Details page appears for the selected profile.

4. Click copy.

   The profile is copied and listed in the Access Profiles table as Copy of *<selected_access_profile>*.

5. (Optional) Edit any of the fields on any of the tabs, then click Save. For details about the fields, see Create Access Profiles (see page 155).

Delete Access Profiles

You can delete any existing Access Profile you no longer want assign to servers.

To delete an Access Profile

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the Access Profiles link (below the main tabs).

   The Access Profiles page appears and displays the existing profiles.

3. Click the check box next to the profile you want to delete, then select Delete Profile from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   A message confirms the deletion and the profile is removed from the Access Profiles table.

Export Access Profiles

You can export an Access Profile as a JAR file to use in another instance of CA Configuration Automation.

To export an Access Profile

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the Access Profiles link (below the main tabs).

   The Access Profiles page appears.

3. Click the check box next to the profile you want to export, then click Select Actions and select Export Access Profiles.

   The File Download dialog appears.

4. Click Save.

   The Save As dialog appears and the export JAR file is assigned a default name using the following format:

   AccessProfile_Export_*<year>*_*<month>*_*<date>*_*<hour>*_*<minutes>*_*<seconds>*.jar

   For example: AccessProfile_Export_2009_08_09_04_20_00.jar

5. Edit the file name if desired, select the location to save the file, and then click Save.

   The profile is exported to the selected location.

View the Server Log

CA Configuration Automation maintains a server log that documents each server-based transaction. It can easily be accessed using the UI for each server being managed by CA Configuration Automation.

To view a server's activity log

1. Click the Management link, then the Servers tab.

   The Servers tab page appears.

2. Click the name of the server whose activity log you to want to view in the Server table.

   The selected server's Server Details page appears and displays the Server tab.

3. Click the Log tab.

   The Log table appears and displays all activity for the selected server. The default view is to show the most recent server activity in the first row, and the oldest activity in the last row. You can click the sort icons at the top of sortable columns to present the table data differently.

4. (Optional) Create a filter to view specific server events captured in the log as described in Filter Table Views.

   Note: The Filter functionality on this page contains the following additional fields only found on Log pages:

   Start Time

   Specifies the time to begin searching for server activity. For example, if you want to see all server activity during the week you were away on vacation, you would specify the day you left as the start time by clicking the the Start Time check box, clicking the calendar icon, and selecting the day and time you left.

   End Time

   Specifies the time to end searching for server activity. Continuing the previous example, if you want to see all server activity during the week you were away on vacation, you would specify the day you returned as the end time by clicking the the End Time check box, clicking the calendar icon, and selecting day and time you returned.

View and Edit Server Details

You can view or edit details about a server.

To edit or view server details

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Do one of the following:

   - Click the name of the server you want to edit or view details about in the Server table.

     The selected server's Server Details page appears in the same browser and displays the Server tab.

   - Ctrl+Click in the Server Name field of the server you want to edit or view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.

4. Edit the information on the Server tab as appropriate, then click Save.

   For a description of the fields, see Add Servers (see page 118).

   The Server Detail page and the Server table are updated with your changes.

In addition to the details contained on the Server Details page, you can view server details on the following pages by clicking the corresponding link on the Server tab:

- Network Adapters
- Hardware
- Applications
- Services/Daemons
- Open Ports

You can view and edit the following server-specific objects from the other tabs on the server's Details page:

- Components (displays the components that are discovered by the server's Management Profile)
- All Components (displays all components that are discovered on the server)

  Note: These two tabs contain a Missing column that specifies whether a component is missing when rediscovered.

- Relationships
- Snapshots
- Services
- Server Groups
- Agent
- Log

These pages and tabs are described in the sections that follow.

View Server Virtualization Details

You can view relationships between elements in a virtual environment.

To view server virtualization details

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Do one of the following:

   - Click the name of the server you want to view details about in the Server table.

     The selected server's Server Details page appears in the same browser and displays the Server tab.

   - Ctrl+Click in the Server Name field of the server you want to view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.

4. Click the Relationships tab, then the Virtualization link.

   The Virtualization table appears and displays details about the selected server's virtual environment relationships.

View Network Adapter Details

You can view details about a server's Network Adapters.

To edit or view network adapter details

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Do one of the following:

   - Click the name of the server you want to edit or view details about in the Server table.

     The selected server's Server Details page appears in the same browser and displays the Server tab.

   - Ctrl+Click in the Server Name field of the server you want to edit or view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.

4. Click the Network Adapter link.

   The Network Adapters table appears and displays details about the selected server.

5. (Optional) Click a link in the IP Address column.

   The network configuration details appear for the selected network adapter.

View Hardware Details

You can view details about the hardware on which a server runs.

To view a server's hardware details

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Do one of the following:

   - Click the name of the server you want to edit or view details about in the Server table.

     The selected server's Server Details page appears in the same browser and displays the Server tab.

   - Ctrl+Click in the Server Name field of the server you want to edit or view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.

3. Click the Hardware link.

   The Hardware table appears and displays details about the hardware on which the selected server runs.

   Note: Due to limitations of the information available for WMI access to virtual machines running Windows operating systems, the Processor Logical Count field shows the number of virtual processors. It does not consider the number of cores per processor or hyper-threading capability of the processor.

View Application Details

You can view details about the applications installed on a server.

To view a server's application details

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Do one of the following:
   - Click the name of the server you want to edit or view details about in the Server table.

     The selected server's Server Details page appears in the same browser and displays the Server tab.
   - Ctrl+Click in the Server Name field of the server you want to edit or view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.
4. Click the Applications link.

   The Applications table appears and displays details about the applications installed on the selected server.

View Service and Daemon Details

You can view details about the services and daemons associated with a server.

To view service and daemon details

1. Click the Management link.

   The Management panel appears.
2. Click the Servers tab.

   The Servers tab page appears.
3. Do one of the following:
   - Click the name of the server you want to edit or view details about in the Server table.

     The selected server's Server Details page appears in the same browser and displays the Server tab.
   - Ctrl+Click in the Server Name field of the server you want to edit or view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.
4. Click the Services/Daemons link.

   The details about the services and daemons associated with the selected server appear.

View Open Port Details

You can view the ports that are not being used servers that are being managed by CA Configuration Automation.

To view servers' open ports

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Do one of the following:

   - Click the name of the server you want to edit or view details about in the Server table.

     The selected server's Server Details page appears in the same browser and displays the Server tab.

   - Ctrl+Click in the Server Name field of the server you want to edit or view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.

4. Click the Open Ports link.

   The Opens Ports table appears and displays the details about the unassigned ports on the selected server.

View Relationship Details

You can view virtual, static, and dynamic relationships between a selected server and the other servers with which it communicates. When you display details about a server, the Relationships table lists all the types of communications for the selected server. The selected server can be listed in the Server Name 1 or Server Name 2 column depending on the types of communication listed in the Communication Type column.

The communication type is configured on the Communication Mapping page as described in View and Edit Communication Mappings.

To display a server's relationship details

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Do one of the following:

   - Click the name of the server you want to edit or view details about in the Server table.

     The selected server's Details for Server page appears in the same browser and displays the Server tab.

   - Ctrl+Click in the Server Name field of the server you want to edit or view details about.

     The selected server's Server Details page appears in a new browser and displays the Server tab.

3. Click the Relationships tab, then click one of the following links:

   - Virtualization

     The Virtual Environment Relationships table appears.

   - Communication

     The Communication Relationships table appears.

   - Configuration

     The Configuration Relationships table appears.

     Configuration relationships resolve both the server name and IP address.

     - If the target server is in the list of CCA managed servers, then the target server information is populated with IPv4 and IPV6 addresses.
     - If the target server is not in the list of CCA managed servers, then the target server information is resolved based on "reverse lookup" and populated with IPv4 and IPV6 addresses.
     - If the target server is not in the list of CCA managed servers and not resolvable using "reverse lookup" then IP information remains empty.

4. Click one of the links in the following columns to display details about the specific relationship:

   - Virtual Environment (Virtual Environment Relationships table)
   - Communication Type (Communication Relationships table)
   - Parameter Name (Configuration Relationships table)

   You can also do the following:

   - Click the Relationships link to return to the Relationships table.
   - Click the Servers link to return to the Servers table.

Managing Server Components

CA Configuration Automation enables you to manage software components from a Components table that is available on the each server's Server Details page. The table shows only the components that are installed on the selected server. A similar view is available for service-centric management of software components from the Service Details page.

You can perform the following component management operations from the Server Details, Components page:

- View Components by Server (see page 177)
- Refresh Components by Server (see page 178)
- Delete Components from Servers

View Components by Server

The Components and All Components pages display a table of the software components for the selected server. From these pages, you can navigate to a tree view of configuration attributes and settings for the components.

Follow these steps:

1. Click the Management link, and then click the Servers tab.
2. Click a server in the Server Name column.

   The Servers Details page opens for the selected server.
3. Click one of the following tabs:

   Components

   Displays the components that the Management Profile for the selected server discovers.

   All Components

   Displays the components that the product has discovered on the selected server.

4. Complete one of the following actions:

   - Click the component link in the Name column.

     The component detail page appears with the Component tab displayed in the right pane.

   - Click the check box next to one or more components, and then select View Components from the Select Actions drop-down list.

     The View Components and Configurations window opens in a new browser page. The window displays the Components pane on the left, and the Component tab on the right.

   In both views, the left pane displays a tree that uses the following icons to identify elements in the tree:

    Server

    Blueprints

    Primary folders

    Subfolders and component folders

    Parameters

   Note: The Management Profile defines the component for the server.

5. To see the folders and configuration elements under each server or component, click any plus sign (+) to expand the tree.

   The details of the selected element appear in the right pane.

Refresh Components by Server

The Refresh Components option obtains the most current server component data.

Note: The basic state of a component does not change as a result of a refresh operation—an Inventoried component remains an Inventoried component and a Managed component remains a Managed component. The state of a component changes only as a result of the Management Profile directing components to be managed or as a result of server-related discovery of a component.

To refresh components on a server

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the name of the server on which you want refresh a component in the Server table.

   The selected server's Server Details page appears and displays the Server tab.

4. Click the Components tab.

   The Components table appears.

5. Click the check box next to the components you want to refresh, then select the Refresh Components option from the Select Actions drop-down list.

   The Current Activity column in the Components table displays Component Refresh in Progress, during the refresh operation. The Refresh Date/Time column displays the current date when the operation is complete.

Delete Components from Servers

You can delete server components that you no longer want to mange in CA Configuration Automation.

To delete one or more components from a server

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the name of the server from which you want delete a component in the Server table.

   The selected server's Server Details page appears and displays the Server tab.

4. Click the Components tab.

   The Components table appears.

5. Click the check box next to the components you want to delete, select Delete Components from the Select Actions drop-down list, and then click OK to confirm the deletion.

   Note: You cannot delete a component that is part of a service. You must delete the component from the service before deleting it from the server.

   The selected components are deleted.and removed from the Components table.

Managing Server Snapshots

A server snapshot is a point-in-time copy of a server being managed by CA Configuration Automation. The Snapshot tab of the Server Details page displays a list of all existing server snapshots. By default the list sorts the snapshots chronologically with the oldest snapshot at the bottom and the newest snapshot at the top.

You can perform the following snapshot management operations from Select Actions or Table Actions drop-down list on the Server Details, Snapshots tab page:

- View Server Snapshots (see page 180)
- Run Change Detection (see page 123)
- Set a Server Snapshot as the Baseline (see page 181)
- Set Server Snapshots as the Gold, Silver, or Bronze Standard (see page 182)
- Remove a Baseline Designation (see page 184)
- Remove a Gold, Silver, or Bronze Standard Designation (see page 184)
- Delete Server Snapshots (see page 185)
- Export Server Snapshots (see page 186)
- Import Server Snapshots (see page 186)

View Server Snapshots

You can view any server snapshot that you have created.

To view a server snapshot

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the name of the server you want to view snapshot details about in the Server table.

   The selected server's Server Details page appears and displays the Server tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the server snapshots you want to view, then select View Snapshot/Components.

   The View Components and Configurations page is displayed. For information about this page, see View Components and Configurations.

You can also create a filter to search server snapshots.

1. Select a filter option from the Column drop-down-list to search server snapshots in the Filters section. The filter options are as follows:
   - Blueprint Name
   - Created By
   - Designation
   - Snapshot Description
   - Snapshot Name
   - Snapshot Origin
2. Select an option from the Value drop-down list or enter a value.
3. Click And or Or to refine the filter by selecting another option in the second pair of Column and Value fields.
4. Enable Start Date/Time or the End Date/Time check box to filter by date and time.
5. Click Go.

   The server snapshots are listed based on the filter criteria.

Set a Server Snapshot as the Baseline

CA Configuration Automation enables you to designate one server snapshot as the *Baseline*—the snapshot that is used as the reference for server Change Detection Across Time operations.

Notes:

- A similar snapshot feature enables you to designate multiple server snapshots as the Gold, Silver, or Bronze Standard. These snapshots are used as the reference for server comparison operations.
- The same server snapshot can be designated as the Baseline and as the Gold, Silver, or Bronze Standard.
- Baseline designations can also be made for service snapshots.

To designate a server snapshot as the Baseline

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the name of the server from which you want delete a snapshot in the Server table.

   The selected server's Server Details page appears and displays the Server tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the server snapshot you want to designate as the Baseline, then select Set as Baseline from the Select Actions drop-down list.

   The Designation column displays Baseline for the corresponding snapshot. This snapshot is used when the Baseline snapshot option is selected as the source or target snapshot as described in Run Server Change Detection (see page 123).

Set Server Snapshots as the Gold, Silver, or Bronze Standard

CA Configuration Automation enables you to designate server snapshots as the Gold, Silver, or Bronze Standard. These snapshots are used as the reference for server comparison operations.

The gold, silver, and bronze designations do *not* necessarily imply a hierarchy where the Gold Standard is a higher standard than the Silver Standard, and the Silver Standard is a higher standard than the Bronze Standard. While they can be implemented this way, they are designed to give you three different standards to apply to the server snapshots used as the reference for comparison operations.

To increase the flexibility, you can designate more than one server snapshot as the Gold, Silver, or Bronze Standard in the case where you want a specific component of the snapshot to be used for the comparison. For example, you could have a Gold Standard for the operating system, another Gold Standard for the database, a Bronze Standard for Linux mail servers, and another Bronze Standard for Windows mail servers.

Notes:

- Because snapshots are differentiated by timestamp and not by name, you can create multiple snapshots with same name and designate them as the Gold, Silver, or Bronze Standard.
- If you designate multiple snapshots as the Gold, Silver, or Bronze Standard, be careful when naming and creating the description of the snapshots to reflect how they are used.

- A similar snapshot feature enables you to designate one server snapshot as the *Baseline*—the snapshot that is used as the reference for Change Detection operations. There can be only one Baseline snapshot per server.
- The same server snapshot can be designated as the Gold, Silver, or Bronze Standard and as the Baseline.
- Gold, Silver, or Bronze Standard designations can also be made for service snapshots.

To designate server snapshots as the Gold, Silver, or Bronze Standard

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the name of the server you want to view or edit snapshot details about in the Server table.

   The selected server's Server Details page appears and displays the Server tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the server snapshots you want to designate as the Gold, Silver, or Bronze Standard, then select one of the following options from the Select Actions drop-down list:

   - Set as Gold Standard
   - Set as Silver Standard
   - Set as Bronze Standard

   The Designation column in the Snapshots table displays Gold Standard, Silver Standard, or Bronze Standard for the corresponding snapshots. These snapshots are used when the Gold, Silver, or Bronze Standard snapshot option is selected as the source or target snapshot as described in Compare Servers or Components.

Remove the Baseline Designation from a Server Snapshot

You can remove a Baseline designation from a server snapshot when you no longer want to use it as the reference for server Change Detection Across Time operations.

To remove a Baseline designation from a server snapshot

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the name of the server whose snapshot is designated as the Baseline.

   The selected server's Server Details page appears and displays the Server tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the server snapshot whose Designation column displays Baseline, then select Remove Baseline from the Select Actions drop-down menu.

   You are prompted to confirm the removal.

6. Click OK to confirm the removal.

   The Baseline designation is removed from the server snapshot.

Remove the Gold, Silver, or Bronze Designation from a Server Snapshot

You can remove a Gold, Silver, or Bronze Standard designation from a server snapshot if you no longer want it to be the reference for server comparison operations.

To remove a Gold, Silver, or Bronze Standard designation

1. Click the Management link.

   The Management panel appears.

2. Click the Servers tab.

   The Servers tab page appears.

3. Click the name of the server whose snapshot is designated as the Gold, Silver, or Bronze Standard.

   The selected server's Server Details page appears and displays the Server tab.

4. Click the Snapshots tab.

   The Snapshots table appears.

5. Click the check box next to the server snapshot whose Designation column displays Gold, Silver, or Bronze Standard, then select one of the following options from the Select Actions drop-down menu:
   - Remove Gold Standard
   - Remove Silver Standard
   - Remove Bronze Standard

   You are prompted to confirm the removal.
6. Click OK to confirm the removal.

   The designation is removed from the server snapshot.

Delete Server Snapshots

You can delete server snapshots that you no longer need.

To delete a server snapshot

1. Click the Management link.

   The Management panel appears.
2. Click the Servers tab.

   The Servers tab page appears.
3. Click the name of the server from which you want delete a snapshot in the Server table.

   The selected server's Server Details page appears and displays the Server tab.
4. Click the Snapshots tab.

   The Snapshots table appears.
5. Click the check box next to the server snapshots you want to delete, select Delete Snapshots from the Select Actions drop-down list, and then click OK to confirm the deletion.

   Note: You cannot delete a Gold, Silver, or Bronze Standard or Baseline snapshot. You must remove the designation before deleting these snapshots.

   The selected server snapshots are deleted.

Export Server Snapshots

You can export a server snapshots to use in another instance of CA Configuration Automation or to archive outside of the CA Configuration Automation Database.

To export a server snapshot

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click the name of the server (in the Server Name column) whose snapshot you want to export.

   The Server Details page appears.

3. Click the Snapshots tab.

   The Snapshots tab page appears.

4. Click the check box next to one or more server snapshots you want to export, click the Select Actions drop-down list, and then select Export Snapshots.

   The File Download dialog appears.

5. Click Save.

   The Save As dialog appears and the export file is assigned the default name ServerSnapshot_Export_<*timestamp*>.jar

6. Edit the file name if desired, select the location to save the file, and then click Save.

   The server snapshot definition is exported to the selected location.

Import Server Snapshots

You can import a server snapshot from a JAR file exported from another instance of CA Configuration Automation.

To import a server snapshot

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.

2. Click the name of the server (in the Server Name column) into which you want to import a snapshot.

   The Server Details page appears.

3. Click the Snapshots tab.

   The Snapshots tab page appears.

4. Click the Table Actions drop-down list, then select Import Snapshots.

   The Import Snapshots dialog appears.

5. Enter the following information in the corresponding field, then click OK:

   JAR File to Import

   Specifies the name of the JAR file that contains the server snapshot you want to import. You can click Browse to navigate to the file.

6. Click one of the following buttons:

   Import All

   Imports all of the snapshots in the JAR file.

   Import On Selected

   Displays a dialog where you can select the snapshots in the JAR file to import.

   The file is imported and the profiles appear in the Snapshots table.

Add Servers to Services

You can add a server to an existing service from the following locations:

- Server Details page
- Service Details page

From the Server Details Page

Follow these steps:

1. Click the Management link.
2. Click the Servers tab.
3. Click the name of the server to which to add a service in the Server table.
4. On the Server Details page, click the Services tab.

   The Available Services and Selected Services columns appear.

   Note: If you cannot select the listed services in the Available Services column, the server is not in the Managed state. Set the server state to Managed before you try to add it to a service.

5. Double-click one or more services in the Available Services column to which to add the selected server.

   The product moves the selected server to the Selected Services column.

   You can also add or remove services from the server as follows:

   - To move a service from either column to the opposite column, click the service and then click the left or right single arrow.
   - To move all the services from either column to the opposite column, click the double left or right arrows.

6. Click Save.

   The product adds the selected servers to the service.

From the Service Details Page

Follow these steps:

1. Click the Management link.
2. Click the Services tab.
3. Click the name of the service to which to add a server in the Services table.
4. On the Service Details page, click the Servers tab.

   The Available Servers and Selected Servers columns appear.

   Note: The Available Servers column only displays servers in the Managed state. Set the server state to Managed before you try to add a service.

5. Add or remove servers from the service as follows:

   - To move one or more servers to the Selected Servers pane, select them from the Available Servers column and then click the down arrow.
   - To move one or more servers from the Selected Servers pane to the Available Servers pane, click the up arrow.

6. Click Save.

   The product adds the selected servers to the service.

View All Services of a Server

The server can be part of a service either directly or indirectly through a Server Group. You can view all such services from Server Details page.

Follow these steps:

1. Click the Management link.
2. Click the Servers tab form the Management panel.
3. Click the name of the server you want to view to a service in the Server table.
4. Click the Services tab in the Server Details page.

   The Available Services and Selected Services appear. The Consolidated List of Services section displays all the services, and the corresponding server group name in the All Services table. The displayed services are related to a server either directly or indirectly through a server group. The Server Group can be either normal server group or a dynamic server group. If the Server Group column is empty, it means that the server is directly part of the corresponding service.
5. Add or remove one or more services for this server from the Available Services or Selected Services columns and click Save.

   All Services table is refreshed with the new services that are added or removed.

Add Servers to a Server Group

You can add a server to an existing server to an existing server group manually or automatically using the Dynamic Server Group option.

To add a server to an existing server group

1. Click the Management link, then click the Servers tab.

   The Servers tab page appears.
2. Click the Server Groups link.

   The Servers Groups page appears.
3. Click the name of the server group to which you want to add to a server in the Server Group table.

   The selected server's Server Group Details page appears and displays the Server tab.

4. Do one of the following:
   - Double-click one or more servers in the Available Servers column.

     The selected server is moved to the Selected Service column.
   - Click the Dynamic Server Group check box and create a filter to add managed servers of a specific type (for example Windows 2008 Servers) when they are discovered by, or imported into CA Configuration Automation.
5. Click Save.

   One of the following happens:
   - The selected server is either added to the server group.
   - The filter is created to add servers to the group dynamically when they are discovered or imported.

Chapter 6: Software Management

This section contains the following topics:

Software Overview

The Software tab lists the components and applications that the product discovered for all the servers in your enterprise. Use CA Configuration Automation to manage software components from a Components table and applications from the Applications table.

You can complete the following management operations from the Components page:

- View Components (see page 192)
- Delete Components (see page 193)

View Components

The Software tab Components option displays the components that CA Configuration Automation discovers on all servers in a table view or an expandable tree view.

Follow these steps:

1. Click the Management link, and then click the Software tab.

   A table lists the available components.

2. Complete one of the following actions:

   - Click the component link in the Name column.

     The component detail page appears with the Component tab displayed in the right pane.

   - Click the check box next to one or more components, and then select View Components from the Select Actions drop-down list.

     The View Components and Configurations window opens in a new browser page. The window displays the Components pane on the left, and the Component tab on the right.

   In both views, the left pane displays a tree that uses the following icons to identify elements in the tree:

    Server

    Blueprints

    Primary folders

    Subfolders and component folders

    Parameters

3. To see the folders and configuration elements under each server or component, click any plus sign (+) to expand the tree.

   The details of the selected element appear in the right pane.

Delete Components

You can delete server components that you no longer want to manage in CA Configuration Automation.

Follow these steps:

1. Click the Management link.
2. Click the Software tab.
3. Click the check box next to the components to delete.
4. Select Delete Components from the Select Actions drop-down list, and then click OK to confirm the deletion.

   Note: You cannot delete a component that is part of a service from the server. Delete the component from the service before you try to delete it from the server.

   The product deletes the selected components and removes them from the Components table.

View Applications

The Software tab Applications option displays the applications that CA Configuration Automation discovers using NDG discovery on all servers. The product displays all the available applications in a table view or an expandable tree view.

Follow these steps:

1. Click the Management link, and then click the Software tab.
2. Click the Applications link.

   The product displays the applications available on the servers. In the table and tree views, the left pane displays a tree that uses the following icons to identify elements in the tree:


Server


Application name


Subfolders and component folders

Chapter 7: Network Management

The Network tab provides access to network-based discovery operations for locating servers on the network segments in your enterprise.

You can perform the following network management operations from the Network tab page:

- Create Network Profiles
- Create Network Scan Policies
- Create Credential Vault Profiles
- Create Notification Profiles (see page 98)

These topics are described in the sections that follow.

Network Profiles

Network Management is used to configure Network Discovery to discover entities on your network and populate the CA Configuration Automation Database with them. Network Discovery can be configured to perform simple entity-level server discoveries that include Server Name, IP address, and operating system classification. Alternatively, Network Discovery can be used to gather more detailed information about the servers such as virtualization environment characteristics, installed applications, and server-to-server communication relationships.

These server details can then be used to decide whether a specific server should be managed by CA Configuration Automation.

Network Discovery uses Network Profiles to identify the following:

- Which Network Discovery Gateway (NDG) server performs the discovery
- Which subnets or computers to target (or exclude)
- The associated Network Scan Policy, which configures the discovery methodology to be used during the discovery process
- The associated Credential Vault, which defines a list of credentials to be used during the discovery process
- How frequently and at what time the network discovery is performed
- What kind of notifications (if any) are sent and to whom

The Network Profiles page displays a table containing all the Network Profiles that are available. Additionally, you can create and manage Network Profiles from this page.

Create Network Profiles

You can create Network Profiles to manage network discovery operations.

Follow these steps:

1. Click the Management link, then the Network tab.

   The Network tab page appears.

2. Click the Network Profile link.

   The Network Profiles page appears.

3. Select Create Network Profile from the Table Actions drop-down list.

   The Profiles page of the Network Profiles wizard appears.

4. Enter the following information in the corresponding field, then click Next:

   Name

   Defines the name of the Network Profile.

   Description

   Describes the function of the profile.

   Network Discovery Gateway

   Defines the computer that performs the Network Discovery. The drop-down list displays all computers where NDG is installed.

   Network Scan Policy

   Defines which Network Scan Policy to use. Network Scan Policies define which discovery engine and which methodology the network discovery uses. The Network Scan Policy drop-down list contains entries from the table on the Network Scan Policy tab page. You can view a description of each policy on the tab page. You can view a description of the predefined policies that CA Configuration Automation installs in View Network Scan Policies (see page 225).

   The main types of Network Scan Policies are:

   - ARP Cache Scan
   - DNS Scan
   - Cloud Service Scan
   - IPv6 Local Link Scan
   - Netflow Analysis
   - Packet Analysis

- Pingsweep Scan
- TCP Connect Scan

Each policy type also has scan options (for example, Pingsweep Scan with Softagent or Pingsweep Scan with Softagent, but without Server Relationships). The fields that are displayed in Step 5 vary depending on which Network Scan Policy you select.

Credential Vault

Defines which Credential Vault Profile to use for network access. The Credential Vault drop-down list contains the profiles that you created. The drop-down list also contains a Use Default option that assigns the default Credential Vault Profile for the Network Profile you are creating.

Scan Type

Specifies whether the scan processes the scan request with IPv4 or IPv6. For example, if you only select the IPv4 check box, the product uses IPv4 to process the scan. The Discovery operation that uses this profile discovers both IPv4 and IPv6 networks.
Default: Both

Network Realm

Defines the realm in environments that have multiple, private networks. These private networks are independent of each other, which can cause conflicts when you try to discover and manage servers with duplicate IP addresses. To identify private networks uniquely, assign each a Network Realm string.

Note: You can customize the Server table on the Servers tab to display the Network Realm column (as described in Filter Table Views (see page 33)). The column displays the name of the Network Realm that is associated with each server. To modify the realm, click the link in the Server Name column to display the Server Details page, and then select a Network Realm or enter a new name.

Set Server State as Managed

Specifies whether discovered servers are automatically set to the Managed state when they are added to the Servers table.

Access Profile

Defines which Access Profile to use to access the discovered servers. When you select an Access Profile, the Test Server for the CA Configuration Automation Agent field becomes active.

Test Server for CA Configuration Automation Agent

Specifies whether to verify that a CA Configuration Automation Agent is installed on the discovered server. An Access Profile is required for the test operation to be performed.

Management Profile

Defines which Management Profile is assigned to the discovered servers.

The Inclusions page opens unless you selected one of the IPv6 Local Link Scan policies, in which case the Schedule page opens. If the Schedule page opens, continue with Step 8.

5. Complete the following fields if they appear on the Inclusions page (not all options appear for all Network Scan Policies):

Target Host Names - Add New Host Name

Defines one or more target servers for the discovery. Enter the server name in the Add New Host Name field, then click the right arrow (>). The server appears in the Selected Host Names column.

Define at least one Target Host Name or Target IP Address for a Network Scan Policy that displays these fields.

Target Host Names - Add From File...

Defines the .csv file from which to import target servers.

Follow these steps:

a. Click Add From File... in the Target Home Names column.

b. On the Add Servers From File dialog, click Choose File, browse to the CSV file location, and then click Open.

c. Click the File Delimiter drop-down list and then select either Comma or Tab.

d. Click OK.

The product adds the servers that are listed in the file to the Selected Host Names field. Click the left arrow (<) to remove unwanted servers.

Target IP Addresses - Add New IP Address

Defines one or more target IP addresses for the discovery. Enter the IP address in the Add New IP Address field, then click the right-facing arrow. The server appears in the Selected IP Addresses column.

Target IP Addresses - Add From File...

Defines the .csv file from which to import the target IP addresses.

Note: The IP address import supports the use of an asterisk (*) as a wild card and subnet reference characters. For example:10.10.10.* or 10.10.10.0/24

Follow these steps:

a. Click Add From File... in the Target IP Addresses column.

b. On the Add IP Addresses From File dialog, click Choose File, browse to the CSV file location, and then click Open.

c. Click the File Delimiter drop-down list then select either Comma or Tab.

d. Click OK.

The product adds the IP addresses that are listed in the file to the Selected IP Addresses field. Click the left arrow (<) to remove unwanted IP Addresses.

Target TCP Ports - Add New TCP Ports

Defines the TCP ports that are monitored during a discovery. The product discovers only the network traffic that uses these ports.

Enter the port number in the Add New TCP Port field, then click the right-facing arrow. The port number appears in the Selected Ports column. During the TCP Connect scan, only the selected ports are probed during the discovery scan to detect the open ports.

Note: If you define the inclusion or exclusion ports for the TCP Connect scan, the default connection timeout value is 1 millisecond to detect the open ports.

For an NDG server, add the following registry parameter to configure the connection timeout to a higher value:

- (32-bit machines) HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Network Discovery Gateway\TcpConnectTimeout
- (64-bit Machine) HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\Network Discovery Gateway\TcpConnectTimeout

Note: To detect the open ports accurately, set the timeout value that is based on the ping response time between the NDG server and target machine.

IPv6 Subnet

Specifies an IPv6 subnet hierarchy by clicking the up and down arrows in the following fields:

Global Routing Prefix Length

Defines the number of bits that precede the subnetID in the IPv6 address. If the global routing prefix length is less than 64, define at least one subnet level in the remaining fields.

Default: 48

Bits Per Level

Associates the bits in the subnetID with subnet levels in an IPv6 subnet hierarchy.

Note: The bits are left justified. The first subnet level comprises the leftmost *n* bits of the subnetID. A corresponding filter in scan policies lets you use these bits to filter scan requests.

Lower Bound

Identifies the lower bound of the range of values to be included in the subnet hierarchy.

Upper Bound

Identifies the upper bound of the range of values to be included in the subnet hierarchy. This value cannot be larger than the maximum number that the number of bits that are reserved for this level can represent.

Add Level

Defines another IPv6 subnet level. Click Add Level and then enter the Bits Per Level, Lower Bound, and Upper Bound for this level. You can use the bits higher than the number specified in the Global Routing Prefix Length field to define the subnet. If it is set to the default (48), you can assign the next 16 bits to subnets.

Remove Level

Removes the first (top) level of subnet filtering.

Note: When NDG discovers relationships, the product collects a relationship if either server in the relationship is targeted using the specified inclusion criteria. The product creates a corresponding server entry in the CA Configuration Automation Database if the selection criteria does not include the second server, but it does not exercise the Soft Agent options. Instead, the product performs only an entity-level discovery for the second host (that is, it gathers the host name, IP, and operating system classification).

6. On the Exclusion page, repeat Step 5, but specify the host servers, IP addresses, and port number to *exclude* from the discovery.

   Note: When NDG discovers relationships, the product does *not* collect the relationship if either server in the relationship is targeted using the specified exclusion criteria.

7. Click Next.

8. On the Schedule page, select one of the following values from the Frequency drop-down list:

   Not Scheduled

   Specifies that the profile does not run automatically. You can run the profile manually or you can schedule it later.

   Once

   Specifies that the profile runs automatically one time. Specify when to run the profile in the Time field if you select this option.

   Minutes

   Specifies that the profile runs at specific intervals (in minutes). Define the following properties if you select this option:

   - **Start Time:** Set the time at which to start running the profile.
   - **Begin Date:** Set the first date on which to run the profile.
   - **End Date:** Set the last date on which to run the profile.
   - **Recur every # minutes:** Set the interval at which to run the profile.

   For example, to run the profile every 10 minutes starting at 11:00 p.m., specify a Start Time of 11:00:00PM and specify Recur every 10 minutes. The profile runs at 11:00 p.m., 11:10 p.m., 11:20 p.m., 11:30 p.m., and so on. The next profile starts when the current profile completes.

   Hourly

   Specifies that the profile runs at specific intervals (in hours). Define the following properties if you select this option:

   - **Start Time:** Set the time at which to start running the profile.
   - **Begin Date:** Set the first date on which to run the profile.
   - **End Date:** Set the last date on which to run the profile.
   - **Recur every # hours:** Set the interval at which to run the profile.

For example, to run the profile every four hours starting at 11:00 p.m., specify a Start Time of 11:00:00PM and specify Recur every 4 hours. The profile runs at 11:00 p.m., 3:00 a.m., 7:00 a.m., 11:00 a.m., 3:00 p.m., and so on. The next profile starts when the current profile completes.

Note: If the Start Time has already passed in the current day, the profile runs immediately and then the product resumes the specified recurring schedule.

Daily

Specifies that the profile runs at specific intervals (in days). Define the following properties if you select this option:

- **Start Time:** Set the time at which to start running the profile.
- **Begin Date:** Set the first date on which to run the profile.
- **End Date:** Set the last date on which to run the profile.
- **Recur every # days:** Set the interval at which to run the profile.

Weekly

Specifies that the profile runs at specific intervals (in weeks). Define the following properties if you select this option:

- **Start Time:** Set the time at which to start running the profile.
- **Begin Date:** Set the first date on which to run the profile.
- **End Date:** Set the last date on which to run the profile.
- **Days:** Set the days on which the profile runs every week.

Monthly

Specifies that the profile runs at specific intervals (in months). Define the following properties if you select this option:

- **Start Time:** Set the time at which to start running the profile.
- **Begin Date:** Set the first date on which to run the profile.
- **End Date:** Set the last date on which to run the profile.
- **Recur every # months:** Set the interval at which to run the profile on the specified days.

9. Define the notification that is used when the profile runs in the following fields:

   Notification Profile

   Defines which notification profile to use when discovery with this profile runs as scheduled. For information about creating notification profiles, see Create Notification Profiles (see page 98).

   Subject

   Defines the subject line of the email that the selected notification profile sends.

10. Click Finish.

    The product creates, enables, and adds the profile to the Network Profile table.

Enable Network Profiles

You must enable Network Profiles before they can be run manually or scheduled.

To enable one or more Network Profiles

1. Click the Management link, then the Network tab.

   The Network tab page appears.

2. Click the Network Profile link (below the main tabs).

   The Network Profiles page appears.

3. Select the check box next to one or more disabled profiles (they display No in the Is Enabled column), then select Enable Profile from the Select Actions drop-down list.

   The profile is enabled and the Is Enabled column displays a Yes for the selected profiles.

Run a Network Profile Manually

You can manually run an enabled Network Profiles if you want to perform an immediate discovery.

To run a Network Profile manually

1. Click the Management link, then the Network tab.

   The Network tab page appears.

2. Click the Network Profile link (below the main tabs).

   The Network Profiles page appears.

3. Select the check box next to a profile that displays Yes in the Is Enabled column, then select Run Network Profiles from the Select Actions drop-down list.

   You are prompted to confirm you want to run the profile.

4. Click OK.

   The profile runs and a Network Discovery is performed.

Rerun Profiles

You can exclude or include servers with the Soft Agents details during a discovery.

Follow these steps:

1. Click the Management link and then click the Network tab.
2. Click the Network Profile link.
3. Select a profile that displays Yes in the Is Enabled column, then select Rerun Profiles from the Select Actions drop-down list.

   You are prompted to confirm you want to run the profile.

   The servers with Soft Agent details are excluded, or the servers without the Soft Agent details are included in the network discovery.

Disable Network Profiles

You can disable Network Profiles when you want to prevent them from being run, but do not want to delete them in case you want to use them in the future.

To disable one or more Network Profiles

1. Click the Management link, then the Network tab.

   The Network tab page appears.

2. Click the Network Profile link (below the main tabs).

   The Network Profiles page appears.

3. Select the check box next to one or more enabled profiles (they display Yes in the Is Enabled column), then select Disable Profile from the Select Actions drop-down list.

   The profile is disabled and the Is Enabled column displays No for the selected profiles.

Delete Network Profiles

You can delete profiles that you no longer want to run or schedule.

To delete Network Profiles

1. Click the Management link, then the Network tab.

   The Network tab page appears.

2. Click the Network Profiles link (below the main tabs).

   The Network Profiles page appears.

3. Click the check box next to one or more profiles that you want to delete, then select Delete Profiles from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK to confirm the profile deletion.

   The selected profiles are deleted and removed from the Network Profiles table.

Import Network Profiles

You can import a Network Profile as a Java Archive (JAR) file from another instance of CA Configuration Automation.

To import a Network Profile

1. Click the Management link, then either the Network tab.

   The Network tab page appears.

2. Click the Network Profiles link (below the main tabs).

   The Network Profiles page appears.

3. Click Table Actions, then select Import Network Profiles.

   The Import Network Profiles dialog appears.

4. Enter or select the following information in the corresponding field:

   JAR File to Import

   Specifies the name of the JAR file that contains the Network Profile you want to import. You can click Browse to navigate to the file.

   Overwrite Existing Network Profiles

   Specifies whether the file being imported overwrites a file with the same name. Select this option if you want the changes made to the profile on another instance of CA Configuration Automation to be retained.

5. Click one of the following buttons:

   Import All

   Imports all of the Network Profiles in the JAR file.

   Import On Selection

   Displays a dialog where you can select the Network Profiles in the JAR file to import.

   The file is imported and the profiles appear in the Network Profiles table.

Export Network Profiles

You can export an Network Profile as a JAR file to use in another instance of CA Configuration Automation.

To export an Network Profile

1. Click the Management link, then either the Network tab.

   The Network tab page appears.

2. Click the Network Profiles link (below the main tabs).

   The Network Profiles page appears.

3. Click the check box next to one or more profiles you want to export, then click Select Actions and select Export Network Profiles.

   The File Download dialog appears.

4. Click Save.

   The Save As dialog appears and the export JAR file is assigned a default name using the following format:

   NetworkProfile_Export_*.jar

   For example: NetworkProfile_Export_2009_12_29_04_20_00.jar

5. Edit the file name if desired, select the location to save the file, and then click Save.

   The profile is exported to the selected location.

Network Scan Policies

Network Scan Policies specify which discovery engine is used and how the engine searches the network during discovery operations. You can create user-defined policies or use the predefined policies automatically installed by CA Configuration Automation. Network Scan Policies are assigned to Network Profiles, which are then assigned to networks.

Note: The predefined policies that were installed with CA Configuration Automation can be identified by the name *system_user* in the Created By column. If the default profiles have been modified, the Modified By column also displays a user name.

These policies and their descriptions can also be viewed in the UI, but because the policies can be edited, they are described in this document as they are originally installed.

CA Configuration Automation includes the following predefined Network Scan Policies:

ARP Cache Scan with Softagent

Specifies that the ARP Cache engine begins discovery with the supplied Gateway Router and recursively discovers routers in the network. It leverages SNMP to interrogate the ARP cache of each router in order to discover both the underlying entities as well as other routers in the network. This Scan Policy implements the option to ask the engine to interrogate the ARP cache of the underlying entities in order to provide a more exhaustive discovery of the network (at the cost of time). This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

ARP Cache Scan with Softagent (Routers Only)

Specifies that the ARP Cache engine begins discovery with the supplied Gateway Router and recursively discovers routers in the network. It leverages SNMP to interrogate the ARP cache of each router in order to discover both the underlying entities as well as other routers in the network. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

ARP Cache Scan with Softagent, but without Server Relationships

Specifies that the ARP Cache engine begins discovery with the supplied Gateway Router and recursively discovers routers in the network. It leverages SNMP to interrogate the ARP cache of each router in order to discover both the underlying entities as well as other routers in the network. This Scan Policy implements the option to ask the engine to interrogate the ARP cache of the underlying entities in order to provide a more exhaustive discovery of the network (at the cost of time). This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials. Note however, in this case, the Softagent processing is configured to ignore network connections and the detection of open ports.

ARP Cache Scan without Softagent

Specifies that the ARP Cache engine begins discovery with the supplied Gateway Router and recursively discovers routers in the network. It leverages SNMP to interrogate the ARP cache of each router in order to discover both the underlying entities as well as other routers in the network. This Scan Policy implements the option to ask the engine to interrogate the ARP cache of the underlying entities in order to provide a more exhaustive discovery of the network (at the cost of time).

ARP Cache Scan without Softagent (Routers Only)

The ARP Cache engine begins discovery with the supplied Gateway Router and recursively discovers routers in the network It leverages SNMP to interrogate the ARP cache of each router in order to discover both the underlying entities as well as other routers in the network. However, this Scan Policy restricts the ARP cache interrogation to the routers, and not their underlying entities in order to provide a more efficient network discovery (at the cost of exhaustive detail).

Cloud Service scan with Softagent

The Cloud Service engine performs discovery by connecting to the Cloud service provider against each of the possible IP address or host name targets. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials. Note however, in this case, the Softagent processing is configured to gather network connections which include the detection of open ports and the detection of remote machines that have established connections with the server currently being discovered. Also note that these remote machines are added to the CCA DB so that relationships can be conveyed between the two servers.

Note: Open the firewall for the WMI and SSH to perform the soft agent discovery when you select the *Cloud Service Scan with Soft agent* policy to discover the servers Ec2 cloud.

Cloud Service scan without Softagent

The Cloud Service engine performs discovery by connecting to the Cloud service provider against each of the possible IP address or host name targets.

DNS Scan with Softagent

Specifies that the DNS engine interrogates a given DNS server in order to discover all of the defined entities within a given DNS domain. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

DNS Scan with Softagent, but without Server Relationships

Specifies that the DNS engine interrogates a given DNS server in order to discover all of the defined entities within a given DNS domain. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials. Note however, in this case, the Softagent processing is configured to ignore network connections and the detection of open ports.

DNS Scan without Softagent

Specifies that the DNS engine interrogates a given DNS server in order to discover all of the defined entities within a given DNS domain.

IPv6 Local Link Scan with Softagent

In an IPv6 network, the Local Link engine discovers all entities on the local segment of the network. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

IPv6 Local Link Scan with Softagent, but without Server Relationships

Specifies that in an IPv6 network, the Local Link engine will discover all entities on the local segment of the network. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials. Note however, in this case, the Softagent processing is configured to ignore network connections and the detection of open ports.

IPv6 Local Link Scan without Softagent

Specifies that in an IPv6 network, the Local Link engine discovers all entities on the local segment of the network.

Netflow Analysis with Softagent (15 minutes)

Specifies that the Netflow Analysis engine leverages Netflow feeds from routers in order to passively monitor and analyze network traffic to identify entities, their applications, and their inter-machine relationships. In order to perform a continuous scan, this policy can be scheduled to run every 15 minutes. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

Netflow Analysis without Softagent (15 Minutes)

Specifies that the Netflow Analysis engine leverages Netflow feeds from routers in order to passively monitor and analyze network traffic to identify entities, their applications, and their inter-machine relationships. In order to perform a continuous scan, this policy can be scheduled to run every 15 minutes.

Packet Analysis with Softagent (15 Minutes)

Specifies that the Packet Analysis engine leverages packet sniffing methodology to passively monitor and analyze network traffic in order to identify entities, their applications, and their inter-machine relationships on the local segment of the network. In order to perform a continuous scan, this policy can be scheduled to run every 15 minutes. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

Packet Analysis without Softagent (15 Minutes)

Specifies that the Packet Analysis engine leverages packet sniffing methodology to passively monitor and analyze network traffic in order to identify entities, their applications, and their inter-machine relationships on the local segment of the network. In order to perform a continuous scan, this policy can be scheduled to run every 15 minutes.

Pingsweep Scan with Softagent

Specifies that the Pingsweep engine performs discovery based on brute force methodology, generating ICMP Ping requests against each of the possible IP address or host name targets. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

Pingsweep Scan with Softagent, but without Server Relationships

Specifies that the Pingsweep engine performs discovery based on brute force methodology, generating ICMP Ping requests against each of the possible IP address or host name targets. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials. Note however, in this case, the Softagent processing is configured to ignore network connections and the detection of open ports.

Pingsweep Scan without Softagent

Specifies that the Pingsweep engine performs discovery based on brute force methodology, generating ICMP Ping requests against each of the possible IP address or host name targets.

TCP Connect Scan with Softagent

Specifies that the TCP Connect engine performs discovery based on brute force methodology, generating TCP Connect requests against each of the possible IP address or host name targets. This engine is particularly useful in environments where ICMP Ping requests are blocked by firewalls. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials.

TCP Connect Scan with Softagent. but without Server Relationships

Specifies that TCP Connect engine performs discovery based on brute force methodology, generating TCP Connect requests against each of the possible IP address or host name targets. This engine is particularly useful in environments where ICMP Ping requests are blocked by firewalls. This scan policy also directs the use of the Softagent technology in order to provide a more detailed discovery of each entity for which it has credentials. Note however, in this case, the Softagent processing is configured to ignore network connections and the detection of open ports.

TCP Connect Scan without Softagent

Specifies that the TCP Connect engine performs discovery based on brute force methodology, generating TCP Connect requests against each of the possible IP address or host name targets. This engine is particularly useful in environments where ICMP Ping requests are blocked by firewalls.

If you use any of the predefined Softagent policies, or create a custom policy that uses Softagent technology for Linux or UNIX, you must edit the ssh_config file to include the following parameter:

```
PasswordAuthentication yes
```

By default, the ssh_config file is installed in the following locations:

- (HP) /etc/opt/ssh

(Linux, Solaris, and AIX) /etc/ssh

Create Network Scan Policies

You can create the Network Scan policies to configure the discovery methodology, and then define options to locate servers and software components on your networks.

Follow these steps:

1. Click the Management link and then click the Network tab.
2. Click the Network Scan Policies link.
3. On the Network Scan Policies page, select Create Policy from the Table Actions drop-down list.

4. On the Create Network Scan Policy wizard Policy page, enter the following information and then click Next.

   Name

   Defines the policy name.

   Description

   Describes the policy purpose and usage.

5. On the Discovery Engine page, select a discovery engine (for example, DNS or PingSweep), and then complete the fields for the selected option.

   DNS

   Specifies that the scan uses the Domain Name System (DNS) of hierarchical naming and numbering. DNS locates servers, services, or other network-connected resources.

   Complete the following fields:

   DNS Server IP Address

   Defines the IP address of the server that provides the DNS name resolution for the domain and the name servers of any subordinate domains.

   Domains

   Defines the domains that the profile is responsible for scanning. Enter a domain name in the Add New Domain field, and then click the right-facing arrow to move it to the Selected Domains field.

   Retries

   Defines how many Simple Network Management Protocol (SNMP) queries the product makes to an IP address before it fails while attempting to classify the operating system.

   Default: 1

   Timeout (in milliseconds)

   Defines how many milliseconds the SNMP query waits for a response before it fails.

   Default: 1000

   Note: Configure your DNS server so it allows Zone Transfers from the designated NDG server.

Cloud Service

NDG uses the cloud service to discover servers on the cloud environment.

Provide the cloud service discovery engine specifications:

Engine Instances

Defines how many discovery engine instances run during the discovery.

Default: 10

Retries

Defines how many times the discovery pings an IP address before it fails.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds a request waits for a response before it fails.

Default: 1000

PingSweep

Sends ICMP ECHO requests to determine which in a range of IP addresses maps to live hosts. If a specified address is live, the request returns an ICMP ECHO reply. The scan uses the reply to identify servers, services, or other network-connected resources.

Complete the following fields:

Engine Instances

Defines how many discovery engine instances run during the discovery.

Default: 10

Burst Size

Defines how many packets the product sends each second to the IP address.

Default: 32

Retries

Defines how many times the discovery pings an IP address before it fails.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds a request waits for a response before it fails.

Default: 2000

SNMP Classification Specifications

Defines the properties for monitoring the network devices and their functions.

Complete the following fields:

Retries

Defines how many SNMP queries the product makes to an IP address before it fails while attempting to classify the operating system.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds the SNMP query waits for a response before it fails.

Default: 1000

TCP Connect Scan

Determines the port availability through a TCP handshake connection. The scan uses an available port to identify network-connected servers, services, or other resources.

Complete the following fields:

Engine Instances

Defines how many discovery engine instances run during the discovery.

Default: 10

Retries

Defines how many SNMP queries the product makes to an IP address before it fails while attempting to classify the operating system.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds the SNMP query waits for a response before it fails.

Default: 1000

ARP Cache

Specifies whether the scan uses SNMP to interrogate the ARP Cache of routers. The ARP Cache of routers locates servers, services, or other network-connected resources.

Complete the following fields:

Engine Instances

Defines how many discovery engine instances run during the discovery.

Default: 10

Gateway IP Address

Defines the IP address of the computer that is the gateway for translating communication protocols.

Discover Router Information Only

Specifies whether to restrict the discovery process to the ARP Cache of routers.

Selected: The product only discovers the ARP Cache of routers.

Cleared: The product discovers the ARP Cache of all network resources.

Retries

Defines how many SNMP queries the product makes to an IP address before it fails while attempting to classify the operating system.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds the SNMP query waits for a response before it fails.

Default: 1000

Packet Analyzer

Analyzes packet data on the network, passively collects the IP traffic relationships, and identifies servers, services, or other network-connected resources.

Complete the following fields:

Execution Time

Defines how many days, hours, and minutes the scan runs.

Default: 15 minutes

Engine Instances

Defines how many discovery engine instances run during the discovery.

Default: 10

Cache Purge Frequency (in hours)

Defines how many hours elapse before the scan operation clears the cache. This scan type maintains a cache of discovered servers so that it does not continuously rediscover recently discovered servers.

Default: 8

Collect Network Statistics

Specifies whether the scan collects packet count summaries for the discovered relationships. To indicate the relationship strength, the packet count summary determines whether the servers exchanged a few or thousands of packets.

Selected: The scan collects packet count summaries.

Cleared: The scan does not collect packet count summaries.

Default: selected

Statistics Reporting Interval (in minutes)

Defines how many minutes elapse between network statistics collection operations.

Default: 15 (if you selected the Collect Network Statistics check box).

Discover Relationships

Specifies whether the scan discovers the relationships between network resources.

Selected: The scan discovers the relationships between network resources.

Cleared: The scan does not discover the relationships between network resources.

Default: selected

Relationship Packet Threshold Count

Defines the minimum number of packets the product requires to determine whether a relationship exists.

Default: 10 (requires that the Discover Relationships check box is selected).

Retries

Defines how many SNMP queries the product makes to an IP address before it fails while attempting to classify the operating system.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds the SNMP query waits for a response before it fails.

Default: 1000

NetFlow

Passively collects the IP traffic relationships and identifies servers, services, or other network-connected resources using the data feed from a NetFlow-enabled router.

Complete the following fields:

Execution Time

Defines how many days, hours, and minutes the scan runs.

Default: 15 minutes

Listen Port

Defines the NetFlow discovery engine port number.

Default: 9991

Note: Configure your router so it sends the NetFlow feed to the specified port on the designated NDG server.

Cache Purge Frequency (in hours)

Defines how many hours elapse before the scan operation clears the cache, and how often the operation rediscovers servers. This scan type maintains a cache of discovered servers so that it does not continuously rediscover recently discovered servers.

Default: 8

Discover Relationships

Specifies whether the scan discovers the relationships between network resources.

Selected: The scan discovers the relationships between network resources.

Cleared: The scan does not discover the relationships between network resources.

Default: selected

Relationship Packet Threshold Count

Defines the minimum number of packets the product requires to determine whether a relationship exists.

Default: 10 (requires that the Discover Relationships check box is selected).

Aggregate Records

Specifies whether the product collects network statistics for a discovered relationship in a single record.

Selected: The product collects network statistics in a single record.

Cleared: The product does not collect network statistics in a single record.

Default: cleared

Aggregation Interval

Specifies for how many minutes the product aggregates network statistics.

Default: 10 (requires that the Aggregate Records check box is selected).

Retries

Defines how many SNMP queries the product makes to an IP address before it fails while attempting to classify the operating system.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds the SNMP query waits for a response before it fails.

Default: 1000

Local Link

Discovers the servers on the local network segment using IPv6.

Complete the following fields:

Retries

Defines how many times the scan attempts to locate an IP address before it fails.

Default: 1

Timeout (in milliseconds)

Defines how many milliseconds the scan waits for a response before it fails.

Default: 2000

SNMP Retries

Defines how many SNMP queries the product makes to an IP address before it fails while attempting to classify the operating system.

Default: 1

SNMP Timeout (in milliseconds)

Defines how many milliseconds the SNMP query waits for a response before it fails.

Default: 1000

6. Click Next.
7. On the Discovery Options page, complete the following fields so the NDG can scan ports explicitly:

   VMware Web Services Port

   Defines the port that communicates with the VMware server.

   Default: 443

   Microsoft SCVMM Port

   Defines the port that communicates with the Microsoft System Center Virtual Machine Manager (SCVMM) server.

   Default: 8100

8. Select the Perform Soft Agent Probe check box if you want the agent-based discovery benefits without deploying an agent.

   The Soft Agent Probe uses supplied credentials to access the WMI services on Windows computers. The Soft Agent Probe uses SSH on UNIX and Linux target computers.

   If you select Perform Soft Agent Probe for Linux or UNIX, edit the ssh_config file to include the following parameter:

   ```
   PasswordAuthentication yes
   ```

   By default, the ssh_config file is installed in the following locations:

   HPUX

   ```
   /etc/opt/ssh
   ```

   Linux, Solaris, and AIX

   ```
   /etc/ssh
   ```

If you select the Perform Soft Agent Probe check box, complete the following fields:

Network Configuration

Specifies whether to discover network configuration settings.

Applications

Specifies whether to discover application configuration settings.

Virtual Environment

Specifies whether to discover servers and configuration settings for virtualized environments.

Restrict Discovery to Targeted Servers for Communications Relationships

Specifies whether to exclude the servers that the product discovers in the communication relationships. When you select this option, the product discovers the servers included in the network profile inclusion list.

Hardware

Specifies whether to discover hardware components.

Network Connections

Specifies whether to discover established network connections and open ports.

Select the check box, then click the Include Exclude Ports link to include or exclude specific ports during a network discovery.

Inclusions tab: In the left pane, double-click a mapped port to include it during a network scan.

Exclusions tab: In the left pane, double-click a mapped port to exclude it during a network scan.

Discover SAN Infrastructure and Relationships

Specifies whether to discover storage devices and storage managers and their relationships.

Enable use of Telnet

Lets you use Telnet to run a network discovery for UNIX and Linux server access when SSH-based discovery fails. Telnet discovery uses the same credentials as SSH discovery.

Note: Because the Telnet standards do not include encryption, the product communicates the user credentials from the credential vault in clear text.

Enable use of sudo

Specifies whether you can access and gather information from the remote UNIX and Linux servers with the sudo command. The sudo command lets the users that are defined in the /etc/sudoers configuration file run commands as if they had different (often unlimited, as for the root user) permissions.

If you enable sudo, comment the Default requiretty entry in the /etc/sudoers file as follows:

```
# Default requiretty
```

For more information, see Configuring sudo for UNIX and Linux Softagent Discovery (see page 405).

SSH Port

Specifies the port that the product uses for the SSH communications.

SSH Mode

Species one of the following modes:

- SSH with Credentials
- SSH with Key File and Credentials

  Note: If you select this option and SSH key file authentication fails, the product continues scanning with the UNIX credentials from the Credential Vault.

User Name

Defines the user that the product uses for the key file authentication.

Private Key File

Defines the private key file for the SSH authentication. Create the public and private key files with puttygen.exe or a similar utility. Copy the private key to the NDG Server that your CA Configuration Automation Server uses for discovery.

Note: For more information, see Create an SSH Key-based Network Scan Policy (see page 222).

Public Key File

Defines the public key file for the SSH authentication. Create the public and private key files with puttygen.exe or a similar utility. Copy the public key to the NDG Server that your CA Configuration Automation Server uses for discovery.

Note: For more information, see Create an SSH Key-based Network Scan Policy (see page 222).

Passphrase

Defines an optional key file protection passphrase. Associate this passphrase with the key files when you create them.

Enable user of SSH Proxy

Specifies the use of SSH Proxy.

Complete the following SSH Proxy fields:

Proxy Server

Defines the proxy server name or IP address.

Proxy Port

Defines the proxy server listening port.

9. Click Finish.

   The product creates the policy and adds it to the Network Scan Policies table.

Create an SSH Key-based Network Scan Policy

If you want to use SSH to access remote servers during network discovery, you can configure a CA Configuration Automation Network Scan Policy to use a public/private key pair to secure communications.

To generate the key pair files

1. Log on to a computer where SSH is installed, open a command window, and navigate to the ssh/bin directory.
2. Issue the following command to generate the public/private key pair:

   ```
   ssh-keygen -t rsa
   ```

   The following prompt appears:

   ```
   Generating public/private rsa key pair.
   Enter file in which to save the key (home/Administrator/.ssh/id_rsa):
   ```

3. Press Enter to accept the default names (id_rsa.pub and id_rsa).

   You are prompted to enter a passphrase:

   ```
   Enter passphrase (empty for no passphrase)
   ```

4. Enter a password (passphrase) or press enter to proceed without the password protecting the key pair.

   You are prompted to confirm the passphrase:

   ```
   Enter same passphrase again
   ```

5. Enter the password again, then press enter.

   The following confirmation appears:

   ```
   Your identification has been saved in /home/Administrator/.ssh/id_rsa
   Your public key has been saved in /home/Administrator/.ssh/id_rsa.pub
   The key fingerprint is:
   45:gd:b1:3e:c0:92:18:44:7b:e6:tc:d5:m1:6c
   ```

6. Copy the public and private keys (id_rsa.pub and id_rsa) to the NDG Server used by your CA Configuration Automation Server for discovery operations.

   You can copy the files to any folder.

7. Copy the public key (id_rsa.pub) to the computer that is the target of your discovery (that is, the server you want to discover and manage using CA Configuration Automation) using one of the following methods:

   - Use secure copy from the command line to add the id_rsa.pub key to the authorized_keys file, for example:

     `scp id_rsa.pub root@targethost:~/.ssh/authorized_keys`

   - If secure copy is not available, you must either FTP or copy the file to the target host and then manually append the contents of the id_rsa.pub key to the authorized_keys file.

   Note: The target server must have the SSH server software installed.

To create a Network Scan policy that uses the key pair files

1. Perform steps 1 through 6 as described in Create Network Scan Policies.

   The Discovery Options page appears.

2. Provide the following ports for the NDG to the scan ports explicitly:

   VMware Web Services Port

   Specifies the port to communicate with the VMware server.

   Default: 443

   Microsoft SCVMM Port

   Specifies the port to communicate with the Microsoft *System Center Virtual Machine Manager* (SCVMM) server

   Default: 8100

3. Click the Perform Soft Agent Probe check box, then specify the following discovery options and SSH parameters:

   Network Configuration

   Specifies whether to discover network configuration settings.

   Applications

   Specifies whether to discover application configuration settings.

Virtual Environment

Specifies whether to discover servers and configuration settings of virtualized environments.

Hardware

Specifies whether to discover hardware components.

Network Connections

Specifies whether to discover established network connections and open ports.

Select the check box, then click the Include Exclude Ports link to include or exclude specific ports during a network discovery.

Inclusions tab: In the left pane, double-click a mapped port to include it during a network scan.

Exclusions tab: In the left pane, double-click a mapped port to exclude it during a network scan.

SSH Port

Specifies the port that is used for the SSH communications.

Default: 22

SSH Mode

Species one of the following modes: SSH with Credentials or SSH with Key File and Credentials. Select SSH with Key File and Credentials.

User Name

Specifies the user name that is used for the key file authentication.

Private Key File

Specifies the location and private key file to use for the SSH authentication. Enter the path to the private key file (id_rsa) on the NDG Server used by your CA Configuration Automation Server (step 6 in the previous procedure).

Public Key File

Specifies the location of public key file to use for the SSH authentication. Enter the path to the public key file (id_rsa.pub) on the NDG Server used by your CA Configuration Automation Server (step 6 in the previous procedure).

Passphrase

Specifies an optional key file protection passphrase. This passphrase must be associated with the key files when they are created (step 4 in the previous procedure). Leave this field blank if you did not create a passphrase.

If you click the Enable use of SSH Proxy check box, you can specify the following SSH Proxy options:

Proxy Server

Specifies the name or IP address of the proxy server.

Proxy Port

Specifies the listening port of the proxy server.

4. Click Finish.

   The policy is created and appears in the Network Scan Policies table.

5. Create a Network Discovery Profile that uses the new Network Scan Profile as described in Create Network Profiles.

6. Perform a Discovery of the target servers using the Network Profile.

   Softagent data about the target servers is discovered and available in CA Configuration Automation.

View Network Scan Policies

Network Scan Policies specify which discovery engine is used and how the engine searches the network. You can view the user-defined policies and predefined policies on the Network Scan Policies page.

To view predefined and user-defined Network Scan Policies

1. Click the Management link, then the Network tab.

   The Network tab page appears.

2. Click the Network Scan Policies link (below the main tabs).

   The Network Scan Policies page appears and displays all the Network Scan Policies.

   Note: The predefined policies that were installed with CA Configuration Automation can be identified by the name system_user in the Created By column. If the default profiles have been modified, the Modified By column also displays a user name.

Import Network Scan Profiles

You can import a Network Scan Profile as a Java Archive (JAR) file from another CA Configuration Automation instance.

Follow these steps:

1. Click the Management link, then click the Network tab.

2. On the Network Profiles tab, click the Network Scan Profiles link.

3. On the Network Scan Profile page, click Table Actions, then select Import Network Scan Profiles.
4. On the Import Network Scan Profiles dialog, complete the following fields:

   JAR File to Import

   Defines the name of the JAR file that contains the Network Scan Profile to import. Click Browse to navigate to the file.

   Overwrite Existing Network Scan Profiles

   Specifies whether to overwrite a file with the same name. Select this option to retain the profile from another CA Configuration Automation instance.
5. Click one of the following buttons:

   Import All

   Imports all of the Network Scan Profiles in the JAR file.

   Import On Selected

   Displays a dialog on which to select specific Network Scan Profiles to import from the JAR file.

   The application imports the file and the Network Scan Profiles table displays the profiles.

Edit Network Scan Policy Details

You can edit any existing Network Scan Policy.

To edit Network Scan Policies

1. Click the Management link, then the Network tab.

   The Network tab page appears.
2. Click the Network Scan Policies link (below the main tabs).

   The Network Scan Policies page appears.
3. Click the name of the policy you want to edit.

   The Details page for the selected policy appears.
4. Edit the fields on any of the tabs as appropriate, then click Save. For a description of the fields, see Create Network Scan Policies.

   The policy is updated with your edits.

Credential Vault Profiles

Credential Vaults store a set of credentials that are used during the discovery process. Credential Vaults include credentials for accessing computers using SNMP (versions 1, 2, and 3), as well as the standard UNIX and Windows user-based credentials. In order to discover VMware environments, Credential Vaults also include credentials to access VMware web services.

Note: Within any particular type of access, the order of precedence is important. Access is attempted by first using the credentials at the top of the list and then working down the list in the configured order.

Create Credential Vault Profiles

Use the following procedure to create Credential Vault Profiles for server and storage device access in your enterprise when you run discovery operations.

Follow these steps:

1. Click the Management link, then click the Network tab.
2. On the Network tab page, click the Credential Vault link below the main tabs.
3. On the Credential Vault Profiles page, select Create Credential Vault Profile from the Table Action drop-down list.
4. On the Create Credential Vault Profile wizard, complete the following fields, then click Next:

   Name

   Defines the name of the profile.

   Description

   Describes the purpose of the profile.

   Default

   Specifies whether the profile is used as the default profile for discovery. You can designate only one profile as the default Credential Vault Profile.
5. On the Windows page of the wizard, complete the following fields, click Test to ensure the credentials are valid, then click Add:

User ID

Specifies the user name the profile uses to connect to computers using a Windows operating system. It can also be used to access virtual Microsoft Hyper-v and operating system-based Windows virtualization (VMware Server, VMware Workstation, Microsoft Virtual Server, or Microsoft Virtual PC).

Password

Defines the password that is associated with the specified User ID.

Verify Password

Requires the user to retype the password that is associated with the specified User ID.

Description

Describes the purpose of the credentials.

The wizard adds the credentials to the Windows Credentials table.

6. Click Next.
7. On the UNIX page of the wizard, complete the following fields, click Test to ensure the credentials are valid, then click Add:

User ID

Defines the user name the profile uses to connect to computers that use a UNIX operating system. The profile also uses this ID to access the following supported virtual environments:

- Citrix Xen
- IBM HMC
- VMware ESX
- Solaris Zones

Password

Defines the password that is associated with the specified User ID.

Verify Password

Requires the user to retype the password that is associated with the specified User ID.

Description

Describes the purpose of the credentials.

The wizard adds the credentials to the Unix Credentials table.

8. Click Next.

9. On the SNMP v1v2 wizard page, complete the following fields, then click Add:

   SNMP Community Name

   Defines the name of the SNMP v1 or v2 community. An SNMP community groups the devices and management stations that run SNMP to help define where to send data.

   Default: public

   Description

   Describes the purpose of the community.

   The wizard adds the new community below the public community in the SNMP v1 & v2 Credentials table.

10. Click Next.

11. On the SNMP v3 wizard page, complete the following fields, click Test to ensure the credentials are valid, then click Add:

    User ID

    Defines the user name the profile uses to connect to the SNMP engine.

    Authentication Protocol

    Specifies one of the following authentication protocol algorithms:

    None

    The profile does not authenticate the client.

    MD5

    The profile uses the MD5 (Message Digest) challenge and response mechanism to authenticate the client.

    SHA

    The profile uses the Secure Hash Algorithm (SHA) mechanism to authenticate the client.

    Authentication Password

    Defines the password that is associated with the specified User ID.

    Verify Authentication Password

    Requires the user to retype the password that is associated with the specified User ID.

Privacy Protocol

Specifies the privacy protocol the profile uses for the selected authentication protocol.

None

The profile does not use a privacy protocol.

AES

The profile uses the Advanced Encryption Standard (AES) privacy protocol.

DES

The profile uses the Data Encryption Standard (DES) privacy protocol.

Privacy Passphrase

Defines the privacy passphrase that is associated with the user.

Verify Privacy Passphrase

Requires the user to retype the privacy passphrase.

Description

Describes the purpose of the credentials.

The wizard adds the credentials to the SNMP v3 Credentials table.

12. Click Next.
13. On the VMware wizard page, complete the following fields, click Test to ensure the credentials are valid, then click Add:

User ID

Defines the user name the profile uses to connect to computers that use VMware vCenter or VMware ESX. This value establishes the credentials that you use to access these server types with the VMware web client.

Password

Defines the password that is associated with the specified User ID.

Verify Password

Requires the user to retype the password that is associated with the specified User ID.

Description

Describes the purpose of the credentials.

The wizard adds the credentials to the VMware Credentials table.

14. Click Next.

15. On the Red Hat Enterprise Virtualization wizard page, complete the following fields, click Test to ensure the credentials are valid, then click Add:

    User ID

    Defines the user name the profile uses to connect to computers that use Red Hat Enterprise Virtualization.

    Password

    Defines the password that is associated with the specified User ID.

    Verify Password

    Requires the user to retype the password that is associated with the specified User ID.

    Description

    Describes the purpose of the credentials.

    The wizard adds the credentials to the Red Hat Enterprise Virtualization Credentials table.

16. Click Next.

17. On the NetApp wizard page, complete the following fields, click Test to ensure the credentials are valid, then click Add:

    User ID

    Defines the user name the profile uses to connect to NetApp OnCommand.

    Password

    Defines the password that is associated with the specified User ID.

    Verify Password

    Requires the user to retype the password that is associated with the specified User ID.

    Description

    Describes the purpose of the credentials.

    The wizard adds the credentials to the NetApp Credentials table.

18. Click Next.

19. On the SMI-S wizard page, complete the following fields, click Test to ensure the credentials are valid, then click Add:

    User ID

    Defines the user name the profile uses to connect to Storage Management Initiative – Specification (SMI-S) devices.

    Password

    Defines the password that is associated with the specified User ID.

Verify Password

Requires the user to retype the password that is associated with the specified User ID.

Description

Describes the purpose of the credentials.

The wizard adds the credentials to the SMI-S Credentials table.

20. On the AWS EC2 wizard page, complete the following fields:

Access Key ID

Defines the Access Key ID of the cloud service when you sign up for a cloud service account.

Secret Access Key

Defines the secure key of the cloud service account.

Signature Method

Specifies the algorithm the cloud service uses to calculate the signature on cloud.

Signature Version

Specifies the version of the cloud service signature method.

Web Service Version

Specifies the cloud Service version.

Description

Describes the purpose of the credentials.

Note: We support the HTTP protocol to access EC2 web services.

21. Click the Test button on the right corner to test the cloud credentials and click Ok. Complete the following field:

AWS EC2 End Point

Describes the AWS EC2 End Point name to which you send the HTTP request. The Network Discovery Gateway field value is populated.

22. (Optional) Click the Add button to provide the cloud credentials of the other EC2 cloud environment to discover the servers. The AWS EC2 credentials are save in the database.

23. Click Finish.

The wizard creates the profile and adds it to the Credential Vault Profiles table.

Set a Credential Vault Profile as the Default

You can designate one Credential Vault Profile as the default profile for discovery operations that specify the Default profile be used.

To designate a Credential Vault Policy as the default

1. Click the Management link, then click the Network tab.

   The Network tab page appears.

2. Click the Credential Vault link (below the main tabs).

   The Credential Vault Profiles page appears and displays the existing profiles.

3. Click the check box next to the profile you want to designate as the default profile, then select Set As Default Profile from the Select Action drop-down list.

   Yes appears in the Is Default column for the selected profile.

Delete Credential Vault Profiles

You can delete a Credential Vault Policy if you no longer need it.

To delete a Credential Vault Policy

1. Click the Management link, then click the Network tab.

   The Network tab page appears.

2. Click the Credential Vault link (below the main tabs).

   The Credential Vault Profiles page appears and displays the existing profiles.

3. Click the check box next to one or more profiles you want to delete, then select Delete Profiles from the Select Action drop-down list.

   Note: You cannot delete a profile if it has been designated as the default profile (that is, No appears in the Is Default column for the selected profile). You must designate another profile as the default before deleting the profile.

   The selected profiles are deleted.

Chapter 8: Storage Management

This section contains the following topics:

Discover and Manage Storage Devices

A configuration manager can:

- Discover storage systems on a storage area network (SAN)
- Manage the discovered storage systems from the Storage tab page in the CA Configuration Automation Server UI

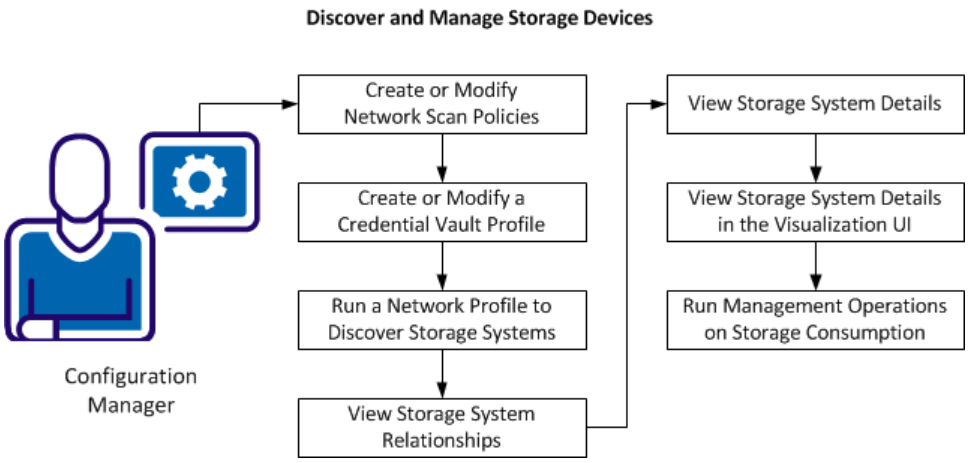

A configuration manager can perform the following tasks to discover and manage storage devices:

1. Create or Modify Network Scan Policies (see page 236)
2. Create or Modify a Credential Vault Profile (see page 236)
3. Run a Network Profile to Discover Storage Systems (see page 238)
4. View Storage System Relationships (see page 238)
5. View Storage System Details (see page 239)
6. View Storage System Details in the Visualization UI (see page 242)
7. Run Management Operations on Storage Consumption (see page 242)

Create or Modify Network Scan Policies

You can create custom Network Scan Policies or you can modify existing policies to discover storage devices and storage relationships.

Note: CA Configuration Automation includes the following predefined Network Scan Policies to discover storage devices on your network:

- Pingsweep Scan with Softagent - SAN
- TCP Connect Scan with Softagent - SAN

If you create custom policies or you modify existing policies, consider using the - *SAN* naming convention. The naming convention lets you easily identify the Network Profiles with policies configured to discover storage devices.

Follow these steps:

1. In the CA Configuration Automation Server UI, navigate to the Network Scan Policies page:
   a. Click the Management panel.
   b. Click the Network tab.
   c. Click the Network Scan Policies link.
2. On the Network Scan Policies page, create a policy:
   a. Select Create Network Scan Policy from the Table Actions drop-down list.
   b. Create a Network Scan Policy.

      Note: For more information, see the CA Configuration Automation Online Help.
   c. On the Discovery Options page, click the Discover SAN Infrastructure and Relationships check box.
3. On the Network Scan Policies page, modify an existing policy:
   a. Click an entry in the Policy Name column.
   b. On the details page for the selected policy, click the Discovery Options tab.
   c. On the Discovery Options tab, click the Discover SAN Infrastructure and Relationships check box.

Create or Modify a Credential Vault Profile

You can create a storage-specific Credential Vault Profile to store the access credentials that are required to discover storage devices and storage relationships.

Note: The Credential Vault Profile requires access to the SAN Manager software. CA Configuration Automation gets storage information from the SAN Manager for all devices it manages.

Follow these steps:

1. In the CA Configuration Automation Server UI, click the Management panel, the Network tab, then the Credential Vault link.
2. On the Credential Vault Profiles page, create a profile:
   a. From the Table Actions drop-down list, select Create Credential Vault Profiles.
   b. Create a Credential Vault Profile as described in the CA Configuration Automation Online Help.
   c. On the NetApp tab, enter the credentials that are required to access the NettApp OnCommand server, then click Test.

      A confirmation appears if the credentials are valid.
   d. Click Add to add the credentials to the profile.
   e. On the SMI-S tab, enter the credentials that are required to access the SMI-S provider, then click Test.

      Note: EMC Clariion and IBM Storwize V7000 are the only SMI-S storage devices that CA Configuration Automation supports.

      A confirmation appears if the credentials are valid.
   f. On the Windows tab, enter the credentials that are required to access any Windows servers that access storage from the SAN storage devices, then click Test.

      A confirmation appears if the credentials are valid.
   g. On the UNIX tab, enter the credentials that are required to access any UNIX (or VMware ESX) servers that access storage from the SAN storage devices, then click Test.

      A confirmation appears if the credentials are valid.
3. Click Add to add the credentials to the profile.
4. Click Finish to create the profile.
5. On the Credential Vault page, modify an existing profile:
   a. Click an entry in the Profile Name column.
   b. On the details page for the selected profile, click the NetApp or SMI-S tab, edit or add the appropriate credentials, and click Test.

      A confirmation appears if the credentials are valid.
   c. Click Add to add the credentials to the profile, then click Finish to save the profile.

Discover Storage Systems

To discover storage devices and relationships, run a Network Profile manually. When the discovery is complete, view the results to verify that the profile discovered the expected devices. When you are satisfied with the results, schedule the profile to run at regular intervals.

Follow these steps:

1. In the CA Configuration Automation Server UI, navigate to the Network Profiles table:
   a. Click the Management panel.
   b. Click the Network tab.
   c. Click the Network Profiles link.
2. In the Network Profiles table, click the check box next to one or more profiles to run.

   Note: Network Profiles used to discover storage devices and relationships must use a Network Scan Policy that has the Discover SAN Infrastructure and Relationships check box selected. The policy that is assigned to the Network Profile appears in the Network Scan Policy column. The predefined policies that are configured for storage discovery use the - SAN naming convention.
3. Select Run Profiles from the Select Actions drop-down list.

   The Network Policy runs and the results appear in the locations described in the next three sections.

View Storage System Relationships

You can view storage system relationships in the following CA Configuration Automation Server UI locations:

Storage Systems page

Displays relationship details from a storage system context as described in View Storage System Details (see page 239).

Storage Relationships page

Displays all storage system relationships including storage system manager relationships. The Storage System Managers table identifies which servers manage the related storage systems. These servers run storage system-specific software (either NetApp, EMC, or IBM) that lets storage area network (SAN) administrators configure the storage system.

Follow these steps:

1. On the CA Configuration Automation Server UI, click the Management panel, then the Storage tab, then the Relationships link.

   The Relationships page displays the following tables:

   Storage System Managers

   Displays the relationships between the server where the storage manager software is installed (Server Name) and the storage systems it manages (Storage System Name). Other columns display the Manufacturer and the Model Type.

   Storage Relationships

   Displays the relationships between the Server that uses the storage system (Server Name), Disk Name (if there is a physical disc), Server LUN Identifier, Storage Infrastructure (iSCSI or Fibre Channel Protocol), Storage System Name, and LUN Name.

2. In the Storage System Managers table, do one or both of the following tasks:
   - To display the Server Details page for the storage manager host server, click an entry in the Server Name column.
   - To display the Storage System Details page, click an entry in the Storage System Name column.
3. In the Storage Relationships table, do one or both of the following tasks:
   - To display the Server Details page for the storage system host server, click an entry in the Server Name column.
   - To display the Storage System Details page, click an entry in the Storage System Name column.

View Storage System Details

CA Configuration Automation displays storage system details in the following places:

- Storage Systems page
- Storage System Details page
- Server Details page

Note: The storage details on the Server Details page can be used in management operations (for example, Change Detection and Rule Compliance) as described in Run Management Operations on Storage Devices (see page 242). The other pages contain read-only details.

Follow these steps:

1. On the CA Configuration Automation Server UI, click the Management panel, the Storage tab, and then the Storage Systems link.

   The Storage Systems page displays the following details:

   - Name
   - Unique Identifier
   - Serial Number
   - Manufacturer
   - Model/Type
   - Storage Capacity (GB)
   - Fail Over System
   - Discovered Time

2. Click the storage system name in the Name column.

   The System Details tab displays the following information:

   - Hardware (a summary of the Storage System table)
   - Storage Processors (Specific to the EMC storage system)
   - iSCSI Identifiers
   - Fibre Channel World Wide Node Names (WWNNs)

3. Click the LUNs tab.

   The LUNs tab displays the following information that varies as per the selected storage system:

   - LUN Name
   - Storage Infrastructure
   - Serial Number
   - Network Address Authority
   - Storage Capacity (MB)

4. Click the Relationships tab.

   The Relationships tab displays the following information for the selected storage system:

   - Storage System Manager Server Name (this value links to the Server Details page; continue to Step 5 for the storage details available on the Server Details page)
   - Server Name (this value links to the Server Details page; continue to Step 5 for the storage details available on the Server Details page)

- Disk Name
- Server LUN Identifier
- Storage Infrastructure
- LUN Name

Note: To display relationships for *all* storage systems as described in View Storage System Relationships (see page 238), click the Relationship link on the Storage System tab.

5. On the Server tab of the Server Details page, do one or both of the following tasks:
   a. Click the Network Adapters link to view the following Storage area network adapter details.
      - iSCSI
      - Fibre Channel
   b. Click the Hardware link to view the following details:
      - Processor
      - Memory
      - BIOS
      - Physical Disks
      - Logical Partitions
      - File Systems
      - CD and DVD Drives
      - Tape Drives
6. On the Relationships tab of the Server Details page, click the Storage link to view the following storage details:
   - Disk Name
   - Server LUN Identifier
   - Storage Infrastructure
   - Storage System Name
   - LUN Name

View Storage System and Storage Manager Relationships in the Visualization UI

You can display a graphical representation of storage systems and their relationships to servers and storage system managers in the Visualization UI.

Follow these steps:

1. On the CA Configuration Automation Server UI, click the Management panel, then click the Storage tab.
2. From the Storage Systems table on the Storage tab page, select one or more storage systems.
3. From the Select Actions drop-down list, select Visualization, then select one of the following options:

   Server Storage Management Relationships

   Displays the storage system and the server where the storage management software is installed. The relationship type for this option is always Manages.

   Server Storage Relationships

   Displays the storage system and all servers with a relationship to the storage system. The relationship types for this option can be:

   - Manages
   - Uses (as in *Uses Storage*)

   The storage system and related servers appear in the Visualization UI.

Run Management Operations on Storage Consumption

You can perform the following management operations on servers that consume storage from storage systems identified in the Server Details page:

- Run Discovery
- Run Management Profile
- Take Snapshot
- Run Change Detection
- Run Rule Compliance

Follow these steps:

1. On the Management panel, click the Servers tab, then click an entry in the Server Name column.

   Note: You can also display the Server Details page from various locations on the Storage tab, as described in View Storage System Details (see page 239).

2. On the Server Details page, do one of the following tasks:

   - Click the Relationships tab, then click the Storage link
   - Click the Network Adapters link
   - Click the Hardware link

3. Click Actions (on the right side of the UI, above the Server Details table), select Management Actions, then select an operation.

Chapter 9: Blueprint Management

A *Blueprint* is the abstract or metadata definition of a software component. Blueprints contain directives describing how to locate an installed component instance and how to account for and understand all of the component's elements. These elements include the installed file system, registry variables, data in databases, and other runtime assets. Once located, Blueprints apply notes and rules to each element enabling even deeper understanding of their function and constraints.

CA provides a library of predefined Blueprints covering most common software components (such as web servers, operating systems, and databases). The predefined Blueprints appear in the Blueprints table and are identified as being created by the internal user *system_user.* You can also create your own custom Blueprints. Custom Blueprints are identified as being created by the user name of the person who creates them.

Create Blueprints

CA Configuration Automation lets you build and maintain custom blueprints. Although simple blueprints can be easy to build, detailed and complex blueprints require careful planning and testing.

Follow these steps:

1. Click Management in the upper right of the main product page, and then click the Blueprints tab on the upper left.

   The Blueprints pane opens, listing all existing blueprints.

2. Select Create Blueprint from the Table Actions drop-down list in the Blueprints pane.

   The Create Blueprint wizard opens.

3. Complete the fields (see page 456) on the Blueprint page Component Blueprint pane, and then click Next.

   The Discovery Methods page opens. It displays the File Indicators pane and the Add New Search Options pane.

4. Complete the Search Options fields (see page 458)on the Add New Search Options pane.

5. Click Add Directory/File, and then complete the File Indicators fields (see page 458) in the Add New File pane.

6. Click the Registry Indicators link.

   The Registry Indicators pane displays with Registry Indicators expanded and the \HKEY_LOCAL_MACHINE element selected. The \HKEY_LOCAL_MACHINE pane displays the Search Options fields.

7. Complete the Search Options fields (see page 460) in the \HKEY_LOCAL_MACHINE pane, and then click Save.

8. Click Add Registry Value/Key.

9. Complete the Registry Indicators fields (see page 460) in the Add New Registry Value/Key pane.

10. Click the Network Probes link.

    The Network Probes pane displays with the Network Probes element selected. The Add New Network Probe pane displays the Network Probe fields.

    On servers without CA Configuration Automation agents, discovery operations can use network probes to:

    - Scan ports
    - Collect the responses from the scanned ports
    - Determine the type of service that is active on a scanned port by comparing the responses against expected expressions

11. Complete the Network Probe fields (see page 461), and then click Next.

    The Discovery Verification Rules page displays with the Discovery Verification Rules element selected in the left pane. The Add New Discovery Verification Rule pane displays the Discovery Verification Rule fields.

    The product runs the verification rules during discovery to verify that the discovered components are correctly identified. In some cases, the file and registry indicators cannot determine the existence of installed components. Similarly, the file and registry are sometimes unable to distinguish between two components with similar indicators. If a verification rule fails, the component discovery fails.

12. Complete the Discovery Verification Rule fields (see page 462), and then click Next.

    The product saves the Discovery Verification Rule values, and then displays the Management page. The $(Root) folder is selected in the File Management pane, and the $(Root) pane displays the File Management Options fields. The Management page links the following pages:

    - File Filters and Attributes
    - Registry Management
    - Registry Filters and Attributes
    - Data Management

    These pages let you define important file attributes, registry entries, and database elements that are associated with this managed component.

    If no files or registry entries are defined, the product manages all files under the component root directory and registry entries under the registry. If a component has a limited number of files or registry entries, allow all files and registry entries under the root to manage them. However, for a complex component with many files, specify only the important directories, files, and registry entries on which to focus. Identifying specific files and registry entries from the Management page lets you refine the managed component view.

13. Complete the File Management Options fields (see page 467) in the Management page $(Root) pane.

14. (Optional) Click Add Directory, complete the Directory fields (see page 468) on the Add New Directory pane, and then click Save.

    The product adds the directory below the $(Root) directory in the File Management pane.

15. (Optional) Click Add File, complete the File fields (see page 468) on the Add New File pane, and then click Save.

    The product adds the file below the $(Root) directory in the File Management pane.

16. (Optional) Select a node in the File Management pane and repeat Steps 14 and 15 as appropriate to create subdirectories and other nested files.

    The product adds the new elements below the selected node in the File Management pane.

17. Click the File Filters and Attributes link.

    The File Filters and Attributes pane displays the directory and file structure that you created in Steps 14 through 16 below the $(Root) folder.

18. Define filters and attributes for the blueprint file.

19. Click the Registry Management link, complete the Registry Management fields (see page 470), and then click Save.

20. (Optional) Click Add Key, complete the Key fields on the Add New Key pane, and then click Save.

    The product adds the key below the $(RegistryRoot) directory in the Registry Management pane.

21. (Optional) Click Add Value, complete the Value fields on the Add New Value pane, and then click Save.

    The product adds the file below the $(RegistryRoot) directory in the Registry Management pane.

22. (Optional) Select a node in the Registry Management pane and repeat Steps 20 and 21 as appropriate to create other nested keys and values.

    The product adds the new keys or values below the selected node in the Registry Management pane.

23. Define registry filters and attributes for the blueprint. (see page 252)

    The product adds the key below the $(RegistryRoot) directory in the Registry Management pane.

24. Click the Data Management link.

    The left pane displays the Data Management folder, and the right pane displays the Database page.

25. Complete the fields (see page 474) on the Database page, and then click Save.

    The database appears in the Data Management pane.

26. Click Next.

    The Component Parameters and Variables page opens.

27. Complete the fields on the Component Parameters and Variables page, and then click Save.

28. Click Next.

    The Configuration - File Parsing page opens.

29. Complete the fields (see page 486) on the Configuration - File Parsing page, and then click the Configuration Executables link.

30. Complete the fields (see page 487) on the Configuration Executables page, and then click Save.

31. Click the Configuration Data link, and then complete the Database field, which defines the database that the blueprint uses. The drop-down list displays the databases that you created in Step 23.

32. Click Save. The Configuration Data tree in the left pane displays the database.

33. Complete the fields (see page 492) on the Add Query pane, and then click Save. The Configuration Data tree in the left pane shows the query.

34. Click the File Structure Data link, complete the fields (see page 493) on the File Structure Class tab, and then click Save. The File Structure Class tree in the left pane shows the structure class.
35. Click the Precedence tab, click Add Group or Add Parameter, and then complete the displayed fields (see page 494).
36. Click Next, complete the fields (see page 497) on the Macros page, and then click Next.

    The Component Grouping Options page opens. The page contains options that let you nest components in a service for display to emphasize the relationships between them. For example, when a component depends on subordinate components in the primary component file system root, you can use nesting to enforce the parent-child relationship.

    Oracle databases, for example, typically install a Java Runtime Engine and an Apache Web server in the installation directory. The product expresses the relationship between the utility components and the Oracle database by nesting the JRE and Apache in the Oracle component.
37. Complete the fields (see page 498) on the Component Grouping Options page, and then click Finish.

    The product creates the blueprint, which then appears in the Blueprint table.

Define Blueprint File Filters and Attributes

Follow these steps:

1. Select the $(Root) folder.

   The Precedence table lists the directories in $(Root).

2. (Optional) Select a column, and then select Move Up or Move Down from the Select Actions drop-down list to change the directory precedence.

   Consider order when you apply meta-links to File Structure Classes (descriptions, rules, filters, and categories) on the parsed data (parameters and groups) and overlays. The Change Detection, the Compare, and the Rule Compliance operations consider the order to process the content to match correct values.

   For example, consider the following parameter definitions in a Blueprint File Structure Class, each with its own rules, category filters, and weights:

   ```
   ab.*
   a.*
   .*
   ```

   - All parameters that start with *a* get their respective filters.
   - All parameters that start with *ab* get their respective filters. In this case, the product overrides the preceding *a.*.*
   - All parameters that do not start with *a* or *ab* get the respective filters that .* specifies.

   Therefore, define more specific parameters first, and then define more generic parameters.

3. (Optional) Repeat Step 2 for the files that the Files table lists.
4. Assign a category or create a filter for the selected directory or file.
5. Click a file or directory, and then double-click one or more options in the Available Categories column.

   The product adds your selections to the Selected Categories column.

6. Double-click one or more options in the Available Filters column.

   The product adds your selections to the Selected Filters column.

7. Select the file or directory Weight, and then click Save.
8. (Optional) Click the Manage Filters tab.

   The Manage Filters tab lists the sub folders and files of the selected folder. The Manage Filters tab lets you apply or remove the Never Run Change Detection, and Time Variant filters to specific folders, its sub folders, and files.

9. (Optional) Select the sub folders and files from the list, and then select one of the following actions from the Select Actions drop-down:

   - Set Never Run change Detection

     Apply the Never Run change Detection filter to the selected folders, its sub folders, and files that are defined in the File Filters and Attributes pane.

   - Set Time Variant

     Apply Time Variant filter to selected folders, its sub folders, and files that are defined in the File Filters and Attributes pane.

   - Remove Never Run change Detection

     Remove the Never Run change Detection filter from selected folders, its sub folders, and files that are defined in the File Filters and Attributes pane.

   - Remove Time Variant

     Remove the Time Variant filter from selected folders, its sub folders, and files that are defined in the File Filters and Attributes pane.

   A message confirms whether the Never Run Change Detection and Time Variant filters are updated for the folder, its sub folders, and files.

   Note: If you add new sub folders or files to a folder, the existing filter does not apply to the newly added sub folder or files. Apply a new filter to the newly added sub folder or file when required.

10. Click the Rules tab.

    The Rules tab defines rules that constrain file and directory values in the Managed - File System overlay. The rules include both explicit constraint rules that you create and predefined, implicit constraint rules. For example, if you specify a value or data type for an element, CA Configuration Automation automatically creates an implicit Check Default or Verify Data Type rule.

11. Complete the fields (see page 468) on the Rules tab.

Define Blueprint Registry Filters and Attributes

Follow these steps:

1. Click the Registry Filters and Attributes link.

   The Registry Filters and Attributes pane displays the $(RegistryRoot) folder and the $(RegistryRoot) pane displays the Precedence tab. The tab shows existing keys and values in the corresponding tables.

2. (Optional) Select a row by which to set the directory precedence, and then select Move Up or Move down from the Select Actions drop-down list.

   The product reorders the directories. The importance of sequencing keys and values is similar to the sequencing described in Step 18.

3. Click Add Key, and then complete the Key pane Name (regex) and Description fields.

4. Double-click one or more options in the Available Categories column. The product adds them to the Selected Categories column.

5. Double-click one or more options in the Available Filters column. The product adds them to the Selected Filters column.

6. Complete the remaining fields in the Key pane, and then click Save.

7. In the Registry Filters and Attributes pane, select the key to which to assign a value.

8. In the right pane, click Add Value, and then complete the Value pane Name (regex) and Description fields.

9. Double-click one or more options in the Available Categories column to add them to the Selected Categories column.

10. Double-click one or more options in the Available Filters column to add them to the Selected Filters column.

11. Complete the remaining fields in the Value pane, and then click Save.

    The product adds the key below the $(RegistryRoot) directory in the Registry Management pane.

Browse Servers to Locate Blueprint Elements

You can use the Browse button while creating or editing Blueprints to do the following on managed Windows servers:

- Find files and directories.
- Find registry keys, values, and data.
- Add found elements to the Blueprint.

The Browse button is located in the following locations when creating Blueprints in the Blueprint Wizard, or editing them in either the Tree View or Tab View on the Blueprint Details page:

- Discovery Methods page (wizard step 2)
  - File Indicators, Add New Search Options
  - File Indicators, Add New File and Add New Directory
  - Registry Indicators, \HKEY_LOCAL_MACHINE Search Options
  - Registry Indicators, \HKEY_LOCAL_MACHINE Add New Registry Value/Key
- Management page (wizard step 4)
  - File Management, $(Root) File Management Options
  - File Management, $(Root) Add Directory
  - File Management, $(Root) Add File
  - File Filters and Attributes, $(Root)
  - File Filters and Attributes, $(Root) Add Directory
  - File Filters and Attributes, $(Root) Add File
  - Registry Management, $(RegistryRoot), Registry Management Options
  - Registry Management, $(RegistryRoot), Add Key
  - Registry Management, $(RegistryRoot), Add Value
  - Registry Filters and Attributes, $(RegistryRoot)
  - Registry Filters and Attributes, $(RegistryRoot), Add Key
  - Registry Filters and Attributes, $(RegistryRoot), Add Value
- Configuration page (wizard step 6)
  - File Parsing, $(Root)

You can create or edit Blueprints in either of the following views:

- Tab View (this includes the creation Wizard)
- Tree View

If you are creating or editing a Blueprint from the Tab View, clicking the Browse button opens the Browse Server dialog set to either File Browser or Registry Browser depending on which page you click Browse. For example, if you are on the Registry Management page, and click Browse, the Browse Server dialog appears set to Registry Browser. If you are on the File Management page, and click Browse, the Browse Server dialog appears set to File Browser. In the Tab View, the Browse utility is in the context of whichever tab was selected. If you click Browse on a File Management tab, you can only add that type of element.

If you are using the Tree View, the Browse Server dialog appears with two tabs: File and Registry. In the Tree View, Browse is global; you can add any type of element.

To locate and add files, directories, and symbolic links using Browse

1. Click Browse in any of the previously mentioned locations.

   The Browse Servers dialog appears.

2. Select a server to browse from the Server drop-down list.

   Servers are listed in alphabetical order by name. By default, the Server field shows the first server in the drop-down list.

3. Click Test Connectivity to ensure communications with the selected server.

   A message confirms the connection.

4. (Tree View only) Ensure the File Browser tab is selected.

   If the server selected in the Server field is a Windows server, two tabs are displayed: File and Registry.

5. In the File Search Directory field, either enter a path to the file or directory or leave the field blank to use the default root path (/ or \), then click Browse.

   All items found in the specified File Search Directory appear in a tree view.

6. Narrow your search to a directory displayed in the tree:

   a. Select a directory.

   b. Click Set Search to narrow your search to the selected directory.

      The selected directory appears in the File Search Directory field and on top of the tree.

   c. Click Set Root.

      The selected directory becomes the root and appears in the Component Files Root and on top of the tree.

7. Navigate the tree and select an element you want to add to the Blueprint, then do one of the following:

   - If the element is located in the File Structure Class folder, click Parse.

     The contents of the file are parsed into component-specific attributes and parameters that define the contents of file.

     The attributes and parameters are added to the Blueprint and are highlighted in yellow. Elements highlighted in yellow have been added using the Browse Servers option and have not been saved.

   - If the element is located in any other folder, click Blueprint.

     The Blueprint dialog appears. Select one or more of the options that specify how the element is referenced in the Blueprint. The following options vary depending on the selected element: Indicators, Managed Files, Filters, and Configuration Files.

     The selected element is added to the Blueprint and highlighted in yellow. Elements highlighted in yellow have not been saved.

8. Click Close.

   The Browse dialog closes.

9. Click Save for each highlighted element.

   The Blueprint is updated.

To find registry keys, values, and data using Browse

1. Click Browse in any of the previously mentioned locations.

   The Browse Servers dialog appears.

2. Select a server to browse from the Server drop-down list.

   Servers are listed in alphabetical order by name. By default, the Server drop-down list shows the first server in the list.

3. Click Test to ensure Agent communications with the selected server.

   A message confirms the connection.

4. (Tree View only) Click the Registry Browser tab.

   If the server selected in the Server field is a Windows server, two tabs appear: File Browser and Registry Browser.

5. Enter a root key in the Registry Search Key field, or leave the field blank to use the default key (\HKEY_LOCAL_MACHINE), then click Browse.

   All items found in the specified Registry Search Key are displayed in a tree view.

6. Narrow your search to a registry key displayed in the tree as follows:
   a. Select a registry key.
   b. Click Set Search.

      The selected registry key becomes the Registry Search Key in the browser's tree view display.

   c. Click Set Root.

      The selected directory becomes the Component Registry Root in the browser's tree view display.
7. Navigate the tree and select an element you want to add to the Blueprint, then click Blueprint.

   The Blueprint dialog appears.
8. Select one or more of the options that specify how the element is referenced in the Blueprint. The following options vary depending on the selected element: Search Options, Indicators, Managed Registry, Filters, and Configuration Files.

   The selected element is added to the Blueprint and highlighted in yellow. Elements highlighted in yellow have not been saved.
9. Click Close.

   The Browse dialog closes.
10. Click Save for each highlighted element.

    The Blueprint is updated.

Test Discover Component Blueprints

Discovery enables you to inspect servers for the presence of software or software components. The success of Discovery, however, depends on certain key elements being specified properly within the Blueprint.

Test Discovery lets you verify that you have set up many of the Component Blueprint elements correctly while you are developing a new or editing an existing Component Blueprint. You can use Test Discovery to verify such things as indicators, registry elements, and managed and configuration files.

Note: Test Discovery cannot test for software components if the specified servers are not available or none of the components are found.

To test Blueprint Discovery success on a server

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the name of the Blueprint on which you want to test Discovery.

   The Blueprint Details page for the selected Blueprint appears.

3. Click Test Discovery.

   The Test Discovery page appears in a new browser window.

4. Select a managed server from the Server drop-down list.

   By default, the drop-down list initially displays the first managed server in the list.

5. Do one of the following:

   - In the File Search Directory field, enter the full path to the file or directory you want to test discover, then click Discover. If you leave the field blank, the search uses the default root path (/ or \).

     Test Discovery searches for either files or directories down to the level in the file system that matches the path text string you specify in the Search Root and displays the results in a tree view.

   - Click Browse to display the file system structure found on the specified server starting with the default root path (/ or \). Then select the directory and file elements in the tree view down to the location you want to test discover.

     The File Search Directory field displays the directory and file elements as you select them in the in the tree view. When you are satisfied with the contents of the File Search Directory field, click Discover.

   The discovery runs and one of the following appears:

   - An error message appears if the CA Configuration Automation Agent on the selected server is not communicating with the CA Configuration Automation Server.

   - The following message appears if Test Discovery cannot find components using the specified indicators:

     ```
     No components found with Test Discovery
     ```

- The results appear under a unique, system-generated name that identifies the discovery as a test. The name includes a time stamped Blueprint name and version and the word Test as shown in the following example:

  `Apache HTTP Server (UNIX) v1.0.0 (Test)`

- If the following error message appears you must run the Reconcile Servers operation as described in Reconcile Server IPs (see page 137):

  `IP address of host host_name conflicts with another host in the DB - Run Reconcile Servers to correct it`

Edit Blueprints in the Tabbed View

You can view and edit existing Blueprints using the Blueprint wizard. We recommend you copy any predefined Blueprint before you make changes to it.

To edit a Blueprint in the wizard

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears and displays the existing Blueprints.

2. Copy the Blueprint you want to edit as described in Copy Blueprints (see page 261).

   The copy of the Blueprint appears in the Blueprints table.

3. Click the link to the copy of the Blueprint in the Blueprint Name column.

   The Blueprint Details page appears and displays the Blueprint tab.

   Note: This page contains a Save As button (top right) where you can also make a copy of the predefined Blueprint instead of using the Copy Blueprints functionality.

4. Edit any of the fields on the Blueprints tab, or click any of the other seven tabs and edit the appropriate field. You cannot make edits on the Log tab. For a description of the fields on each tab, see Create Blueprints.

5. Click Save.

   A message confirms the Blueprint was updated.

Edit Blueprints in the Tree View

Blueprints can be viewed and edited in a hierarchical tree view instead of using the Blueprint wizard.

To edit Blueprints in the tree view

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears and displays the existing Blueprints.

2. Copy the Blueprint you want to edit as described in Copy Blueprints (see page 261).

   The copy of the Blueprint appears in the Blueprints table.

3. Click the link to the copy of the Blueprint in the Blueprint Name column.

   The Blueprint Details page appears and displays the Blueprint tab.

   Note: This page contains a Save As button (top right) where you can also make a copy of the predefined Blueprint instead of using the Copy Blueprints functionality.

4. Click Tree View (top right).

   A tree view of the Blueprint appears is the Blueprint pane. Details about the selected node appear in the Details pane. For a description of the fields, see Create Blueprint.

5. Click the nodes in the tree view to locate the element you want to edit.

   The editable fields appear in the details pane.

6. Edit the fields as needed, then click Save.

   A message confirms the Blueprint was updated, or, if you used the Save As button, a message confirms the Blueprint was created. The modified copy of the original Blueprint appears in the Blueprints table.

Import Blueprints

An import utility is provided to import new or modified Component Blueprint data into your CA Configuration Automation Server. This enables you to import new and upgraded versions of Component Blueprint files without having to re-install or upgrade your CA Configuration Automation Server installation. CA Configuration Automation provides new and upgraded Component Blueprints as .jar files that you can import.

A similar export utility is provided to export new or modified Component Blueprint data.

To import Component Blueprints

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprints link (below the main tabs)

   The Blueprints page appears and lists all existing blueprints in the Blueprints table.

3. Select Import Blueprints from the Table Actions drop-down list.

   The Import Blueprints dialog appears.

   a. Click Browse and navigate to the .jar file you want to import.

      If the file was exported from a CA Configuration Automation Server, the file was assigned a default name using the following timestamp convention:

      Blueprint_Export_*YYYY_MM_DD_HH_MM_SS*.jar

      Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

      Blueprint_Export_2010_08_01_16_20_00.jar

   b. (Optional) Select the following option.

      Overwrite Existing Blueprint Groups

      Specifies that the version of Blueprint currently on the CA Configuration Automation Server is replaced by the version in the .jar file.

   c. Click one of the following buttons:

      Import All

      Imports all Blueprints in the .jar file.

      Import On Selection

      Displays the Available Imports table which lists the Component Blueprints in the .jar file, and enables you to specify which of them you want to import.

   The import begins and may take a few moments depending on the size of the file. When the import is complete, a status message appears, the imported Blueprints appears in the Blueprints table, and the Modification Date/Time column is updated.

We strongly recommend that you do not make modifications to any of the predefined Component Blueprints delivered with CA Configuration Automation for the following reasons:

- Modifying a predefined Component Blueprint makes upgrading to newer versions of that Component Blueprint difficult in the future.
- Some Component Blueprints share data, including directives and macros, so changing one Component Blueprint can inadvertently affect another.
- CA thoroughly tests the Component Blueprints as delivered. Modifications can break base Component Blueprint functionality.

If customization of a predefined Component Blueprint is required, consider copying the Component Blueprint, then modifying it.

To copy a Component Blueprint

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprints link (below the main tabs)

   The Blueprints page appears and lists all existing blueprints in the Blueprints table.

3. Click the check box next the Component Blueprint that you want to copy, then select Copy Blueprint from the Select Actions drop-down menu.

   The Copy Blueprint dialog appears and displays the following details for the selected Blueprint:

   - Component Blueprint Name
   - Component Version
   - Blueprint Version

4. Modify any or all fields to distinguish the new Blueprint from the predefined version, then click OK.

   The new version of the Blueprint is added to the table. The Created By column displays the login name of the user who created the new Blueprint.

Delete Blueprints

You can permanently delete Blueprints if they are no longer needed.

To delete one or more Blueprints

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprints link (below the main tabs)

   The Blueprints page appears and lists all existing blueprints in the Blueprints table.

3. Click one or more check boxes next the Component Blueprints that you want to delete, then select Delete Blueprints from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   The selected Blueprints are deleted.

Search for Blueprints References

The following blueprint references are maintained in CA Configuration Automation:

- Blueprint Groups
- Compare
- Management Profiles
- Rule compliance results
- Server test discovery
- Server or Service discovered components
- Snapshots
- Server or Service change detection

These references restrict the deletion of a blueprint. The Blueprint Search action lets you search the references of selected blueprints and delete them. You can only one search for blueprint at a time.

Follow these steps:

1. Click the Management link, then click the Blueprints tab.
2. Select a blueprint from the Blueprint page to search for related references.
3. Select Blueprints Search from the Select Actions drop-down list.

   The blueprint references are displayed on the Blueprint Search Results page.

4. (Optional) To delete the reference, select the reference type and then select delete from the Select actions drop-down list.

The deletion of the blueprint reference is based on the reference types. The following list displays the different reference types and the action for the reference.

- Server (Service) Current data—Deletes the component that is discovered for a given blueprint.
- Server (Service) Snapshots—Deletes the entire snapshot with a component that the blueprint uses.
- Management profile—Deletes the blueprint references from the management profile.
- Blueprint Group—Deletes the blueprint references from the blueprint group.
- Blueprint Test Discovery data—Removes the blueprint reference after the test discovery browser window is closed. If Test Discovery blueprint reference is listed in Blueprint Search result page, deleting the reference shows a warning message to close the Test Discovery browser window. Closing the browser window deletes the blueprint reference.
- Discovery in progress data—If the blueprint references selected for deletion contain server discovery data then the references are not deleted. A warning message is displayed about the job being in progress. Refresh the action to delete the reference after the respective job is completed.
- Diff/RC results—The change detection, compare, or rule compliance results of a server that are obtained via Management Profile run are displayed in the search result.

  Note: Diff/RC results of ad-hoc change detection, compare, or rule compliance operations are not considered for the blueprint search references.

Disable or Enable Discovery on Blueprints

Blueprints can be disabled to prevent them from being used during discovery operations. This makes the discovery faster, and enables you to prevent software components that you are not interested in managing from being added to the CA Configuration Automation inventory.

Discovery is enabled by default on all predefined Component Blueprints.

To disable or enable discovery on Component Blueprints

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprints link (below the main tabs)

   The Blueprints page appears and lists all existing blueprints in the Blueprints table.

3. Click the check box next to one or more Component Blueprints whose discovery state you want to change, then select Disable Discovery or Enable Discovery from the Select Actions drop-down list.

   The Discovery Status column is updated to reflect the current state of the associated Blueprint.

Export Blueprints

An export utility is provided to export new or modified Component Blueprint data from one CA Configuration Automation implementation so it can be imported into another. The export utility extracts the data files for the selected Component Blueprint, then compresses and saves the Blueprint data to a .jar file.

To export Component Blueprints

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprints link (below the main tabs)

   The Blueprints page appears and lists all existing blueprints in the Blueprints table.

3. Click the check box next to one or more Component Blueprints that you want to export, then select Export Blueprints from the Select Actions drop-down list.

   A File Download window prompts you for the location where you want to save the export .jar file.

The file is assigned a default name using the following timestamp convention:

Blueprint_Export_*YYYY_MM_DD_HH_MM_SS*.jar

Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

Blueprint_Export_2010_01_06_13_59_57.jar

4. Click Save, specify the location, then click OK.

   Note: If you are using a Windows operating system, Windows may suggest saving the file as a .zip file. Ensure the .jar is included in the file name, and the Save As Type field is set to All Files.

   The Component Blueprint is saved in the specified location.

Create Blueprint Groups

You can organize your Blueprints into logical groupings, for example all UNIX operating systems. Blueprint groups can speed up management operations by letting you select Blueprints as a group instead of selecting them individually.

To create a Blueprint Group

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprint Groups link.

   The Blueprint Groups page appears.

3. Select Create Blueprint Group from the Table Actions drop-down list.

   The Create Blueprint Group page appears.

4. Enter the following information:

   Name

   Specifies the name for the Blueprint Group.

   Description

   Describes the intended purpose or function of the Blueprint Group.

5. Click Next.

   The Blueprints page appears.

6. Double-click one or more Blueprints that you want to add to the group in the Available Blueprints column.

   The Blueprints are moved to the Selected Blueprints column.

7. Click Finish.

   The Blueprint Group is created and appears in the Blueprint Groups table.

View and Edit Blueprint Groups

You can view and edit existing Blueprint Groups.

To view or edit Blueprint Groups

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears and displays the existing Blueprint Groups.

2. Click a link in the Group Name column.

   The Blueprint Groups Details page appears and displays the Blueprint Group tab.

3. Edit the name or description of the group, or click The Blueprints tab.

   The Blueprints page appears.

4. Do one of the following:

   - Double-click one or more Blueprints in the Available Blueprints column to add them to the selected Blueprint Group.
   - Double-click one or more Blueprints in the Selected Blueprints column to remove them from the selected Blueprint Group.

   The Blueprints are moved to the opposite column.

5. Click Save.

   The Blueprint Group is updated.

Import Blueprint Groups

An import utility is provided to import Blueprint Groups into your CA Configuration Automation Server from another instance of CA Configuration Automation. A similar export utility is included to export Blueprint Groups as .jar files that you can import.

To import Blueprint Groups

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprint Groups link (below the main tabs)

   The Blueprint Groups page appears and lists all existing groups in the Blueprint Groups table.

3. Select Import Blueprint Groups from the Table Actions drop-down list.

   The Import Blueprints dialog appears.

4. Click Browse and navigate to the .jar file you want to import.

   When the file was exported from a CA Configuration Automation Server, the file was assigned a default name using the following timestamp convention:

   BlueprintGroup_Export_*YYYY_MM_DD_HH_MM_SS*.jar

   Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

   Blueprint_GroupExport_2010_01_06_13_59_57.jar

5. Click the Overwrite Existing Blueprint Groups option if you want to update the existing groups on your CA Configuration Automation Server, then click one of the following buttons:

   Import All

   Imports all Blueprint Groups in the .jar file.

   Import On Selection

   Displays the Available Imports table which lists the Blueprint Groups in the .jar file, and enables you to specify which of them you want to import.

   The import begins and may take a few moments depending on the size of the file. When the import is complete, a status message appears, the Blueprint Group appears in the Blueprint Groups table, and the Modification Date/Time column is updated.

   Note: The Blueprint Groups table does not include a column that identifies the user who imported the Blueprint Group. You can click the Log link (below the main tabs) to display the log and identify the user who imported the file.

Export Blueprint Groups

An export utility is provided to export Blueprint Groups as .jar files from your CA Configuration Automation Server to be used in another instance of CA Configuration Automation. A similar import utility is included to import Blueprint Groups.

To export Blueprints Groups

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprint Groups link (below the main tabs)

   The Blueprint Groups page appears and lists all existing groups in the Blueprint Groups table.

3. Select Export Blueprint Groups from the Select Actions drop-down list.

   The File Download dialog appears.

4. Click Save, specify where you want to save the .jar file, and then click Save again.

   The file is assigned a default name using the following timestamp convention and saved in the selected location:

   BlueprintGroup_Export_*YYYY_MM_DD_HH_MM_SS*.jar

   Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

   BlueprintGroup_Export_2010_01_06_13_59_57.jar

Delete Blueprint Groups

You can delete one or more existing Blueprint Groups if you no longer need them.

To delete Blueprint Groups

1. Click the Management link, then the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Blueprint Groups link.

   The Blueprint Group page appears.

3. Click the check box next to the Blueprint Groups you want to delete, then select Delete Blueprint Groups from the Select Actions drop-down list.

   The selected Blueprint Groups are deleted.

View and Edit Structure Classes

You can view and edit existing Structure Classes. CA Configuration Automation includes parsers that dissect the classes into component-specific attributes and parameters that, for example, define the specific contents of service configuration files.

Before editing any predefined Structure Class, consider copying it—as described in Copy Structure Classes (see page 271)—and making edits to the copy.

To view or edit Structure Classes

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Structure Classes link.

   The Structure Classes page appears and displays the existing predefined and custom Structure Classes.

3. Click the link in the Structure Class Name column of the class you want to view or edit.

   The Structure Classes Details page appears and displays the File Structure Class tab.

4. Edit any of the fields, then click Save.

   A description of each field can be found in Create Structure Classes (see page 269).

   The Structure Class is updated.

Create Structure Classes

In addition to the predefined classes, you can create custom Structure Classes. CA Configuration Automation includes parsers that dissect the classes into component-specific attributes and parameters that, for example, define the specific contents of service configuration files.

To create a Structure Class

1. Click the Management link, then click the Component Blueprints tab.

   The Blueprints tab page appears.

2. Click the Structure Classes link (below the main tabs).

   The Structure Classes page appears.

3. Select Create Structure Class from the Table Actions drop-down list.

   The Create Structure Class page appears.

4. Enter the following information in the corresponding field:

   Name

   Specifies a unique name for the Structure Class.

   Version

   Specifies the version number for the class.

   Default: 1.0

   Display Name

   Specifies the name that appears in the structure class table. This name does not need to be unique as in can be differentiated when combined with the release number (for example, Display Name 1.0, Display Name 1.1, and so on).

   Description

   Describes the function or purpose of the class.

   Allow Remediation Jobs

   Specifies whether remediation is allowed for this Structure Class.

   Parsers

   Specifies the parser used to parse files associated with this Structure Class.

5. Click Save.

   The new Structure Class is created.

Import Structure Classes

An import utility is provided to import Structure Classes into your CA Configuration Automation Server from another instance of CA Configuration Automation. A similar export utility is included to export Structure Classes as .jar files.

To import Structure Classes

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Structure Classes link (below the main tabs)

   The Structure Classes page appears and lists all existing classes in the Structure Classes table.

3. Select Import Structure Classes from the Table Actions drop-down list.

   The Import Structure Classes dialog appears.

4. Click Browse and navigate to the .jar file you want to import.

   When the file was exported from a CA Configuration Automation Server, the file was assigned a default name using the following timestamp convention:

   FileStructureClass_Export_*YYYY_MM_DD_HH_MM_SS*.jar

   Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

   FileStructureClass_Export_2010_01_06_13_59_57.jar

5. Click the Overwrite Existing Structure Classes option if you want to update the existing version on your CA Configuration Automation Server, then click one of the following buttons:

   Import All

   Imports all Structure Classes in the .jar file.

   Import On Selection

   Displays the Available Imports table which lists the Structure Classes in the .jar file, and enables you to specify which of them you want to import.

   The import begins and may take a few moments depending on the size of the file. When the import is complete, a status message appears, the Structure Classes appears in the Structure Classes table, and the Created By and the Modification Date/Time columns are updated.

Copy Structure Classes

You can copy a Structure Class then modify it instead of making modifications to the predefined Structure Classes shipped with CA Configuration Automation.

To copy a Structure Class

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Structure Classes link (below the main tabs)

   The Structure Classes page appears and lists all existing classes in the Structure Classes table.

3. Click the check box next the Structure Class that you want to copy, then select Copy Structure Class from the Select Actions drop-down menu.

   The Copy Structure Class dialog appears and displays the following details for the selected Structure Class:

   - Name
   - Version
   - Display Name

4. Modify any or all fields to distinguish the new Structure Class from the predefined version, then click OK.

   The new version of the Structure Classes is added to the table.

Delete Structure Classes

You can permanently delete Structure Classes if they are no longer needed.

To delete one or more Structure Classes

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Structure Classes link (below the main tabs)

   The Structure Classes page appears and lists all existing classes in the Structure Classes table.

3. Click one or more check boxes next the Structure Classes that you want to delete, then select Delete Structure Classes from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   The selected Structure Classes are deleted.

Export Structure Classes

An export utility is provided to export Structure Classes from one CA Configuration Automation implementation so it can be imported into another. The export utility extracts the Structure Classes, compresses, encrypts, and saves the Structure Classes data to a .jar file.

To export Structure Classes

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Structure Classes link (below the main tabs)

   The Structure Classes page appears and lists all existing classes in the Structure Classes table.

3. Click the check box next to one or more Structure Classes that you want to export, then select Export Structure Classes from the Select Actions drop-down list.

   A File Download window prompts you for the location where you want to save the export .jar file.

   The file is assigned a default name using the following timestamp convention:

   FileStructureClass_Export_*YYYY_MM_DD_HH_MM_SS*.jar

   Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

   FileStructureClass_Export_2010_01_06_13_59_57.jar

4. Click Save, specify the location, and then click OK.

   Note: If you are using a Windows operating system, Windows may suggest saving the file as a .zip file. Ensure the .jar is included in the file name, and the Save As Type field is set to All Files.

   The Structure Class is saved in the specified location.

View Parsers

The Parsers page displays a list and description of the parsers included with CA Configuration Automation. These predefined parsers can be used to parse most of the common configuration file formats.

To view the available parsers

1. Click the Management link, then click the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Parsers link (below the main tabs).

   The available parsers and are listed and described in the Parsers table.

Manage Global Variables

Global Variables enable you to define a set of name-value pairs that can be used in variable substitution expressions within directives, parameters, and macros. Global variables are independent of the typical server- or service-specific CA Configuration Automation data model, and can be defined manually within the CA Configuration Automation user interface, imported from an existing XML or CSV file, or exported to a CSV file.

CA Configuration Automation displays the global variable repository in a table on the Global Variables page. Global variables require unique variable names and can include only alphanumeric characters, periods (.), parentheses (( )), and underscores (_).

The named values defined in the global variable repository are:

- Stored and evaluated only as strings (they have no specific data type)
- Optional and do not contain default values (not all names require a value and can simply be used to group together other named values)

The Global Variables page lets you create, export, import, and delete the global variables.

To manage the global variables, follow the steps:

1. Click the Management link, then the Blueprints tab.

   The Blueprints tab page appears.

2. Click the Global Variables link (below the main tabs).

   The Global Variables page appears.

To create a global variable:

1. Select Create Global Variable from the Table Actions drop-down list.

   The Create Global Variable dialog appears.

2. Enter the following information in the corresponding field:

   Name

   Identifies the global variable.

   Description

   Describes the purpose of the global variable.

   Value

   Specifies the syntax for addressing global variables for variable substitution using the $(@global_variable_name) format. This expression represents the value of the global variable c at the path /a/b in the Global Variables tree.

   Note: You cannot use variable substitution expressions with name qualifiers that contain colons (:), for example: $(/var1/var2/var3/var3(:2:)/value) is not a valid expression.

3. Click OK to display the new global variable in the tree listing.

To Import the global Variables that have been exported from another instance of CA Configuration Automation:

1. Select Import Global Variables from the Table Actions drop-down list.

   The Import Global Variables dialog appears.

2. Click Choose file and navigate to the exported global variable data file.

   The selected file appears in the Import XML or CSV File field.

3. Click the Ignore Duplicate or Invalid Variables check box to ignore any duplicate or invalid variables in the global variable data file you are importing.

   By default, all entries are imported.

4. Click OK.

   The global variables are imported and appear in the Global Variables table.

To export global variables to an Excel file in a comma separated value (CSV) format:

1. Select Export to Excel from the Table Actions drop-down list.

   The ExportToCsvServlet runs in a web browser, then displays a File Download window.

2. Click Save.

3. The Save As dialog appears with the default file name of GlobalVariables.csv.

4. Edit the default file name, select a location to save the file, and then click Save.

   The CSV file is saved in the specified location.

To delete the global variables:

1. Select one or more global variables that you want to delete.
2. Select Delete Global Variables from the Table Actions drop-down list.
3. Click OK when prompted to confirm the deletion.

   The selected global variables are removed from the table.

Chapter 10: Blueprint Element Reference

The topics in this section provide a reference to the various blueprint elements that can be used to discover and manage software components.

This section contains the following topics:

Understanding the Structure and Contents of a Component Blueprint

The ability to view, add, or modify a Component Blueprint or Component Blueprint elements depends on the CA Configuration Automation roles assigned to your user account or group. See the *CA Configuration Automation Implementation Guide* for detailed information about CA Configuration Automation roles and privileges.

All Component Blueprints are presented in a standardized tree view, consisting of the following organizational folder elements:

- Nesting
- Indicators
- Managed
- Parameters
- Configuration
- Diagnostics
- Documentation
- Runtime
- Utilities

Each element is described in a section that follows.

Nesting

The nesting element lets you emphasize the relationship between components. When, for example, a software component uses and depends on several subordinate components, and those components are installed within the primary component's file system root, nesting can be used to enforce the parent-child relationship between them.

Oracle databases, for example, normally install a Java Runtime Engine and an Apache Web server within their installation directory. The relationship between the utility components and the Oracle database is expressed by nesting the JRE and Apache components within the Oracle component.

Indicators

The Indicators primary folder defines where to look for and how to identify a component on a server.

Because component discovery has been extended to servers without CA Configuration Automation Agents installed, and CA Configuration Automation does not have the registry and file systems available for this kind of discovery, you need to define different sets of indicators to find components on servers with installed agents and to find components on servers without installed agents.

- On servers with CCA Agents, component discovery uses file indicators to search the file system for a specific set of files and directories that indicates the presence of a component, and then determines the corresponding root directory under which the managed files for the component are located.

  For Windows servers, you can alternatively define registry entries as indicators. Note that if you use registry entries as indicators, you must specify the component's root directory in the parameters primary folder.

  Note: While both file system and registry indicators can be used on Windows platforms, we recommend that you define one or the other, not both. You can, however, define more than one set of indicator type, which allows multiple pattern matching to extend the same Component Blueprint across multiple platforms. A good strategy for building an indicator set is to define two to three files/directories or registry keys/values that have a known position relative to one another as defined in Depth From Root or Path From Root. Defining too many indicators can produce undesirable and incorrect results.

- On servers without CCA Agents, component discovery uses network probes to scan TCP ports, collect responses from those ports, and compare the responses against expected expressions to determine the type of service active on that port.

Indicators are not always sufficient to determine the existence of installed components or cannot distinguish the differences between two components with similar indicators. The verification directives folder contains directives that are used to discard components that are partially installed, of the wrong version, or not of interest to a particular service.

Note: Because of the order in which verification directives are executed in component discovery, database and configuration file parameters may not be available.

Adding Indicator Elements

To Add a...	Access and Procedure
File Search Option	Click on the files folder, then click the Add Search Options button.
Directory Indicator	Click on the files > / (file root) folder, then click the Add Directory button.
File Indicator	Click on the files > / (file root) folder, then click the Add File button.
Registry Search Option	Click on the registry folder, then click the Add Search Options button.
Registry Key Indicator	Click on the registry > \ (registry root) folder, then click the Add Key button.
Registry Value Indicator	Click on the registry > \ (registry root) folder, then click the Add Value button.
Network Probes	Click on the service folder, then click the Add Network Probe button.
Verification Directive	Click on the verification directives folder, then click the Add Directive button.

Managed

The Managed primary folder defines important files, registry entries, and database elements that are associated with the managed component.

If the managed folder contains no files or registry entries, the application manages all files under the component root directory and registry entries under the registry root. If a component has a limited number of files or registry entries, it makes sense to let all such elements under the root be managed. However, for a complex component with many files, it can be more useful to identify only the files and registry entries on which to focus. Identifying specific files and registry entries lets you refine the view of the managed component.

The File System Overlay and Registry Overlay folders let you (optionally) customize specific file and registry elements. For example, you can assign the categories, filters, or weights, or you can attach notes and rules to an element.

The data folder provides the database connection information and lets you:

- Define a database to reference elsewhere in the Component Blueprint
- Manage the database metadata, including table definitions and indexes

Important! Include any file that is referenced in the configuration, documentation, or run-time primary folders as a managed file.

Adding Managed Elements

To Add a...	Access and Procedure
Directory	Click on the files > $(Root) folder, then click the Add Directory button.
File	Click on the files > $(Root) folder, then click the Add File button.
File System Overlay Directory	Click on the file system overlay > $(Root) folder, then click the Add Directory button.
File System Overlay File	Click on the file system overlay > $(Root) folder, then click the Add File button.
Registry Key	Click on the registry > $(RegistryRoot) folder, then click the Add Key button.
Registry Value	Click on the registry > $(RegistryRoot) folder, then click the Add Value button.
Registry Overlay Key	Click on the registry overlay > $(RegistryRoot) folder, then click the Add Key button.
Registry Overlay Value	Click on the registry overlay > $(RegistryRoot) folder, then click the Add Value button.
Database	Click on the data folder, then click the Add Database button.
Database Table	Click on a database in the data folder, then click the Add Table button.

Parameters

The Parameters primary folder contains a Directives subfolder that defines and displays details that are critical for locating and identifying the component, such as the file system or registry root, component version, vendor, and database connection information.

Discovery determines the file system root and registry root of a Component Blueprint and displays these parameters in the discovered service tree view. Special Version and NameQualifier parameters, if defined in the Component Blueprint, are displayed after the name in the discovered service tree view. The Component Blueprint parameter NameQualifier is defined as $(Product Name) SP$(Service Pack) and parameter Version is defined as $(RegistryRoot)\Windows NT\CurrentVersion\CurrentVersion.

Parameter directives can be used for variable substitution into diagnostics, utilities, and the definition of rules and other parameters. Wherever the string $(ParameterName) is entered as a value, the actual value of that parameter is substituted.

Adding Parameter Elements

To Add a...	Access and Procedure
Parameter Directive	Click on the directives folder, then click the Add Directive button.

Configuration

The Configuration primary folder defines how to find and interpret the component configuration information. You can find the configuration information in managed files or databases or you can derive it from the output of executables.

The Structure Classes folder defines how to interpret the configuration files, database queries, and executables. The information includes the semantics of each potential value in the configuration data set:

- Data typing
- Default values
- Enumerated values
- Qualifiers
- Categories
- Filters
- Weights
- Rules

Think of a structure class as the metadata.

The application locates, parses, and interprets for viewing and comparative analysis the files that the Configuration Files folder identifies.

The data folder identifies the queries or stored procedures to run to extract configuration data and how to interpret the results for comparative analysis.

The executables folder defines scripts and directives that can extract the configuration information from a server and how to interpret the results for comparative analysis.

Adding Configuration Elements

To Add a...	Access and Procedure
Structure Classes Class	Click on the structure classes folder, then click the Add Class button.
Structure Classes Parameter	Click on a class under the structure classes folder, then click the Add Parameter button.
File	Click on the files > $(Root) folder, then click the Add File button.
Structure Classes Group	Click on a class under the structure classes folder, then click the Add Group button. You can add additional groups and parameters to a group using the Add Group and Add Parameter buttons that display when you select a group. You can also make a copy of a group and all of its associated parameters using the Group Copy button.
Database	Click on the data folder, then click the Add Database button.
Database Query	Click on a database under the data folder, then click the Add Query button.
Executable Directive	Click on the executables folder, then click the Add Directive button.

Utilities

The Utilities primary folder contains executable files, macros, and scripts that can be used to perform common administrative tasks, for example, programs or scripts that start or stop the component or that provide additional information about a server or service, such as viewing system information, memory statistics, or disk volume statistics.

Adding Utility Elements

To Add a...	Access and Procedure
File	Click the file's $(Root) folder, then click the Add File button.
File Usage	Click a file in the file's $(Root) folder, then click the Add File Usage button.
Macro	Click on the macros folder, then click the Add Macro button.

Macro Step	Click on a macro under the macros folder, then click the Add Step button.

Diagnostics

The Diagnostics primary folder defines executable files and macros used for diagnosing, troubleshooting, and fixing component problems.

Any executable, script, or batch file that can run on a server can be defined here and made available to CA Configuration Automation users with the appropriate access control role. Macros provide a way to include frequently used scripts and troubleshooting tools that can help diagnose problems specific to the servers containing the data being managed by its Component Blueprint.

Adding Diagnostic Elements

To Add a...	Access and Procedure
File	Click on the files > $(Root) folder, then click the Add File button.
File Usage	Click on a file in the files > $(Root) folder, then click the Add File Usage button.
Macro	Click on the macros folder, then click the Add Macro button.
Macro Step	Click on a macro under the macros folder, then click the Add Step button.

Runtime

The Runtime primary folder defines component run-time files and helps you locate and view them quickly. The run-time files (for example, log files) typically change frequently. To exclude the file contents from comparative analysis, define the files in the Runtime primary folder.

Adding Runtime Elements

To Add a...	Access and Procedure
File	Click on the files > $(Root) folder, then click the Add File button.

Documentation

The Documentation primary folder defines how to find the component's managed documentation files (for example, readme files, PDF files, or HTML versions of the product manuals). The Documentation folder also lets you add explicit URL links to additional sources of component information or documentation, such as company intranets or vendor support sites.

Note: Access to referenced URLs is determined by your network configuration.

Adding Documentation Elements

To Add a...	Access and Procedure
File	Click on the files > $(Root) folder, then click the Add File button.
URL	Click on the URLs folder, then click the Add URL button.

Category Descriptions

The Category field lets you organize Component Blueprint elements.

Category	Description
Administration	Administrative settings that have to do with the general availability and management of the component. Examples of administrative settings would be how to back up, when to flush a cache, or how many times to retry something.
Configuration	Configuration settings for the component, except for those that are better described by Administration, Log and Debug, Network, or Performance. Examples of configuration settings would be a component alias name or default web page.
Documentation	Elements that document the component behavior or act as a guide to users, for example manuals, readme files, FAQs, or online help pages.
Log And Debug	Elements that have to do with setting up such things as log locations, log levels, debug output, or diagnostic variable types.
Network	Elements that represent or indicate a network-related setting for the component. An example would be the port settings for SNMP. If an element can be categorized as both Network and Security (for example, enabling LDAP Authentication), use Security as the category.
Other	CA internal use only.

Performance	Settings that are known to seriously impact performance and are generally a specific subset of configuration parameters. Examples of performance settings would be number of threads or number of concurrent users.
Product Info	General (and usually static) information about the product. Examples of static component information would be licensing, installation location, vendor, or module name.
Resources	Component resources. Examples of component resources would be storage, memory and cache allocation or size, or CPU. Note that this static resource category is different than the Transient category, which relates more to real-time information.
Security	Elements that represent security-related settings and are generally a specific subset of configuration parameters. Examples of security settings would be authentication types, enabling authentication, encryption settings, directory browsing, SSL, or HTTPS.
Transient	Elements that change with some regularity. Examples of transient elements would be server states (for example, up, down, running, or stopped), current number of connected clients, current number of threads, or current disk utilization.
Versioning and Patches	Elements that indicate the version or patch levels of the product.

Filter Descriptions

The Filter field lets you mark Component Blueprint elements for exclusion from key CA Configuration Automation operations and results like Change Detection and Rule Compliance.

Filter	Description
Component Specific	Identifies an element that is specific to a single component instance, for example, installation root and Service Server Name. The purpose is to identify and exclude certain component-specific elements, which are already known to be different, from Change Detection operations and results.
Service Specific	Identifies an element that is specific to a service, for example server names and installation roots. The purpose is to identify and exclude certain service-specific elements, which are already known to be different, from Change Detection operations and results.

Server Specific	Identifies an element that is specific to a single server, for example Server name and IP address. The purpose is to identify and exclude certain server-specific elements, which are already known to be different, from Change Detection operations and results.
Never Run Change Detection	Identifies an element that should be permanently excluded from all types of Change Detection operations and results. A temporary directory, log files in the managed folder, or anything that is known to be transient are examples of elements that you want to identify and always exclude from Change Detection operations and results.
Never Run Rule Compliance	Identifies an element that should be permanently excluded from all types of Rule Compliance operations and results due to their inconsequential or variable nature. A temporary directory, a known old configuration file, a template, or an example file are examples of elements you want to identify and always exclude from Rule Compliance operations and results.
Time Variant	Identifies an element that is known to change over time, but not necessarily across servers or across services, for example, log files, process start times, and registry event counters. The purpose is to identify and exclude parameters that are time variant, which are already known to be different, from Change Detection operations and results.

POSIX 1003.2-1992 Pattern Matching

POSIX pattern matching expressions are descriptions that help you find files and directories when the exact file or directory name is not known. CA Configuration Automation uses POSIX pattern matching expressions to search for files and directories during discovery and during refresh operations.

Syntax of POSIX Pattern Matching Expressions

File Name Matching Expression	**Description**
Characters	
unicodeChar	Matches any identical Unicode character.
?	Matches any single character.
*	Matches any string, including the null string (zero length).

Character Classes	
[abc]	Matches any character listed within the square brackets (simple character class).
[a-zA-Z]	Matches any range of characters listed within the square brackets (character class with ranges).
[^abc]	Matches any range of characters *except* those listed within the square brackets (negated character class).
Standard POSIX Character Classes	
[:alnum:]	Matches alphanumeric characters
[:alpha:]	Matches alphabetic characters
[:blank:]	Matches space and tab characters
[:cntrl:]	Matches control characters
[:digit:]	Matches numeric characters
[:graph:]	Matches characters that are printable and visible. (A space is printable but not visible, while an a is both.)
[:lower:]	Matches lowercase alphabetic characters
[:print:]	Matches printable characters (characters that are not control characters)
[:punct:]	Matches punctuation marks (characters that are not letters, digits, control characters, or space characters)
[:space:]	Matches characters that create space, such as tab, space, and formfeed
[:upper:]	Matches uppercase alphabetic characters
[:xdigit:]	Matches characters that are hexadecimal digits

Variable Substitution

Variable substitution allows the value of any element managed by CA Configuration Automation to be used as a parameter within a Component Blueprint directive. CA Configuration Automation defines an expression syntax to identify elements in the blueprint. Once identified, the element's value is extracted and used in place of the expression when the directive is executed. Substitution can be applied to any attribute of a directive. These include values, default values, paths, file names, parameters, environment variables, regular expressions, queries, and column names.

Managed elements whose values are addressable by variable substitution syntax include:

- Discovery parameters
- Registry variables
- Configuration values (file, database, or executable)
- File and directory attributes
- Managed data elements (schema metadata)
- Service and component attributes

Expression Types

Variable substitution expressions have the following forms:

Parameter Substitution

Provides access to the values of parameters that are defined in the current component.

Object Substitution

Lets you address the value of any element in any service.

Global Variable Substitution

Lets you address the values from the CA Configuration Automation global variable repository.

The following sections describe these substitution types.

Parameter Substitution

Parameter substitution expressions have the form:

`$(`*`VariableName`*`)`

VariableName

Identifies a discovery parameter that is defined in the current component. The name is case-sensitive, so it must match the parameter name exactly. You can embed the parameter expressions in string literals or you can use them standalone. You can define multiple substitutions in the same expression, and you define them recursively. For example, the substitution value can be a string in the parameter expression. If so, the product evaluates the substitution value recursively.

Example:

If you define the following Discovery parameters:

```
User=info
Domain=ca.com
v1=$(User)
v2=$(Domain)
```

the following parameter substitution expression:

```
$(User)@$(Domain) [$(v1) at $($(v2))]
```

returns the following result after evaluation:

```
info@ca.com [info at ca.com]
```

Object Substitution

Object substitution expressions define the path to an object in the CA Configuration Automation managed element tree. When an object is identified, the product returns the object value as the result of the expression. Alternatively, the product can return an attribute of the object, as the example in the Component-Scoped Object Substitution Expressions shows. If no object matches the expression, the product returns a null value.

You can define object substitution expressions in the scope of a service, in the current component, or globally.

Service-Scoped Object Substitution Expressions

A service-scoped object substitution expression must specify a component that can exist in the service. Service-scoped object expressions have the form:

```
${Component[ComponentName,ElementType[ElementName or Identifier,
...]]}
```

Braces { }

Distinguish the object syntax from the parameter expression syntax.

ComponentName

Defines either the name of a single component or a list of components that is delimited with the | character. The delimited list variant lets the object expression resolve a value when the service database contains multiple components.

Examples:

To use a delimited list variant where the component could be either SQL Server or Oracle:

```
${Component[Microsoft SQL Server|Oracle 8i
Server,Parameter[DatabaseUser]]}
```

To access the root parameter of a component:

```
${Component[CCA Server,Parameter[Root]]}
```

You can also select components by the Component Blueprint category:

```
${ComponentCategory[Relational Databases,Parameter[DatabaseUser]]}
```

The Component Blueprints page lists the valid category names:

- Application Platforms
- CA Software
- Clustering
- Compliance
- Custom Components
- Directory Servers
- Enterprise Applications
- Imported Components
- IT Management Systems
- Messaging Systems
- Network Devices
- Operating Systems
- Operating Systems - Limited
- Relational Databases
- Server Components
- Storage Managers
- Utilities
- Virtualization
- Web Servers

Component-Scoped Object Substitution Expressions

Component-scoped object expressions have the form:

```
${ElementType[ElementName or Identifier, ...]]}
```

Example:

```
${FileSet[$(Root),Directory[admin/logs,File[filter.log,Attribute[s
ize]]]]}
```

Globally Scoped Object Substitution Expressions

Globally scoped object expressions can access information from any service in a single CCA Database. Globally scoped object expressions have the form:

```
${Service[ServiceBlueprintName(ServiceName),Component[ … ]]}
```

Example:

To get the CA Configuration Automation mail from a configuration parameter from a service other than CA Configuration Automation:

```
${Service[CCA(MyCCA),Component[CCA Server,Configuration
[*,Files[*,Directory[lib,File[cca.properties,FileStructure[*,NVFil
e
[com.ca.mail.from]]]]]]]}
```

Elements and Attributes that are Available for Object Substitution Expressions

The following tree enumerates:

- The specific element types
- The relationships of specific element types in the tree
- The object values
- The available attributes (in parentheses)

You can use the strings in this example in an object substitution expression to build a path to an object in the tree.

```
Component [name or id]
   (module_id, mod_name, mod_desc, mod_version, platform_id
   mod_instance_type, mod_instance_of, release_version, mod_state,
   created_by, creation_time, server_id, server_name, domain_name,
   ip_address, mac_address, server_state,
   cc_agent_yn,cc_agent_port,
   cc_agent_protocol, os_type, os_version, processor, platform_name)
Parameter [parameter name]
Files [$(Root)]
   Directory [directory name or path (a/b/c)]
      (name, mtime, ctime, owner, perm, bytes, depth, files,
      directories)
   Directory ...
   File
   File [file name]
      (name, mtime, size, owner, perm, prodver, filever, ctime)
   Registry [*]
      RegKey [keyname or path (a\b\c)]
      (name, value)
   RegKey ...
   RegValue [name]
   (name, value)
   Configuration [*]
      Files [*]
       File [name]
        FileStructure
         GroupFileBlock [name]
         GroupFileBlock [name(value)] where value is the group
      block's
                  value, name qualifier, or name qualifier child
      value.
         GroupFileBlock ...
          NVFileBlock [name]
          NVFileBlock [name]
                  (description, view, weight, password, folder)
   Database [name]
      ResultSet [name]
```

```
        (name, type, query, queryType, description)
    DataRow [name]
    DataCell [name]
        (name, value)
DatabaseKey [name]
        (name, description, key, keyValues, column)
ExecutablesFileSystem [*]
    File [name]
    FileStructure
    GroupFileBlock [name] or GroupFileBlock [name(value)] where
    value
        is the group block's value, name qualifier, or name
    qualifier
        child value.
    GroupFileBlock ...
    NVFileBlock [name]
        (description, view, weight, password, folder)
Database [database name]
    DataBaseAccessSpec
        (server, user, password, driver, databaseName,
    databaseContext, env)
    Table [table name]
        (name, description, rowcount)
    Column [column name](name, description, length, nullable,
    default, ordinal, precision)
    Index [index name]
        (name, sort, unique, description)
    Column [column name]
        (name, description, length, nullable, default, ordinal,
    precision)
```

Global Variable Substitution

Global variable substitution expressions have the form:

$(*GlobalVariableName*)

GlobalVariableName

Defines a valid path in the CA Configuration Automation global variable repository. The name is not case-sensitive. You can embed the global variable expressions in strings or you can use them standalone. You can define multiple substitutions in the same expression, and you can define them recursively. For example, if the substitution value is a string in the parameter expression, the application evaluates the substitution value recursively.

Example:

For a global variable repository with the following structure:

```
Global Variables
   Site
      Phoenix
        Main: x4000
        Fire: x4911
      Tucson
        Main: x5000
        Fire: x5911
```

the following global variable substitution expression:

```
$(/Site/Tucson/Main)
```

returns the following result after evaluation:

```
x5000
```

Interpret As Descriptions

Interpret As provides a hint to CA Configuration Automation about a configuration parameter string format and how it is intended for use by the associated component. The application uses context-sensitive parsers to inspect interpreted parameter values, which lets the application extract multiple subvalues from complex parameter strings.

For example, if CA Configuration Automation can interpret and extract the following value as a JDBC URL, it can extract the database type, server, port, and database name:

```
jdbc:oracle:thin:@dbserver:1521:MYDBNAME
```

In addition to enabling context-sensitive parsing, interpretation also lets the application derive relationships. Using the extracted server in the example above, the application can establish a relationship between the current server and the server dbserver. To establish a relationship, use the Relationship Key.

The application allows only one interpretation for each value, and many values have no interpretation (leave such values uninterpreted). If more than one interpretation applies (for example, File Name and File Name or Path), use the one that most accurately describes the field. For example, if a field is defined as a file name (with no path), select File Name. If the field can be a file name, path, or partial path, select File Name or Path. The application includes the following Interpret As selections:

Database Name

The value is the name of a database in a database server.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Database Table

The value is the name of a database table in a database. The database table can contain a schema prefix.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Date

The value is a date in any format.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Date And Time

The value is a date that is combined with the time of day.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Description

The application interprets the value as descriptive text.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Directory Name

The value is only a directory with no path.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Directory Name or Path

The value is a directory name, path, or partial path.

Because the application can derive the Directory Reference relationships from this interpretation, you can set Relationship Key to Yes.

Email Address

The value is the destination for an email message. The interpreter looks for one or more email addresses in the value string.

File Name

The value is only a filename with no path.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

File Name or Path

The value is a file name, path, or partial path. Many named values allow any of these interpretations.

Because the application can derive the File Reference relationships from this interpretation, you can set Relationship Key to Yes.

Server Name or IP Address

The value is (or contains) an IP address or server name. Use this interpretation only when no port number is defined in the value. If the value contains a port number, use Server Name and Port.

The application can recognize server names and IP addresses that are embedded in larger strings.

Because the application can derive the Server Reference relationships from this interpretation, you can set Relationship Key to Yes.

Define a Server Reference relationship as a relationship key only if the referenced server is considered a dependency of the current server.

Server Name and Port

The value is (or contains) a server name or IP address and port number. A colon (:) must separate the server and port number.

The application can recognize server names, IP addresses, and port numbers that are embedded in larger strings.

Because the application can derive the Server Reference relationships from this interpretation, you can set Relationship Key to Yes.

Specify a Server Reference relationship as a relationship key only if the referenced server is considered a dependency of the current server.

Java Class Name

The value is a Java class name. It can be a class name, a package name, or a fully qualified class name with a package prefix.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

JDBC URL

The value defines a JDBC URL. The table shows the supported formats.

Because the application can derive the Server Reference relationships from this interpretation, you can set Relationship Key to Yes.

JDBC URLs almost always define an important relationship. You should typically identify them as Relationship Keys.

LDAP Path

The value defines a path to an LDAP subtree.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

LDAP Entry

The value is the name of an LDAP directory entry or the full path to an entry.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Network Domain

The value is a network domain (not including the server name). For example:

`ca.com.`

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Network Protocol

The value defines an IP protocol, such as TCP, UDP, FTP, SNMP, or SMTP.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Password

The value is a password.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Registry Key Name

The value is only a registry key name with no path.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Registry Key Path

The value is the full path (starting with \) to a registry key.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Registry Value Name

The value is only a registry value name with no path.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Registry Value Path

The value is the full path (starting with a \) to a registry value.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

SNMP Community String

The value specifies an SNMP community string. For example:

```
public
```

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

SNMP OID

The value specifies an SNMP object ID. For example:

```
1.3.6.1.4.1.18071.1.1.1
```

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

TCP Port Number

The value is a TCP port number (not UDP or unspecified).

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Time Interval

The value is an interval of time.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Time Of Day

The value is the time of day.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

UDP Port Number

The value is a UDP port number (not TCP or unspecified).

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

URL

The value specifies a URL, including the following protocols: file, http, https, ftp, jrmi, jmx:rmi, iiop, gopher, news, telnet, mailto, jnp, t3, and ldap.

The interpreter decomposes the URL and makes parts of it available through custom methods.

Because the application can derive the Server Reference relationships from this interpretation, you can set Relationship Key to Yes.

Define a URL relationship as a relationship key only if:

- The URL contains a server name
- The named server always defines a dependency between the current server and the named server.

User Group

The value is a user group.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

User Name

The value is a user name.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Version String

The value is any string that can be interpreted as a version.

Because no relationships are derived from this interpretation, always set the Relationship Key to **No**.

Web Service URL

The value specifies a URL that identifies a web service that a component uses.

Because the application can derive the Server Reference relationships from this interpretation, you can set Relationship Key to Yes.

The JDBC URL relationship interpretation supports the following formats:

Database Name	URL Pattern
SQL Server2005 SQL Server 2008	jdbc:sqlserver://<host>:<port>;databasename=<database>; SendStringParametersAsUnicode=false

Database Name	URL Pattern
Oracle 9, 10, and 11	■ jdbc:oracle:thin:@$(host):$(port):$(database) ■ jdbc:oracle:thin:@<host>:<port>:<database> ■ jdbc:oracle:thin:@<host>:<port>/<database> ■ jdbc:oracle:thin:@((DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=<host1>)(PORT=<port1>)(ADDRESS=(PROTOCOL=TCP)(HOST=<host2>)(PORT=<port2>))(FAILOVER=ON)(LOAD_BALANCE=OFF)(CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=<database>)) ■ jdbc:oracle:thin:@((DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=<host>)(PORT=<port>)))(CONNECT_DATA=(SERVICE_NAME=<database>))) ■ jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=<host>)(PORT=<port>)))(CONNECT_DATA=(SERVICE NAME=<database>))) ■ jdbc:bea:oracle://<host>:<port>
Informix	jdbc:informix-sqli://<host>:<port>/<database>:informixserver=<serverName
DB2	jdbc:db2://<host>:<port>/<database>
Sybase 11 and 15	jdbc:sybase:Tds:<host>:<port>/<database>
MYSQL	■ jdbc:mysql://<host>:<port>/<database> ■ jdbc:mysql://<host>:<port>
Postgres	jdbc:postgresql://<host>:<port>/<database>
HSQLDB	jdbc:hsqldb:hsql://<host>:<port>
ODBC	jdbc:odbc:<database>
Cloudscape	jdbc:cloudscape:<database>
Java DB (Derby)	■ jdbc:derby://<host>:<port>/<database> ■ jdbc:derby://<host>:<port>/<database>;create=true
Ingres	jdbc:ingres://<host>:<port>/<database>
Pointbase	jdbc:pointbase:server://<host>:<port>/<database>
Generic	jdbc:<xyz>:server://<host>:<port>/<database>

Support for multiple Relationships

Multiple relationships are created for the following URL patterns:

- If the JDBC URL Patterns contains multiple host/port combinations.

  For example:
  jdbc:oracle:thin:@((DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=<host1>)(PORT=<port1>)(ADDRESS=(PROTOCOL=TCP)(HOST=<host2>)(PORT=<port2>))(FAILOVER=ON)(LOAD_BALANCE=OFF)(CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=<database>))

  The following relationships are created:

  - Server "X" uses "host1" "port1" "database"
  - Server "X" uses "host2" "port2" "database"

- If JDBC URL Patterns are separated with comma (,) or pipe (|).

  For example:

  - jdbc:oracle:thin:@<host1>:<port1>:<database1>,
    jdbc:oracle:thin:@<host2>:<port2>:<database2>

    The following Relationships are created:

    - Server "X" uses "host1" "port1" "database1"
    - Server "X" uses "host2" "port2" "database2"
  - jdbc:oracle:thin:@<host1>:<port1>:<database1>|jdbc:oracle:thin:@<host2>:<port2>:<database2

    The following Relationships are created:

    - Server "X" uses "host1" "port1" "database1"
    - Server "X" uses "host2" "port2" "database2"

Regular Expressions

Regular expressions are pattern descriptions that enable sophisticated matching of strings. CA Configuration Automation uses regular expressions to:

- Search files for matching strings
- Verify that the string conforms to certain patterns, such as email addresses
- Extract string values from large blocks of text

If you need more information about the concepts behind regular expressions, there are many sources on the Web. For example, the Google directory about computer programming languages includes a helpful section about regular expressions: http://directory.google.com/Top/Computers/Programming/Languages/Regular_Expressions/FAQs,_Help,_and_Tutorials.

Regular Expression Syntax

The following table shows the supported syntax for regular expressions in CA Configuration Automation.

Regular Expression	Description
Characters	
unicodeChar	Matches any identical Unicode character.
\ (backslash)	Used to quote a meta-character or to process a special character as normal or literal text. For example, * makes the asterisk a normal text character, instead of a wildcard character.
\\	Matches a single \ character.
\0nnn	Matches a specified octal character.
\xhh	Matches a specified 8-bit hexadecimal character.
\uhhhh	Matches a specified 16-bit hexadecimal character.
\t	Matches an ASCII tab character.
\n	Matches an ASCII newline character.
\r	Matches an ASCII return character.
\f	Matches an ASCII form feed character.
Character Classes	
[abc]	Matches any character between the square brackets (simple character class).
[a-zA-Z]	Matches any range of characters between the square brackets (character class with ranges).
[^abc]	Matches any range of characters *except* those between the square brackets (negated character class).
Standard POSIX Character Classes	
[:alnum:]	Matches alphanumeric characters.
[:alpha:]	Matches alphabetic characters.
[:blank:]	Matches space and tab characters.

[:cntrl:]	Matches control characters.
[:digit:]	Matches numeric characters.
[:graph:]	Matches characters that are printable and visible. (A space is printable but not visible, but an *a* is both.)
[:lower:]	Matches lowercase alphabetic characters.
[:print:]	Matches printable characters (characters that are not control characters).
[:punct:]	Matches punctuation marks (characters that are not letters, digits, control characters, or space characters).
[:space:]	Matches characters that create space (for example, tab, space, and formfeed characters).
[:upper:]	Matches uppercase alphabetic characters.
[:xdigit:]	Matches characters that are hexadecimal digits.
Non-Standard POSIX-Style Character Classes	
[:javastart:]	Matches the start of a Java identifier.
[:javapart:]	Matches part of a Java identifier.
Predefined Classes	
. (period)	Matches any character other than newline.
\w	Matches a “word” character (alphanumeric plus “_”).
\W	Matches a non-word character.
\s	Matches a whitespace character.
\S	Matches a non-whitespace character.
\d	Matches a decimal digit.
\D	Matches a non-digit character.
Boundary Matches	
^ (caret)	Matches only at the beginning of a string.
$ (dollar sign)	Matches only at the end of a string.
\b	Matches any character that is at the beginning or end of a word boundary.
\B	Matches any character that is not at the beginning or end of a word boundary.

Greedy Closures (Also known as *quantifiers*. For more information, see the Note that follows this table.)	
A*	Matches A zero or more times.
A+	Matches A one or more times.
A?	Matches A one or zero times.
A{*n*}	Matches A exactly *n* times.
A{*n*,}	Matches A at least *n* times.
A{*n*,*m*}	Matches A at least *n* but not more than *m* times.
Reluctant Closures (Also known as *quantifiers*. For more information, see the Note that follows this table.)	
A*?	Matches A zero or more times.
A+?	Matches A one or more times.
A??	Matches A zero or one times.
Logical Operators	
AB	Matches A followed by B.
A|B	Matches either A or B.
(A)	Parentheses are used to group subexpressions.
Back References (Reaches back to what a preceding grouping operator matched and uses it again to match something.)	
\1	Back reference to first parenthesized subexpression match.
\2	Back reference to second parenthesized subexpression match.
\3	Back reference to third parenthesized subexpression match.
\4	Back reference to fourth parenthesized subexpression match.
\5	Back reference to fifth parenthesized subexpression match.
\6	Back reference to sixth parenthesized subexpression match.
\7	Back reference to seventh parenthesized subexpression match.
\8	Back reference to eighth parenthesized subexpression match.
\9	Back reference to ninth parenthesized subexpression match.

Note: All closure operators (+, *, ?, {*m*,*n*}) are "greedy" by default. That is, they match as many elements of the string as possible without causing the overall match to fail. To use a "reluctant" (non-greedy) closure, follow it with a ? (question mark).

Java Plug-ins Supplied with CA Configuration Automation

Optionally, Java plug-ins implementing the com.ca.catalyst.object.CCICatalystPlugin interface can filter directive values. You can develop the plug-ins and then add them to the CLASSPATH or use one of the following plug-ins that are supplied with CA Configuration Automation:

com.ca.catalyst.plugin.CCParameterRuleFilter(*pattern*)

Formats the Version parameter. This filter only recognizes and alters directives named Version.

For example, if the Version value initially extracted from a file is 530, specifying the CCParameterRuleFilter(#.#.#) plug-in translates the Version to 5.3.0. Specifying the CCParameterRuleFilter(#.##) translates the Version to 5.30.

CCMatch(*regex*)

CCMatchAnywhere(*regex*)

Returns "true" or "false" values by matching the specified regular expression with the directive value.

- If the regular expression exactly matches the entire value, CCMatch() returns "true."
- If the regular expression matches anywhere in the value, CCMatchAnywhere() returns "true."

The application interprets the regular expressions with DOTALL and MULTILINE mode enabled. DOTALL mode implies that the regular expression character '.' matches all characters, including end of line characters. MULTILINE mode implies that the regular expression characters '^' and '$' delimit lines, instead of the entire value beginning to end.

CCReplaceAll("*regex*","*replacement*")

CCReplaceFirst("*regex*","*replacement*")

Replaces the portions of the value that match the regular expression.

Surround the regular expression and the replacement with quotation marks, and separate them with a comma. To replace the regex value with a carriage return, use the \n special character in the replacement string. For example, CCReplaceAll(" ","\n") replaces all spaces with carriage returns.

CCToUpper

CCToLower

Converts the directive value to upper or lower case.

CCTrim

Removes leading and trailing spaces from the value.

CCExpression

Runs the specified expressions that contain the directive value.

Expressions are written in ECMA-script (JavaScript) and can contain any syntax that is valid in Version 2 of that language. Include the value of the current parameter as $(*VALUE*) in the expression. The parameter must have a value or the application does not call the plug-in.

The application allows variable substitution in the expression. For example:

CCExpression($(*VALUE*)*50)

Multiplies the value by 50.

CCExpression('$(*VALUE*)' == 'XXX')

Returns "true" or "false".

CCExpression(Math.sqrt($(*VALUE*)))

Takes the square root of the value.

CCExpression(function add(a,b){return a+b;};add($(*VALUE*),$(Other));)

Defines and calls functions.

Understanding and Using the Tabular Data Parser

Tabular data is any text that is formatted in rows and columns. Tabular data can also include embedded comments and one or more heading rows.

Examples of Tabular Data

- Output of the netstat command:

```
# netstat –ant

Active Internet connections (servers and established)
Proto Recv-Q Send-Q  Local Address          Foreign Address        State
tcp        0      0  0.0.0.0:512            0.0.0.0:*              LISTEN
tcp        0      0  0.0.0.0:32768          0.0.0.0:*              LISTEN
tcp        0      0  127.0.0.1:32769        0.0.0.0:*              LISTEN
tcp        0      0  0.0.0.0:513            0.0.0.0:*              LISTEN
tcp        0      0  0.0.0.0:2401           0.0.0.0:*              LISTEN
tcp        0      0  0.0.0.0:514            0.0.0.0:*              LISTEN
tcp        0      0  0.0.0.0:9188           0.0.0.0:*              LISTEN
tcp        0      0  127.0.0.1:8005         0.0.0.0:*              LISTEN
```

- Tab-separated output from an Excel spreadsheet:

```
Host        Address          Type          Owner
Bertha      192.168.123.12   Linux         Jerome
Factotum    192.168.123.33   Windows 2008  Bukowski
Terrapin    192.168.124.13   AIX           Hunter
```

CA Configuration Automation provides a Tabular Data Parser that interprets and parses any form of tabular data within configuration files or executables. In addition, you can specify parser options that control the layout of rows and columns within the tabular data set, assign names to columns, eliminate header and comment text, as well as organize the data hierarchically.

You can also use the Tabular Data Parser and specify parser options at the structure class-level, however file- and executable-level assignments take precedence over any assignments made at the structure class level.

Accessing and Using the Tabular Data Parser

To use the Tabular Data Parser, follow these steps:

1. In the class, file, or executable attribute sheet, select Tabular Data Parser from the Parser drop-down list.
2. Define the details of the tabular layout in the Parser Option(s). The following illustration displays the parser details:

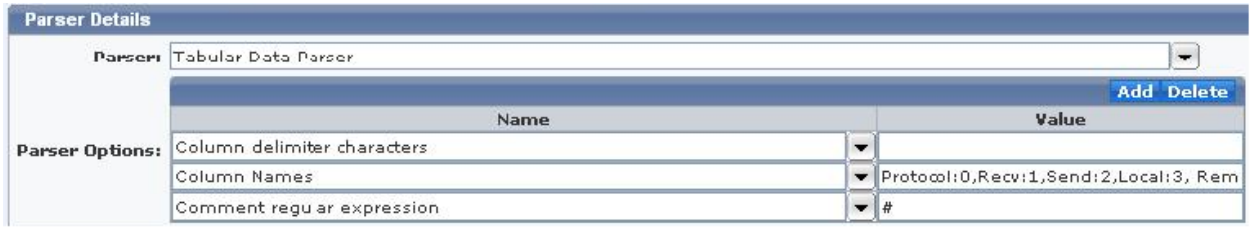

Parser Options are a set of attributes that are listed in a dropdown. Use the Add and Delete to add or delete an Option. Click the Name and Value option to edit the options.

For example, the Parser Options that are listed in the illustration are as follows:

- Column delimiter method
- Column Names
- Comment regular expression

3. Click save.

The Parser Option(s) field is updated.

Parser Options

You can set the following options for the Tabular Data Parser in the Parser Option(s) field.

Note: Enclose all values that you supply for an option in parentheses.

Column delimiter characters=

Defines one or more column delimiters.

If you do not define this option, the application defaults to the tab character. You can also specify more than one column delimiter.

For example, to specify the colon, slash, and comma as potential delimiters:

```
Column delimiter characters=:/,
```

For example, to specify the space, tab, or both as delimiters:

Space only

```
Column delimiter characters=" "
```

Tab only

```
Column delimiter characters="    "
```

Tab and space

```
Column delimiter characters="     "
```

Note: The quotation marks (" ") are only used for illustration. In actual practice, type only the space or tab character without enclosing quotation marks.

Column delimiter method=

Defines how to process consecutive delimiters. You can set this option to "one" or "all."

If you do not define this option, the application defaults to "all" and it processes consecutive delimiters as a single delimiter. For example, if you specify:

```
column delimiter characters=,
column delimiter method=all
```

for the following data:

```
ftp,tcp,udp,,,,xyz
```

the parser returns four columns: ftp, tcp, udp, xyz.

Set the value to "one" to process multiple consecutive spaces as one delimiter or to include data columns even if they have no defined values. For example, if you specify:

```
column delimiter characters=:
column delimiter method=one
```

for the following data:

```
root::0:XDCGBH!:
```

the parser returns the following columns:

```
root, "", 0, XDCGBH!, ""
```

Note: If you specified "column delimiter method=all in the previous example, the parser would return only the following columns:

```
root, 0, XDCGBH!
```

Header count=

Defines the number of lines to ignore at the beginning of the data set. For example, you can use:

```
Header count=2
```

to eliminate the header information from the tabular data results of netstat command.

You can only eliminate complete lines with the Header count= option. The parsed file does not include lines that you remove in the CA Configuration Automation UI.

Comment regular expression=

Defines a regular expression that identifies comments in the data set. For example, if you specify:

```
Comment regular expression=#.*
```

the parser interprets and ignores as comments patterns that start with # (including all other characters to the end of a line). For example:

```
#
# These three lines are removed
#
```

The parser also uses this option to interpret and ignore partial lines. For example:

```
some data   # this comment is also removed
```

The parsed file does not display lines and partial lines that you remove in the CA Configuration Automation UI.

Column Names=

Defines the assignment of names to individual columns of data. The format of the field is a comma-delimited list of names and column index numbers. Column indexing starts at zero. For example:

"Column Names"=Protocol:0,Recv:1,Send:2,Local:3,Remote:4,State:5

The parsed file does not display columns that are that you exclude from the Column Names= option in the CA Configuration Automation UI.

An important aspect of parsing tabular data to the standard CA Configuration Automation internal data format is the structuring of data into Structure Class groups. Groups allow you to assign each row of data a unique qualifier, nest them, and display them hierarchically on the user interface. The Tabular Data Parser uses the first name:index pair that is defined in the Column Names= option to name the group that contains the data rows (the *group pivot*).

With the Column Names= option in the previous example and the netstat command output as input, the application would display the data as follows:

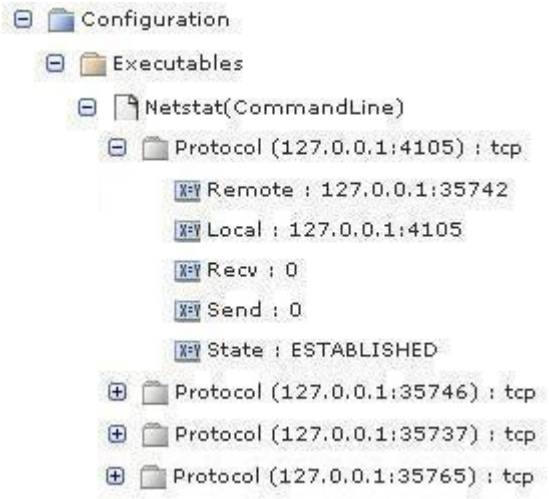

Note: The Structure Class used to interpret the parsed data in this example specifies Local as the qualifier for Protocol.

You can also group multiple columns to form subgroups under the top-level group pivot. For example, to nest the Recv and Send columns by one level in the hierarchy, apply the group modifier as follows:

"Column Names"=Protocol:0,Send/Recv(group):1-2,Local:3,Remote:4,State:5

The data would be displayed as:

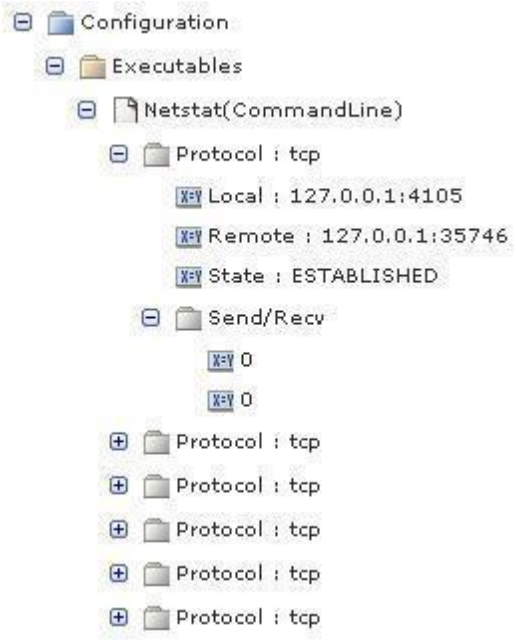

Note: In this example, the nested values are displayed under the named group and do not have names. The lack of names can limit the ability to write Structure Classes that accurately qualify the parent group. To define names for the nested columns, specify name:index pairs for the columns that are nested after the group modifier. For example:

```
"Column Names"=
Protocol:0,Send/Recv(group):1-2,Recv:1,Send:2,Local:3,Remote:4,
State:5
```

The application displays the data as:

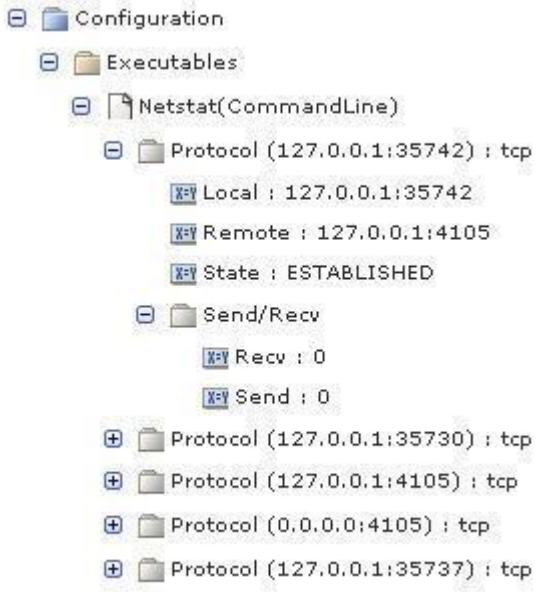

In the name:index pair specification, you can define groups as ranges of columns, lists of individual columns, or a combination. Valid group column formats include:

5-

Column 5 and all successive columns

3-5|7-

Columns 3 through 5 and 7 and all successive columns

3-5|7|9-11

Columns 3 through 5, 7, and 9 through 11

Note: You cannot specify the group option for the first name:index pair because it is the group pivot on which the application builds the hierarchy.

To substitute the value of a column for the name that a name:index pair specifies, use the valueasname modifier. You can specify the valueasname modifier on the group pivot (the first name:index pair). This is a convenient way to display tabular data because one column in a table is commonly a unique key. For example, if you applied the valueasname modifier as follows:

```
"Column
Names"=Protocol(valueasname):0,Send/Recv(group):1-2,Local:3,Rem
ote:4,State:5
```

The application displays the data:

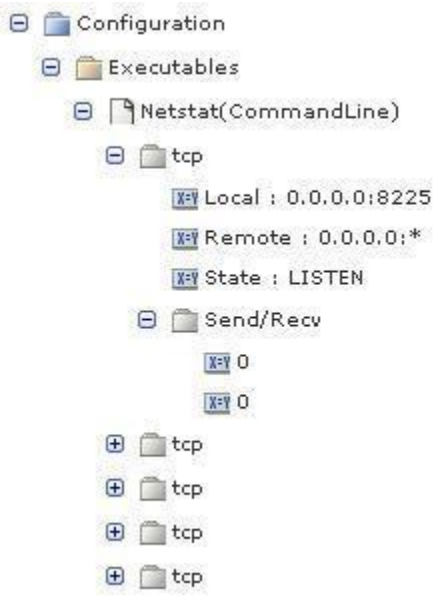

Line continuation regular expression=

Defines a regular expression to identify the line continuation syntax.

For example, use the line continuation character \ to continue a single row of data across multiple lines in the following file:

Name	Phone	Email	Address
Fred	9000	Fred@acme.com	12Jones Cl.
Ame	3002	Ame@acme.com	33535 Eucalyptus Terrace
Mary	7381	Mary@acme.com	31 Main Street

To parse the data correctly, use the following Line continuation regular expression= option:

```
Line continuation regular expression= \\$
```

Note: The first \ escapes the second \ to form a valid regular expression.

Chapter 11: Compliance Management

The Compliance tab page contains links to the following pages:

- Rule Groups
- Jobs
- History
- Log
- Reports

The sections that follow describe the functionality available on each of these pages.

Working with Rule Groups

The Rule Group page enables you to create and manage Rule Groups that organize rules in custom, logical groupings that make sense within your enterprise. Rule groups are primarily used for reporting purposes.

The Rule Groups tab page contains the Rule Groups table that lists the predefined and user-defined rule groups, and enables you to view, create, edit, import, export, and delete Rule Groups.

The Rule Group assignments are made from the following places:

- Servers page in the Run Rule Compliance wizard on the Servers tab page (accessed from the Select Actions drop-down list)
- Services page in the Run Rule Compliance wizard on the Services tab page (accessed from the Select Actions drop-down list)
- Management Options page in the Create a Management Profile wizard

For more information about the predefined rules, see the *Compliance Rules Reference Guide*.

Create Rule Groups

To create Rule Groups

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the Rule Groups link (below the main tabs)

   The Rule Groups page appears and lists all predefined and user-defined rule groups in the table.

3. Select Create Rule Groups from the Table Actions drop-down list.

   The Create Rule Groups wizard appears.

4. Enter the following information in the corresponding field, then click Next:

   Name

   Specifies the name of the Rule Group.

   Description

   Describes the function or purpose of this Rule Group.

   Documentation URL

   Specifies the location of the user-created documentation that describes the rule group.

   The Rules page appears.

5. Click Add Rule.

   The Category page of the Add Rule wizard appears.

6. Select a category from the Category drop-down list or enter a name for a new category, then click Next.

   The Filter page appears.

7. Click Refresh to display all Blueprints.

   The Blueprints appear in Available Blueprints column.

8. Double-click one or more Blueprints to move them to the Selected Blueprints column, then click Next.

   The Rules page appears.

9. Click the plus sign next to a Blueprint in the Rules tree, click the check box next to the rules you want to add to the Rule Group, then click Finish.

   The Rules table appears and displays the rules you selected in the previous step.

10. Click Finish.

    The new Rule Group appears in the Rule Group table.

Import Rule Groups

An import utility is provided to import Rules Groups into your CA Configuration Automation Server from another instance of CA Configuration Automation. A similar export utility is included to export Rule Groups as .jar files.

To import Rule Groups

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the Rule Groups link (below the main tabs)

   The Rule Groups page appears and lists all existing groups in the Rule Groups table.

3. Select Import Rule Groups from the Table Actions drop-down list.

   The Import Rule Groups dialog appears.

4. Click Browse and navigate to the .jar file you want to import.

   When the file was exported from a CA Configuration Automation Server, the file was assigned a default name using the following timestamp convention:

   ExportDatabaseObject_*YYYY_MM_DD_HH_MM_SS*.jar

   Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

   ExportDatabaseObject_2010_01_06_13_59_57.jar

5. Click the Overwrite Existing Rule Groups option if you want to update the existing version on your CA Configuration Automation Server, then click one of the following buttons:

   Import All

   Imports all Rule Groups in the .jar file.

   Import On Selection

   Displays the Available Imports table which lists the Rule Groups in the .jar file, and enables you to specify which of them you want to import.

   The import begins and may take a few moments depending on the size of the file. When the import is complete, a status message appears, the Rule Groups appears in the Rule Groups table, and the Modification Date/Time column is updated.

   Note: The Rule Groups table does not include a column that identifies the user who imported the file. You can click the Log link (below the main tabs) to display the log and identify the user who imported the file.

Export Rule Groups

An export utility is provided to export Rule Groups from one CA Configuration Automation implementation so it can be imported into another. The export utility extracts the Rule Groups, compresses, and saves the Rule Groups data to a .jar file.

To export Rule Groups

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the Rule Groups link (below the main tabs)

   The Rule Groups page appears and lists all existing groups in the Rule Groups table.

3. Click the check box next to one or more Rule Groups that you want to export, then select Rule Groups from the Select Actions drop-down list.

   A File Download window prompts you for the location where you want to save the export .jar file.

   The file is assigned a default name using the following timestamp convention:

   ExportDatabaseObject_*YYYY_MM_DD_HH_MM_SS*.jar

   Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

   ExportDatabaseObject_2010_01_06_13_59_57.jar

4. Click Save, specify the location, and then click OK.

   Note: If you are using a Windows operating system, Windows may suggest saving the file as a .zip file. Ensure the .jar is included in the file name, and the Save As Type field is set to All Files.

   The Rule Group .jar is saved in the specified location.

Delete Rule Groups

You can permanently delete Rule Groups if they are no longer needed.

To delete one or more Rule Groups

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the Rule Groups link (below the main tabs)

   The Rule Groups page appears and lists all existing groups in the Rule Groups table.

3. Click one or more check boxes next the Rule Groups that you want to delete, then select Delete Rule Groups from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   The selected Rule Groups are deleted.

Working with Compliance Jobs

The Compliance Jobs page enables you to create, schedule, view, run, and delete Compliance Jobs.

Create Compliance Jobs

You can create and schedule Compliance Jobs so they run automatically at regular intervals.

To create a Compliance Job

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the Jobs link (below the main tabs)

   The Compliance Jobs page appears and lists all existing Compliance Jobs.

3. Select Create Compliance Jobs from the Table Actions drop-down list.

   The Create Compliance Jobs wizard appears.

4. Enter a name and a description of the job in the corresponding fields, then select one of the following options from the Remediation drop-down list to specify whether you want to use remediation to reset non-compliant values:

   None

   Specifies that non-compliant values appear in the Rule Compliance results, but does not use remediation to reset these values.

   Rule Value Only

   Specifies that remediation is used to reset non-compliant values to the values that are defined in the rules.

   Rule Value or Blueprint Default Value

   Specifies that remediation is used to reset non-compliant values to the values that are defined in the rules. If an explicit rule is not defined for the component, the non-compliant value is reset to the default value defined in the Component Blueprint.

5. To run the remediation job manually from the Rule Compliance Results Page, deselect the Auto Remediate check box.

   Alternately, to run the remediation job automatically and view the results in the Rule Compliance Results Page, select the Auto Remediate check box.

   Note: By default, the Auto Remediate check box is not selected.

6. Click Next.

   The Rule Group page appears.

7. Double-click one or more entries in the Available Rule Groups field, or click the double right-facing arrow to select all of the Rule Groups.

   The selected Rule Groups appear in the Selected Rule Groups field.

8. Click Next.

   The Services page appears.

9. Double-click one or more entries in the Available Services field, or click the double right-facing arrow to select all of the services.

   The selected services appear in the Selected Services field.

10. Click Next.

    The Servers page appears.

11. Double-click one or more entries in the Available Servers field, or click the double right-facing arrow to select all of the servers.

    The selected servers appear in the Selected Servers field.

12. Click Next.

    The Server Groups page appears.

13. Double-click one or more entries in the Available Server Groups field, or click the double right-facing arrow to select all of the groups.

    The selected groups appear in the Selected Server Groups field.

14. Click Next.

    The Schedule page appears.

15. Define the schedule for automatically running the Rule Compliance operation by selecting one of the following from the Frequency drop-down list:

    Not Scheduled

    Specifies that the operation does not run automatically. It can be run manually or scheduled in the future.

    Once

    Specifies that the operation is run automatically one time. If you select this option, you also need to specify when it is run in the Time field.

    Minutes

    Specifies that the operation is run on a recurring basis using at an interval defined in minutes. If you select this option, you also need to specify the following:

    - Start Time—Specify the time the operation starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
    - Begin Date—Specify the date the operation is first run.
    - End Date—Specify the date the operation is run for the last time.
    - Recur every # minutes—Specify the interval at which the operation runs.

    For example, if you want the operation to run every 10 minutes starting at 11:00 p.m., you would specify a Start Time of 11:00:00PM, and specify Recur every 10 minutes. The operation would run at 11:00 p.m., 11:10 p.m., 11:20 p.m., 11:30 p.m., and so on until the end of the hour (midnight in this example). If the current operation has not finished running by the time the next interval occurs, the next run waits until the previous one completes, and then starts.

    Hourly

    Specifies that the operation is run on a recurring basis using at an interval defined in hours. If you select this option, you also need to specify the following:

    - Start Time—Specify the time the operation starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
    - Begin Date—Specify the date the operation is first run.
    - End Date—Specify the date the operation is run for the last time.
    - Recur every # hours—Specify the interval at which the operation runs.

For example, if you want the operation to run every four hours throughout the day starting at 11:00 p.m., you would specify a Start Time of 11:00:00PM, and specify Recur every 4 hours. The operation would run at 11:00 p.m., 3:00 a.m., 7:00 a.m., 11:00 a.m., 3:00 p.m., and 7:00 p.m.. If the current operation has not finished running by the time the next interval occurs, the next run will wait until the previous one completes, and then start. Also note that if the Start Time has already passed in the current day, the operation runs immediately, then resumes the recurring schedule you specify.

Daily

Specifies that the operation is run on a recurring basis using at an interval defined in days. If you select this option, you also need to specify the following:

- Start Time—Specify the time the operation starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
- Begin Date—Specify the date the profile is first run.
- End Date—Specify the date the operation is run for the last time.
- Recur every # days—Specify the interval at which the operation runs.

Weekly

Specifies that the operation is run on a recurring basis using at an interval defined in weeks. If you select this option, you also need to specify the following:

- Start Time—Specify the time the operation starts to run. Start time is always on the hour (for example, 10:00:00PM, 8:00:00AM, and so on).
- Begin Date—Specify the date the operation is first run.
- End Date—Specify the date the operation is run for the last time.
- Recur every # weeks—Specify the interval at which the operation runs.

Monthly

Specifies that the operation is run on a recurring basis using at an interval defined in months. If you select this option, you also need to specify the following:

- Start Time—Specify the time the operation starts to run.
- Begin Date—Specify the date the operation is first run.
- End Date—Specify the date the operation is run for the last time.
- Recur every # months—Specify the interval at which the operation runs.

16. Define the notification that is sent when the operation is run in the following fields:

    Notification Profile

    Specifies the notification profile to use when this operation is run as scheduled. For information about creating notification profiles, see Create Notification Profiles (see page 98).

    Subject

    Specifies the subject line of the email message that is sent by the selected notification profile.

    The schedule for the Rule Compliance Job is defined.

17. Click Finish.

    The job is created and appears in the Compliance Jobs table.

Run Compliance Jobs

You can run any existing Compliance Job manually if you need an immediate compliance update.

To manually run a Compliance Job

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the Jobs link (below the main tabs)

   The Compliance Jobs page appears and lists all existing Compliance Jobs.

3. Click the check box next to the job you want to run, then select Run Compliance Jobs from the Select Actions drop-down list.

   The following message appears: Compliance Jobs started successfully.

   After the job completes, the Rule Compliance results appear.

Delete Compliance Jobs

You can delete any existing Compliance Job if you no longer need it.

To delete a Compliance Job

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the Jobs link (below the main tabs)

   The Compliance Jobs page appears and lists all existing Compliance Jobs.

3. Click the check box next to the jobs you want to delete, then select Delete Compliance Jobs from the Select Actions drop-down list.

   The selected Compliance Jobs are deleted.

Run Compliance Job using Live Browse

To automate the discovery and compliance processes as an administrator, you can:

- Browse the objects (File System, Windows Registry, Windows Services, and Group Policies) in a compliant server.
- Create compliance rules for the selected objects.
- Run the compliance job for the selected service, servers, or server groups.
- Run the remediation job to remediate non-compliant objects on services, servers, or server groups using the Remediation Job tab.

When a compliance job is created using the Live Browse option, the following objects are created:

- Rule Group: Assume that the rule group is created with the name Rule Group 1.
- Blueprint: The new blueprint name is created as <Rule Group 1>_<Milliseconds since January 1, 1970, 00:00:00 GMT>
- File Structure Classes: The new file structure classes are created as '<Rule Group 1>_<Milliseconds since January 1, 1970, 00:00:00 GMT> (Services)' and '<Rule Group 1>_<Milliseconds since January 1, 1970, 00:00:00 GMT> (Group Policies)'

Note: You can filter the live browse objects using the Is Live Browse option from the respective Rule Group, Blueprint, or File Structure Classes tables.

Create a Rule Group using Live Browse

You can browse and create the compliance rules for the objects in the server. After you create the rules, you can run compliance jobs on the selected servers and can remediate the failed rules, if necessary.

Follow these steps::

1. Click Tasks, Run Compliance Job, then click Next.
2. Click Create New Compliance Job and then click Next.
3. Enter the required information, then click Create New Rule Group with Live Browse.
4. Enter the required information, then click Next.
5. Click the Browse button and select the server from Server drop-down list in the Browse Server dialog.
6. Select the required objects, click Add, click Close, then click Save.

   The Rules tab displays for the selected object.
7. On the Rules tab, select the Create Rule option from the Table Actions drop-down list.
8. Complete the information that is required to create the rules for the object, then click Save.

   Note: See Create Rules using Scripts (see page 326) for more information.
9. Select the services, servers, and server groups on which you want to apply the rules.
10. Follow the wizard and complete one of the following actions:
    a. Select Now to run the compliance job immediately
    b. Select a scheduling option.
11. To view the results, click Log and select the Rule Compliance tab.

Remediate Failed Rules for Compliance Jobs

To remediate failed rules for the compliance job, use the Remediation tab. The remediation job runs for the included steps and updates the values that were provided during the rule creation.

Follow these steps:

1. Select the required server, then select the remediation steps to exclude.
2. From the Select Actions drop-down list, select the Delete Steps option.
3. Click Run Remediation Job.

Create Rules using Scripts

You can use the Custom Scripts option to define scripts for Rule Compliance, Remediation, and Undo Remediation. You can include parameters, global and context variables (for example, DiscoveredPath and OSSERVICENAME) in scripts and can replace the variable values at run time.

Follow these steps:

1. Follow the steps 1 through 6 in Create a Rule Group using Live Browse (see page 325).
2. From the Constraint Type drop-down list, select the Custom Script option.
3. Click Editor and add a custom script or update an existing script.

   Similarly, you can use the editor in the Remediation section to add or update remediation scripts.

Example: Create Rule for Group Policy

Assume a requirement to set a Minimum Password Length (for example, 4) that complies with to the Security Settings policy for a server.

Follow these steps:

1. Complete steps 1 through 5 in the Create a Rule Group using Live Browse (see page 325) section.
2. From the Group Policies tab, select Minimum Password Length under Security Settings.
3. Click Add, click Close, and click Save.

   The Rules tab displays for Minimum Password Length.
4. From the Table Actions drop-down list, select Create Rule.
5. Enter a name and value (for example, '4'), then click Next.
6. Select the servers on which to apply the rules.
7. Follow the wizard, select Now to run the compliance job immediately, then click Finish.

   The compliance job runs for the selected servers and displays a success message.
8. Click the success message, select the Tree or Flat Table tab, and view the results.

Modify Rule Groups (Live Browse)

You can modify the following information for a Rule Group:

- Registry
- File System
- Windows Services
- Group Policy
- Parameters

Follow these steps:

1. Click Management, Compliance, and Rule Groups.

   The Rule Groups page displays all predefined and user-defined rule groups.
2. Select a rule group from the list of rule groups.
3. Select the Live Browse tab.

   Note: The Live Browse tab displays when you create the rule group using the live browse option. However, if the rule group is not created using Live Browse, then make the necessary modifications in the Details for Rule Group page.
4. Click the Browse button and select the server from Server drop-down list in the Browse Server dialog.
5. To view the list of objects, click the Browse button.
6. Click + to navigate and select the required object details.
7. (Optional) Click View to see the object information.
8. Click Add.

   A confirmation message appears.
9. Click Close and click Save.
10. Select the Rules tab.
11. Select the Create Rule option from the Table Actions drop-down list.
12. To create the rules for the object, complete the required information and click Save.
13. Save the Rule Group details.
14. (Optional) Repeat step 4 to step 13 to add more objects and rule for the object.

    The updated Rule Group displays in the Rule Group table.

Test Live Browse

For the rule groups created using the Live Browse option, you can verify the objects added for compliance for a server.

Follow these steps::

1. Click Management, Compliance tab, Rule Groups.

   A list of predefined and user-defined rule groups displays.

2. Select a rule group and select the Live Browse tab.
3. Click Test Live Browse.
4. Select a required server and click Test Live Browse.

   The objects that are added for compliance are displayed.

Delete Live Browse Objects

You can delete the live browse objects if you no longer need it.

Follow these steps:

1. Delete Compliance Jobs (see page 324).
2. Delete Rule Group (see page 318).
3. Delete Blueprint references using Blueprint Search option from Blueprints table (see page 262).
4. Delete Blueprint (see page 262).
5. Delete Structure Classes (see page 272).

Working with Compliance History

The Compliance History page enables you to view and delete compliance history entries.

View Compliance History

You can view the history of Compliance Jobs on the Compliance History page.

To view Compliance History

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the History link (below the main tabs)

   The Compliance History page appears and lists all existing history entries in the Compliance History table.

Delete Compliance History

You can permanently delete Compliance History entries if they are no longer needed.

To delete Compliance History entries

1. Click the Management link, then click the Compliance tab.

   The Compliance tab page appears.

2. Click the History link (below the main tabs)

   The Compliance History page appears and lists all existing history entries in the Compliance History table.

3. Click one or more check boxes next the entry that you want to delete, then select Delete History from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   The selected entries are deleted.

Working with Compliance Reports

The Compliance, Reports page contains two versions of a rule compliance report that you can run on either servers or services.

- The Rule Compliance Server report summarizes rule compliance results by component, category, and weight, and then presents the details of each rule violation.
- The Rule Compliance Service report summarizes rule compliance results by server, component, category, and weight, and then presents the details of each rule violation.

For information about scheduling or running the reports, see Report Management (see page 359).

Working with Compliance Exceptions

The Rule Exceptions page lets you perform the following tasks:

- View and manage the exceptions that are associated with Servers and Services.
- Select the rules to add the rule as exceptions on servers and services.
- Delete the exceptions that associated with Servers and Services for the selected rules.
- Select the rules to modify the server or service list.

Add Exceptions

Add rules as exceptions for servers and services.

Follow these steps:

1. Click the Management link, then click the Compliance tab.
2. Click the Rule Exceptions link in the Compliance tab page.
3. The Rule Exceptions page displays the already configured or available exceptions that are associated with servers and services.
4. Select Add Exceptions from the Table Actions drop-down list and do the following:
    a. In the Add Exceptions Rules page, double-click one or more blueprints from the Available Blueprints columns to move them to the Selected Blueprints column, then click Next.
    b. In the Select Structure Classes page, double-click one or more structure classes from the Available classes columns to move them to the Selected classes column, then click Next.

c. In the Select Rules page, select one or more rules that are related to a blueprint or FSC to add an exception, then click Next.

d. In the Servers page, select one or more servers to add rules as an exception, then click Next.

e. In the Services page, select the services to add as an exception.

f. Click Finish.

Rules are added as exceptions for the selected or the servers, and services.

Manage Exceptions

You can delete and modify the servers and services list that are defined on a rule.

Follow these steps:

1. Click the Management link, then click the Compliance tab.
2. In the Compliance tab page, click the Rule Exceptions link.

   The Rule Exceptions page displays the available or already configured exceptions that are associated with servers and services.
3. In the Rule Exception page, modify the existing exceptions for the selected rule in the following way:

   a. Select one or more rules and then select Manage Exceptions from the Select Actions drop-down list.

   b. Select the *Include in Exception*, or *Exclude from Exception* option to add, or remove the servers or service.

   c. (Modify a single rule) Click a rule to add or remove the servers or services from selected servers or services tab.
4. Click Save.

Delete Exceptions

You can select and delete one or more rule exceptions that are no longer used.

Follow these steps:

1. Click the Management link, then click the Compliance tab.
2. In the Compliance tab page, click the Rule Exceptions link.
3. The Rule Exceptions page displays the available or already configured exceptions that are associated with servers and services.
4. Select one or more rules, then select Delete Exceptions from the Select Actions drop-down list to delete the rules.

5. Confirm the deletion and click Ok.

   The exceptions that are associated with servers and services for the select rules are deleted.

Chapter 12: Remediation Management

CA Configuration Automation enables you to change server or service component attributes by running a *Remediation Job*. You can run a Remediation Job manually at any time, or you can schedule a Remediation Job.

The following lists describes how and where Remediation operations can be invoked:

- Scheduled or run manually using a Remediation Profile and Job
- Ad hoc using the View Components and Configurations window
- Ad hoc using Change Detection or Compare
- Ad hoc using Rule Compliance from servers or services management actions
- Ad hoc using a compliance Job
- Automatically using a Compliance Job
- Automatically using Management Profile Rule Compliance

Create Remediation Profiles

CA Configuration Automation enables you to use Remediation to change server or service component attributes by running a Remediation Job or making the change immediately. Remediation changes can be made to the following components:

- Registry keys and values located in the Managed folder.
- Parameter values located in the Configuration folder.
- Values returned by a *macro step* being executed. A macro step is a special set of Component Blueprint directives that perform utility and diagnostic operations. Each macro step returns a status and a value that is either extracted directly from the tree view of a server, component, or service or by Agent interrogation.

Each attribute change added to a Remediation Profile is considered a step. Remediation Profiles can include multiple steps. When the profile is specified by a Profile Job, the steps can be performed on one server, one step performed on multiple servers, or multiple steps performed on multiple servers.

Additionally, macro steps can be added to a Remediation Profile from elements contained in the Component Blueprints, Macro folder. When you click an element in the Macro folder, the Macro attribute sheet is displayed. Macros can be used to perform steps like stopping a service before making Remediation changes, and then restarting it after the changes are made. For example, WebLogic caches its config.xml file when it is running, and then writes the cached file to disk when it is shut down. If you make changes to the config.xml file using CA Configuration Automation without stopping WebLogic, your changes will be overwritten when WebLogic stopped and writes the cached file to disk.

To avoid this problem, you can create a Remediation job that includes a macro that stops WebLogic as the first step, makes the desired changes in subsequent steps, and then runs another macro that starts WebLogic as the final step.

Note: Some configuration files cannot be modified using Remediation (for example, listener.ora). To disable Remediation on these files, an Allow Remediation Jobs option has been added to the Class attribute sheet. By default, these files have this option set to No to prevent Remediation changes. To check if a file can be modified using Remediation, go to Blueprints, Structure Classes, select a file, then select Edit/View Structure Class to display the Class attribute sheet. If the Allow Remediation Jobs option is set to Yes, you can modify the file using Remediation.

The following section describes how to perform a simple, immediate attribute change using Remediation manually. The section that follows it describes creating a multi-step Remediation job that can include macro and Remediation job steps, and can be scheduled and managed as a job.

To create a Remediation Profile

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the Profile link (below the main tabs)

   The Remediation Profile page appears and lists all existing profiles in the Remediation Profiles table.

3. Select Create Remediation Profile from the Table Actions drop-down list.

   The Profile page of the Create Remediation Profile wizard appears.

4. Enter the following information in the corresponding field, then click Next:

   Name

   Specifies the name of the profile

   Description

   Describes the function or purpose of the profile.

   The Steps page appears.

5. Select Create Step from the Table Actions drop-down list.

   The software components known to CA Configuration Automation are listed in the Components pane.

6. Navigate the components tree and click the element whose value you want to remediate.

   The Details pane displays the selected element in the Name field, and Update is displayed in the Action field.

7. Enter the appropriate information in the corresponding field, then click Add Step:

   New Value

   Specifies the value that is assigned to the selected element.

   Change Description

   Describes the change being made by the Remediation Profile.

   Fail if Expected Value Not Equal to Actual Value

   Specifies whether the change is allowed to fail instead of completing when the differences in the values affect the current change.

   This is important because the value stored in the CA Configuration Automation Database may not be the actual value on the server if the value has changed and the component has not been refreshed since the change. Having the operation fail lets you investigate and evaluate the differences before re-running the Remediation Job.

   Stop If Step Failed On Target Server

   Specifies whether the Remediation Job associated with this profile stops if this step fails.

   A message confirms that the step was created.

8. Click the Servers link (above the Details fields).

   A pie chart appears above the Summary table.

   - The pie chart shows the total number of servers where the component is present. It divides the total into segments where the selected configuration setting has the same value. You can mouse-over a segment to display the value associated with the represented servers.
   - The Summary table lists the colors that represent the servers in the pie chart, the value of the selected setting, and the number of servers where the component is present.

9. Do one of the following:

   - Repeat steps 6 through 8 if you want to add additional steps.
   - Click Done Adding Steps.

   The newly-created step appears in the Steps table.

10. (Optional) Reorder the steps using one of the following options:
    - Click the check box next to the step you want to move, then select Move Step Up or Move Step Down from the Select actions drop-down list.
    - Click the check box next to the step you want to move, then click the up- or down-facing arrow to the right of the table to move the step up or down.
11. Click Finish.

    The profile appears in the Remediation Profiles table.

View and Edit Remediation Profiles

You can view and edit the details of a Remediation Profile.

To view and edit Remediation Profiles

1. Click the Management Link, and then the Remediation tab.

   The Remediation tab page appears.
2. Click the Profiles link (under the main tabs).

   The Remediation Profiles page appears and displays the existing profiles in the Remediation Profiles table.
3. (Optional) Click a profile name in the Profiles Name column.

   The Remediation Profile Details page appears. You can edit the details on this page. The fields are described in Create Remediation Profiles (see page 333).

Import Remediation Profiles

You can import a Remediation Profile as a Java Archive (JAR) file from another CA Configuration Automation instance.

Follow these steps:

1. Click the Management link, then click the Remediation tab.
2. On the Remediation tab, click the Profiles link.
3. On the Remediation Profiles page, click Table Actions, then select Import Remediation Profiles.

4. On the Import Remediation Profiles dialog, complete the following fields:

   JAR File to Import

   Defines the name of the JAR file that contains the Remediation Profile to import. Click Browse to navigate to the file.

   Overwrite Existing Remediation Profiles

   Specifies whether to overwrite a file with the same name. Select this option to retain the profile from another CA Configuration Automation instance.

5. Click one of the following buttons:

   Import All

   Imports all of the Remediation Profiles in the JAR file.

   Import On Selected

   Displays a dialog on which to select specific Remediation Profiles to import from the JAR file.

   The application imports the file and the Remediation Profiles table displays the profiles.

Delete Remediation Profiles

You can permanently delete Remediation Profiles if they are no longer needed.

To delete Remediation Profiles

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the Profile link (below the main tabs)

   The Profiles page appears and lists all existing profiles in the Remediation Profiles table.

3. Click one or more check boxes next the profiles that you want to delete, then select Delete Profiles from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   The selected profiles are deleted.

Run Ad Hoc Remediation Jobs from the Component List

While many CA Configuration Automation operations require you to specify a service or server on which to run an operation, an Ad Hoc Remediation Job does not show components in the context of a particular server or service. Instead, it uses a component-centric approach the lists all components currently on at least one server in your enterprise. From this list, you can navigate to a particular configuration parameter and view its value on each server where it appears. From there you can run or schedule a job to remediate one or more of these values.

To run an Ad Hoc Remediation Job

1. Click the Management link, then the Remediation tab.

   The Remediation tab page appears.

2. Click the Ad Hoc link (below the main tabs).

   The Ad Hoc Remediation page appears with the software components listed in the left pane.

3. Click the plus sign next to the component to navigate to the file or registry entry you want to remediate.

   The Details page appears in the right pane.

4. Enter the following information in the appropriate field:

   Action

   Specifies the type of remediation action.

   New Value

   Specifies the value to replace the unwanted value. If you want to see the current values before entering a new value, you can click the Servers link above the Details fields and view the current value on each server where the selected component exists (as described in step 5).

   Change Description

   Describes the purpose of the Remediation Job. This description appears in the Job Description column of the Remediation Jobs table after the jobs is created.

   Fail If Expected Value Not Equal to Actual Value

   Specifies whether you want this job to fail instead of completing when the differences in the values affects the current change you are making.

   This is important because the value stored in the CA Configuration Automation Database may not be the actual value on the server if the value has changed and the component has not been refreshed since the change. Having the operation fail lets you investigate and evaluate the differences before re-running the Remediation Job.

Stop If Pre Macro Execution Failed

Specifies that the job stops if the Pre Macro fails to perform its function. For example, if the Pre Macro was created to stop a service before a remediation change could be made, and it failed, this job would fail.

Pre Macro

Specifies the macro step to run before the remediation occurs. A common example is to create a Pre Macro that stops a service.

Post Macro

Specifies the macro step to run after the remediation occurs. A common example is to create a Post Macro that restarts a service.

5. Click the Servers link (above the Details fields).

   A pie chart appears above the Summary table.

   - The pie chart shows the total number of servers where the component is present. It divides the total into segments where the selected configuration setting has the same value. You can mouse-over a segment to display the value associated with the represented servers.
   - The Summary table lists the colors that represents the servers in the pie chart, the value of the selected setting, and the number of servers where the component is present.

6. Do one or both of the following in the Summary table:

   - Click the check box next to one or more rows in the table that you want to replace the current value with the value you entered in the New Value field (in step 4).
   - Click the link in the Value column to display the Servers table, then select one or more individual servers on which to replace the current value with the value you entered in the New Value field (in step 4).

7. (Optional) Click the Selected Servers link.

   The Selected Servers table appears and displays a summary of the servers selected in step 6.

8. Click Remediate.

   The Remediate dialog appears.

9. Select one of the following options from the Run drop-down list:

   Now

   Specifies that the job starts when you click OK.

   Later

   Displays the Time field where you can specify the date and time to start the job.

10. Select one of the following options from the Notification drop-down list, then enter the Subject line for the notification if you want to override the subject defined in the selected profile.:

    Use Default

    Specifies that the Notification Profile that is designated as the Default is used for this job.

    User-defined profiles

    Specifies the user-defined Notification to use for this job.

11. Click OK.

    One of the following happens depending on your selection in step 9:

    - The job starts. When it is complete, a notification is sent, and the job history appears in the Remediation History table (click the History link on the Remediation tab page to view it).
    - The job is scheduled and appears in the Remediation Jobs table (click the Jobs link on the Remediation tab page to view it).

Create Profile Jobs

You can create *Profile Jobs* that use existing Remediation Profiles to make changes to the following components:

- Registry keys and values located in the Managed folder.
- Parameter values located in the Configuration folder.
- Values returned by a *macro step* being executed. A macro step is a special set of Component Blueprint directives that perform utility and diagnostic operations. Each macro step returns a status and a value that is either extracted directly from the tree view of a server, component, or service or by Agent interrogation.

The Remediation Profile defines what component is changed by the job, and what the new value is. The Profile Job specifies what server, server group, or service is updated with the new value, and when the job is runs.

To create a Profile Job

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the Jobs link (below the main tabs)

   The Remediation Jobs page appears and lists all existing jobsin the Remediation Jobs table.

3. Select Create Profile Job from the Table Actions drop-down list.

   The Job page of the Create Profile Job wizard appears.

4. Enter the following information in the corresponding field, then click Next:

   Name

   Specifies the name of the profile

   Description

   Describes the function or purpose of the profile.

   Profile

   Specifies the Remediation Profile to use for this job.

   Stop If Step Failed On Target Server

   Specifies whether the Remediation Job associated with this profile stops if this steps fails.

   The Services page appears.

5. Double-click one or more services you want to be updated by the job in the Available Services column to move them to the Selected Services column, and then click Next.

   The Servers page appears.

6. Double-click one or more servers you want to be updated by the job in the Available Servers column to move them to the Selected Servers column, and then click Next.

   The Server Groups page appears.

7. Double-click one or more server groups you want to be updated by the job in the Available Server Groups column to move them to the Selected Server Groups column, and then click Next.

   The Schedule page appears.

8. Define the schedule for automatically running the job by selecting one of the following from the Frequency drop-down list:

   Not Scheduled

   Specifies that the job does not run automatically. It can be run manually or scheduled in the future.

   Once

   Specifies that the job is run automatically one time. If you select this option, you also need to specify when it is run in the Time field.

Minutes

Specifies that the job is run on a recurring basis using at an interval defined in minutes. If you select this option, you also need to specify the following:

- Start Time—Specify the time the job starts to run.
- Begin Date—Specify the date the job is first run.
- End Date—Specify the date the job is run for the last time.
- Recur every # minutes—Specify the interval at which the job runs.

For example, if you want the job to run every 10 minutes starting at 11:00 p.m., you would specify a Start Time of 11:00:00PM, and specify Recur every 10 minutes. The job would run at 11:00 p.m., 11:10 p.m., 11:20 p.m., 11:30 p.m., and so on until the end of the hour (midnight in this example). If the current job has not finished running by the time the next interval occurs, the next run waits until the previous one completes, and then starts.

Hourly

Specifies that the job is run on a recurring basis using at an interval defined in hours. If you select this option, you also need to specify the following:

- Start Time—Specify the time the job starts to run.
- Begin Date—Specify the date the job is first run.
- End Date—Specify the date the job is run for the last time.
- Recur every # hours—Specify the interval at which the job runs.

For example, if you want the job to run every four hours throughout the day starting at 11:00 p.m., you would specify a Start Time of 11:00:00PM, and specify Recur every 4 hours. The job would run at 11:00 p.m., 3:00 a.m., 7:00 a.m., 11:00 a.m., 3:00 p.m., and 7:00 p.m.. If the current job has not finished running by the time the next interval occurs, the next run will wait until the previous one completes, and then start. Also note that if the Start Time has already passed in the current day, the job runs immediately, then resumes the recurring schedule you specify.

Daily

Specifies that the job is run on a recurring basis using at an interval defined in days. If you select this option, you also need to specify the following:

- Start Time—Specify the time the job starts to run.
- Begin Date—Specify the date the job is first run.
- End Date—Specify the date the job is run for the last time.
- Recur every # days—Specify the interval at which the job runs.

Weekly

Specifies that the job is run on a recurring basis using at an interval defined in weeks. If you select this option, you also need to specify the following:

- Start Time—Specify the time the job starts to run.
- Begin Date—Specify the date the job is first run.
- End Date—Specify the date the job is run for the last time.
- Recur every # weeks—Specify the interval at which the job runs.

Monthly

Specifies that the job is run on a recurring basis using at an interval defined in months. If you select this option, you also need to specify the following:

- Start Time—Specify the time the job starts to run.
- Begin Date—Specify the date the job is first run.
- End Date—Specify the date the job is run for the last time.
- Recur every # months—Specify the interval at which the job runs.

9. Define the notification that is used when the job is run using the following fields:

   Notification Profile

   Specifies the profile to use when this job is run as scheduled. For information about creating notification profiles, see Create Notification Profiles (see page 98).

   Subject

   Specifies the subject line of the email message that is sent by the job.

10. Click Finish.

    The job is created, enabled, and added to the Remediation Jobs table.

Run Remediation Jobs Manually

You can manually run an existing Remediation Job to locate and respond to unwanted configuration changes.

To run a Remediation Job manually

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the Jobs link (below the main tabs).

   The Remediation Jobs page appears.

3. Click the check box next to one or more Remediation Jobs you want to run manually, then select Run Remediation Jobs from the Select Actions drop-down list.

   You are prompted to confirm that you want to run the job.

4. Click OK.

   A message confirms that the job has started successfully. When the job is complete, a notification is sent as defined in the job's associated Remediation Profile.

View and Edit Remediation Jobs

Remediation Jobs are specialized jobs that are performed to correct unwanted changes to a service or server.

To view and edit Remediation Jobs

1. Click the Management Link, and then the Remediation tab.

   The Remediation tab page appears.

2. Click the Jobs link (under the main tabs).

   The Remediation Jobs page appears and displays all of the existing jobs in the Remediation Jobs table.

3. (Optional) Click a job in the Job Name column.

   The Remediation Job Details page appears. You can edit the Remediation Job from this page as described in Create Profile Jobs. (see page 340)

Delete Remediation Jobs

You can delete Remediation Jobs that you no longer need.

To delete Remediation Jobs

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the Jobs link (below the main tabs).

   The Remediation Jobs page appears with the existing jobs listed in the Remediation Jobs table.

3. Click the check box next to one or more jobs you want to delete, then select Delete Remediation Jobs from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   The selected Remediation Jobs are deleted.

View Remediation History

You can view the history of Remediation Jobs that have been run on the Remediation History page. In addition to the top-level history details about the job, you can navigate to details about the job, the servers associated with the job, and the individual remediation steps performed by the job.

To view Remediation History

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the History link (below the main tabs)

   The Remediation History page appears and lists all existing Remediation Jobs in the Remediation History table.

3. Click a link in the Job Name column.

   Details of the selected job appear on the History tab.

4. Click the Servers tab.

   The servers associated with the Remediation Job appears in the table.

5. Click a link in the Server Name column.

   Details of the selected server appear on the Server tab.

6. Click the Steps tab.

   The steps associated with Remediation Job appear.

Delete Remediation History

You can permanently delete Remediation History entries if they are no longer needed.

To delete Remediation History

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the History link (below the main tabs)

   The Remediation History page appears and lists all jobs that have been run in the Remediation History table.

3. Click one or more check boxes next the jobs that you want to delete, then select Delete Remediation History from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

4. Click OK.

   The selected jobs are deleted.

Rerun Remediation Jobs

You can rerun a Remediation Job from the Remediation History page if you notice that the job failed.

To rerun Remediation Jobs

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the History link (below the main tabs).

   The Remediation History page appears and displays the Remediation Jobs that have been run. You can view the entries in the Status Message column for a Job Failed message, or the Run Success column for a red X indicating that the job did not succeed.

3. Click the check box next to one or more jobs that you want to rerun, then select Run Remediation Again from the Select Actions drop-down list.

   You are prompted to confirm that you want to rerun the job.

4. Click OK.

   The job runs. When it completes, the Remediation History table is updated with the results.

5. (Optional) Click the link in the Job Name column to display details about the job, the server where it was run, and the individual steps performed by the Remediation Job.

Undo Remediation

You can undo unwanted changes made by Remediation Jobs. To help you determine if you want to undo changes, you can view the individual steps performed by a Remediation Job as described in View Remediation History (see page 345).

To undo changes made by Remediation Jobs

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the History link (below the main tabs).

   The Remediation History page appears and displays the Remediation Jobs that have been run.

3. Click the check box next to one or more jobs on which you want to undo changes, then select Undo Remediation from the Select Actions drop-down list.

   You are prompted to confirm that you want to undo the changes made by the job.

4. Click OK.

   The undo operation runs and reverts to the configuration settings that existed before the selected job was last run. The undo operation is logged on the Remediation Log page.

View the Remediation Log

You can view the Remediation Log page to observe all remediation activity.

To view the Remediation Log

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the Log link (below the main tabs).

   The Remediation Log page appears with the existing remediation activity listed in the Remediation Log table.

Run Remediation Reports

CA Configuration Automation includes the following Remediation-specific report templates that can be used to generate reports:

- Remediation History—Provides a historical summary of the Remediation jobs for the selected service, and lists all configuration changes made to components in the service.
- Remediation Jobs—Provides execution-specific Remediation job details including run history, summarized results, and individual step details for each Remediation job instance.

To run Remediation reports

1. Click the Management link, then click the Remediation tab.

   The Remediation tab page appears.

2. Click the Reports link (below the main tabs).

   The Report Templates page appears and displays the Remediation reports.

3. Click the link of the report you want to run in the Report Name column.

   The General tab of Report Details page appears.

4. Run the report as described in Run or Save Report Templates.

Chapter 13: Job Management

The Jobs link on the Services and Servers tab pages provides access to the Jobs table where you can view scheduled jobs. The table displays jobs that are scheduled as part of Management Profiles, Notification Profiles, and Network Profiles.

Note: Remediation jobs do not appear in the Jobs table. Remediation jobs appear in the Remediation Jobs table as described in View Remediation Jobs (see page 344).

View Scheduled Jobs

The Jobs table displays details about all of the scheduled jobs that are associated with Management Profiles and Network Profiles.

To view scheduled jobs, click the Management Link then do one of the following actions:

- Click the Services tab, and then the Jobs link (below the main tabs)
- Click the Servers tab, and then the Jobs link (below the main tabs)
- Click the Compliance tab, and then the Jobs link (below the main tabs)
- Click the Remediation tab, and then the Jobs link (below the main tabs)
- Click the Reports tab, and then the Jobs link (below the main tabs)
- Click the Jobs tab

The scheduled jobs appear in the Jobs table. The Jobs tab maintains a history of completed and failed scheduled jobs. Click the completed or error jobs links to view the logs for the selected job. The completed and the error jobs are archived. The archiving of the jobs is based on the following configuration information available in the Administration tab:

job.archive.threshold

Specifies the number of records to exceed before a completed job history archive is created.

Default: 500

job.archive.skip.records

Specifies the number of records to skip before job history archiving remainder.

Default: 200

job.archive.minimum.records

Specifies the number of records before a completed job history archive is created.

Default: 200

View Catalyst jobs

CA Configuration Automation® Connector Guide

The Catalyst Integration tab displays the catalyst jobs, and the Export Summary tab lists the successfully exported or failed CIs. You can view the export summary details of the executed catalyst jobs from the Management or Administrator link. You can export the failed CIs for a selected Catalyst Job while exporting the data to CA Catalyst.

Follow these steps:

1. From the Management link:
   a. Click the Management link, Management Profiles.
   b. Select a management profile, Enable Integration check box to view the Catalyst Export summary tab.
   c. Select the Catalyst Export summary tab.
2. From the Administrator link:
   a. Click the Administrator link, Catalyst Integration tab, and then the Jobs link.
   b. Select a job from the available catalyst jobs executed catalyst.
   c. Select the Export Summary tab.

   The Profile details (using the Management profile), or the Job details (using the Administrator link) page lists the total count of the exported CIs, successfully exported CIs, and failed CIs as part of the Catalyst job.

   You can run catalyst jobs, refresh the export summary, and export failed summary to an excel file from the Actions drop-down.

   The excel file contains the Summary sheet that displays the Job Name and Job Description. The subsequent worksheets illustrate the failed CIs for the selected Catalyst Job.

   If there are failed CIs, a new worksheet is created for each USM Type or Category. All the worksheets have unique ID along with the USM Type or Category name.

   Note: For further reference, view the table (see page 351) that maps ID with the USM Type or Category name.
3. Click a link in listed in the Catalyst CI count (successful) column of a USM type to view the USM web view.

   Note: The USM web view is successfully launched only when you configure the catalyst.server.name and catalyst.server.port properties.
4. Click a link in the Catalyst CI count (failed) column of a particular USM type.

   The source information of the CIs is listed failed during the export to CA Catalyst when the job was run as shown in the following illustration:

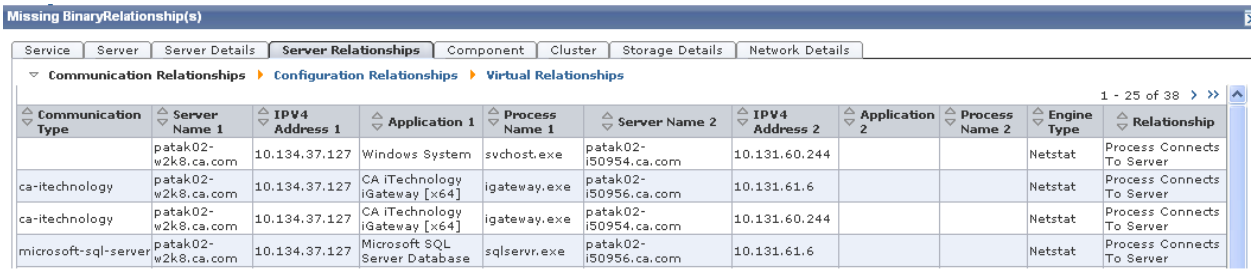

Note: For more information about the failed CIs and how to resolve the errors, see the *CA Configuration Automation Connector Guide*.

Map the ID with the USM Type or Category name

The following table maps the ID with the USM Type or Category name:

ID	GUI Panel Path
1	RELATIONSHIP_Server Details_Person
2	RELATIONSHIP_Server Details_Location
3	RELATIONSHIP_Server Details_Memory
4	RELATIONSHIP_Server Details_Processor
5	RELATIONSHIP_Server Details_OS
6	RELATIONSHIP_Server Details_IPConfig
7	RELATIONSHIP_Server_Compliance Status
8	RELATIONSHIP_Server_Installed Applications
9	RELATIONSHIP_Server_Services and Daemons
10	RELATIONSHIP_Server_Open Ports
11	RELATIONSHIP_Server_Physical Disks
12	RELATIONSHIP_Server_CD or DVD drives
13	RELATIONSHIP_Server_Tape Drives
14	RELATIONSHIP_Server_Port
15	RELATIONSHIP_Server_Port IPConfig
16	RELATIONSHIP_Server Relationships_Communication Relationships

ID	GUI Panel Path
17	RELATIONSHIP_Server Relationships_Configuration Relationships
18	RELATIONSHIP_Server Relationships_Virtual Relationships
19	RELATIONSHIP_Server Relationships_Storage Relationships
20	RELATIONSHIP_Network Details_Port
21	RELATIONSHIP_Network Details_Port IPConfig
22	RELATIONSHIP_Network Details_Cluster
23	RELATIONSHIP_Storage System_Storage LUN
24	RELATIONSHIP_Storage System_Storage Processor
25	RELATIONSHIP _Storage System_Storage Manager
26	RELATIONSHIP_Service
27	RELATIONSHIP_Component
28	RELATIONSHIP_Cluster
29	RELATIONSHIP_Storage Details
30	COMPUTERSYSTEM_Server
31	COMPUTERSYSTEM_Communication Relationships
32	COMPUTERSYSTEM_Configuration Relationships
33	VIRTUALSYSTEM_Server
34	ROUTER_Server
35	SERVICE_Service
36	PROVISIONEDSOFTWARE_Installed Applications
37	PROVISIONEDSOFTWARE_Component
38	COMPLIANCESTATUS_Server
39	COMPLIANCESTATUS_Service
40	LOCATION_Server
41	LOCATION_Service
42	PERSON_Server
43	PERSON_Service
44	OPERATINGSYSTEM_Server
45	PROCESSOR_Server
46	MEMORY_Server

ID	GUI Panel Path
47	MEDIADRIVE_Physical Disks
48	MEDIADRIVE_CD or DVD Drives
49	MEDIADRIVE_Tape Drives
50	MEDIADRIVE_Storage Details
51	BACKGROUNDPROCESS_Services and Daemons
52	BACKGROUNDPROCESS_Open Ports
53	BACKGROUNDPROCESS_Communication Relationships
54	BACKGROUNDPROCESS_Configuration Relationships
55	DISKPARTITION_Logical Partitions
56	PORT_Network Interface Cards
57	PORT_Network Details
58	IPCONFIG_Server
59	IPCONFIG_Network Interface Cards
60	IPCONFIG_Network Details
61	IPCONFIG_iSCSI Initiators
62	FILE_Logical Partitions
63	FILE_Storage Details
64	CLUSTER_Cluster Details
65	CLUSTER_Network Details
66	VIRTUALIZATIONMANAGER_Server
67	HYPERVISORMANAGER_Server
68	STORAGEARRAY_Storage System
69	STORAGEVOLUME_Storage LUN
70	INTERFACECARD_Storage Processor

Test the Connectivity between CA Configuration Automation Server and CA Catalyst Server

The Test Connector Status tab tests the CCA Connector connectivity between CA Configuration Automation Server and CA Catalyst server, and lets you view the server connectivity status. The connectivity status indicators are as follows:

- Responding—CCA Connector is running.
- Not Responding—The CCA connector is not running.

Follow these steps:

1. Click the Administrator link, Catalyst Integration tab, and then the Jobs link.
2. Click the Test Connector Status at the top right corner.
3. Enter the CA Catalyst user name and password.

   The status of the CCA connector connectivity between the CA Configuration Automation Server and CA Catalyst server is displayed.

Chapter 14: Viewing CA Configuration Automation Logs

CA Configuration Automation logs every operation, event, and task that is performed on the Logs tab page. Additionally, some other tab pages have component-specific log pages that display events specific to that tab page. For example, the Blueprints Log page shows Blueprint-specific events including creation, modification, import, export, and so on. These Blueprint-specific events also appear on the main Log tab page.

All CA Configuration Automation log pages have a number of columns in common including the following:

- Date/Time—Displays the timestamp that shows the data and time the event occurred. For example, 2010-04-20-15:18:22
- Message—Displays a system-generated identifier and the confirmation message that describes the event. For example, CCA-SC-5004 File Structure Class ".netwebtruststrclass" updated successfully.
- Log Level—Displays a description of the type of message. The entry can be Informational, Error, Fatal, or Warning.
- Management Server—Displays the name or IP address of the server on which the event occurred.
- Event Type, Event Sub Type, and Job Type—Classifies the event using two or more of these three columns. All events have an Event Type, and Event Sub Type entry. Some events are further classified using a Job Type when appropriate.

  For example, if you create a Network Profile, the Event Type column shows Create, and the Event Sub Type column shows Network Profile. If you schedule a job to run a network scan, the Event Type column shows Run Network Scan, and the Event Sub Type column shows Network Profile, and the Job column shows Network Scan.
- User Identifier—Displays the user who scheduled or performed the task or event.

  Note: Scheduled jobs are owned by an internal user called system_user that owns certain processes and is credited as the creator of some predefined content (for example, predefined Blueprints). The system_user is similar to the Windows user SYSTEM (if you open the Processes tab of the Windows Task Manager you will see SYSTEM in the User Name field for many processes) in that it is an internal user that is not created by an administrator or assigned a user ID. This user does not have an entry in CA EEM, and cannot log into the CA Application Configuration Manager UI.
- Service Name—Displays the name of the service on which the event occurred.

- Server Name—Displays the server that was the target of the event. For example, if you create a snapshot of a server called Bertha from a CA Configuration Automation Server host called Terrapin, Bertha appears in the Server column (and Terrapin appears in the Management Server column).
- Management Profile—Displays the name of the profile assigned to the server or service on which the event occurred.

Turn Off Logging Auto-Refresh

All log pages contain an Auto-Refresh feature that automatically checks for updates. You can disable this functionality if you want to read information already on the page without the interruption of the refresh operation.

To disable Auto-Refresh

1. Click the Management link, then click any of the following:
   - Log tab
   - Services tab, Log link
   - Servers tab, Log link
   - Network tab, Log link
   - Blueprints tab, Log link
   - Compliance tab, Log link
   - Remediation tab, Log link
   - Reports tab, Log link

   The selected Log page appears with Auto-Refresh enabled.
2. Click the Auto-Refresh is On button.

   Auto-Refresh is disabled, and the button displays Auto_Refresh is Off.
3. (Optional) Click Refresh.

   The page refreshes.
4. Restart the CCA Server for the new value to take effect.

Configure Logging Auto-Refresh Interval

When Auto-Refresh is enabled on a Log page, it refreshes the page automatically at the default interval of 50 seconds. You can increase or decrease the refresh rate.

To edit the logging Auto-Refresh interval

1. Click the Administration link, the Configuration tab, and then the Properties link.

   The Properties page appears.

2. Locate the auto.refresh.limit property in the cca group, then click the corresponding Value field.

   An editor appears in the field .

3. Edit the value to increase or decrease the refresh interval, then press Enter.

   The new value is saved and all Log pages refresh at the new rate.

View Archived Logs

When the logs exceed the maximum storage size limit they are archived, and are available in the Log Archives tab. The Log Archives tab lets you view the archived logs, and use the log information for diagnostics, or troubleshooting.

The Log tab in the Management panel displays information about archived logs. To view the archived logs for a specified date and time, use the log archive filter.

Follow these steps:

1. Click the Administration, Diagnostics, Log Archives tab.

   The archived files are displayed with its description, creation time, log size, and the path to locate the logs.

2. Select an archive file link to open or download the archived (CSV) file in Microsoft Excel or Notepad.

3. (Optional) Select Delete Log Archives form the Select Actions to remove the archived logs if it exceeds the storage capacity.

Chapter 15: Report Management

CA Configuration Automation provides predefined reports to obtain detailed information about CA Configuration Automation-managed services, servers, blueprints, software, and data.

Reports are originally installed as report templates which can be run as is, or saved to create a custom instance of the report.

Reports can be scheduled as jobs, run manually, viewed online, printed, or exported, however the underlying data used to generate the report cannot be edited.

In addition to accessing all report templates and custom reports from the Reports tab page, you can also access certain object-specific reports from the other tab pages. For example, you can access server-specific reports and report templates from the Reports link on the Server tab page, and Remediation-specific reports and report templates from the Reports link on the Remediation tab page.

The Reports tab page contains links to the following report pages:

- Templates—Lists the predefined reports
- Saved Reports—Lists the reports customized by opening and saving a report template
- Jobs—Lists all reports that are scheduled or are currently running
- Log—Lists all reporting activity

Run or Save Report Templates

You can manually run a report using any of the predefined report templates from the Report Templates page. You must at least specify a target for the report (for example, a service, server, or Blueprint) if there is an entry in the Target column of the Report Templates table. You can also edit any other settings before generating a report to return more specific information.

If you think that the report settings you specify while generating a manual report are something you could reuse in the future, you can save the settings as a custom report. For example, if there were servers in your organization named Bertha, Minglewood, and Darkstar, you could customize the Change Detection - Servers report template to be a custom report for each of them: Change Detection - Bertha, Change Detection - Minglewood, and Change Detection - Darkstar. The custom reports are displayed on the Saved Reports page.

Note: If you want to schedule a report to run at a specific time you must save it as a custom report.

To run a predefined report template or save a template as a custom report

1. Click the Management link, then click the Reports tab.

   The Report Templates page appears unless you have created custom reports, in which case the Saved Reports page appears.

2. Click the name of the template you want to run or save in Report Name column of the the Report Templates table.

   Note: Some report templates have the same name except for the target identifier (Servers, Service, Virtualization, or Blueprints). This supports either the server- or service-centric methods of configuration management.

   The Details page for the selected template appears with the General tab displayed by default.

3. Edit the following Details fields as appropriate:

   Name

   Specifies the name of the report template, or the report if you save it.

   Description

   Describes the purpose of the report.

   Share the Report With Other Users

   Specifies whether the saved report is available to other users, or only the user who created it.

   Format

   Specifies the output format of the report. The following options are available:

   - Crystal Reports
   - CSV
   - Microsoft Excel
   - Microsoft Excel (Data Only)
   - PDF
   - Rich Text Format

4. Click the check boxes to include or exclude the report options in the Customizable Output Fields area.

   The options with a check are included in the report.

5. Click the Targets tab if it is available for the selected report template, or skip to step 7.

   The Targets tab page appears and displays the available targets that corresponds with the selected report type (Servers and Server Groups, Services, or Blueprints).

6. Double-click any entry in the Available *<object>* column (you can also click an object, then click the right-facing arrow to move the object to the Selected column). For example, if you selected Rule Compliance (Servers) as the report template, the Target tab contains Available Servers and Selected Servers columns.

   The selected object appears in the Selected *<object>* column.

7. (Optional) Click the Filters tab to further control what information is included or excluded from the report.

8. (Optional) Click the Schedule tab to create a schedule for the report to run, and the details of the notification that is sent when the report is generated.

   The options are:

   Not Scheduled

   Specifies that you must run the job manually (by clicking Run/View as described later in this procedure).

   Once

   Specifies that the report runs one time only. When you select this option the Time field appears where you specify the date and time for this job to run.

   Minutes

   Specifies that the report runs at an interval defined by minutes. When you select this option, additional fields appear where you specify the date to start and end, the start time, and the interval in minutes.

   Hourly

   Specifies that the report runs at an interval defined by hours. When you select this option, additional fields appear where you specify the date to start and end, the start time, and the interval in hours.

   Daily

   Specifies that the report runs at an interval defined by days. When you select this option, additional fields appear where you specify the date to start and end, the start time, and the interval in days.

   Weekly

   Specifies that the report runs every seven days on a specific day. When you select this option, additional fields appear where you specify the date to start and end, the start time, and the day.

   Monthly

   Specifies that the report runs on a monthly interval (configurable from 1 through 12) on a specific date (or the last day of the month). When you select this option, additional fields appear where you specify the date to start and end, the interval, the start time, and the date.

Notification Profile

Specifies the notification profile you want to use to generate the notification that is send after the report is generated.

Subject

Specify the Subject of the email message sent by the notification profile.

Scheduled report jobs can be viewed in the following locations:

- Jobs tab page—This page shows all scheduled jobs.
- Jobs link on the Reports tab page—This page shows only scheduled report jobs.

9. (Optional) Click the Destination tab, and select one of the following delivery types from the Delivery Method drop-down list:

None

Specifies that the report is not delivered. Instead, it opens on the client computer after it is generated.

FTP

Specifies that the report is delivered by file transfer protocol (FTP) to the location specified in the following fields:

Server

Specifies the server name where the report is delivered.

User/Login

Specifies a user account on the specified server with write privileges.

Path

Specifies the location on the specified server to save the report.

Password

Specifies the password for the user name specified in the User/Login field.

Retype Password

Ensures that the password was typed correctly in the Password field.

File Name

Specifies the name assigned to the report file. The file name you enter here is appended with the file extension that corresponds with the report format you selected in step 3. The file name is also suffixed with the time stamp when the report was executed. For example, if the file name is Servers_Open_Ports and the selected format is Adobe Acrobat (PDF), the file name will be Servers_Open_Ports2010-04-20-16-19-59.pdf.

File System

Specifies that the report is saved in a non-default directory on the server on which BO is installed as specified in the following fields:

Path

Specifies the directory in which to save the report.

File Name

Specifies the name assigned to the report file. The file name you enter here is appended with the file extension that corresponds with the report format you selected in step 3. The file name is also suffixed with the time stamp when the report was executed. For example, if the file name is Servers_Open_Ports and the selected format is Adobe Acrobat (PDF), the file name will be Servers_Open_Ports2010-04-20-16-19-59.pdf.

10. Do one of the following:
    - Click Save to create a customized report based on the modifications you made to the report template.

      The newly created custom report is displayed in the list of reports on the Saved Reports page.
    - Click Run/View to generate the report based on the current settings in the selected report template.

      The report runs, and then appears in a web browser.

      You can save, print, or export the generated report using the options in the Table Actions drop-down list, and your browser.
11. (Optional) After you run a report, report instance is saved. To view the report, click report name and the Instances tab. The tab lists all the report instances that you executed and the 'shared' report instances executed by others.
    - Click on any report instance name to view the report.
    - Select the check box next to one or more report instances then select Share Instances from the Select Actions drop-down list to enable the report instances to be viewed by other users.
    - Select Remove Sharing from the Select Actions drop-down list to restrict the report instances to be viewed only by the owner of the instance.
    - Select Delete Instances from the Select Actions drop-down list to delete the instances.

Run Saved Reports

The Saved Reports page displays the reports that were saved as custom reports as described in Run or Save Report Templates. You can manually run any saved report even if it is scheduled to run automatically in the future.

To run a saved (custom) report

1. Click the Management link, the Reports tab, and then the Saved Reports link.

   The Saved Reports page appears and lists the custom reports in the Report Name column.

2. (Optional) Click the report name to view details about the report.

   The Details page for the selected report appears.

3. Do one of the following:

   - Click the check box next to one or more reports you want to run, then select Run from the Select Actions drop-down list. Report Job will be initiated and will be listed in the Jobs sub tab on the Reports tab page.
   - Click report name to display report details page and click Run or View button to generate the report based on the current settings in the selected saved report. The report runs, and then appears in a web browser.
   - If report is saved as a job, and runs at the time specified on the Schedule tab on the report's Details page. Once the job is completed, The report instance can be viewed from the instances tab in the details page of the report selected.

   Once the run report or report job is complete, view the the report result instance from the instances tab in the details page of the report selected.

Delete Saved Reports

You can delete any saved report even if it is scheduled to run.

To delete a saved (custom) report

1. Click the Management link, the Reports tab, and then the Saved Reports link.

   The Saved Reports page appears and lists the custom reports in the Report Name column.

2. Click the check box next to one or more reports you want to delete, then select Delete Reports from the Select Actions drop-down list.

   You are prompted to confirm the deletion.

3. Click OK.

   The selected reports are deleted.

Chapter 16: Dashboards and Visualization

CA Configuration Automation includes the following two types of graphical views of elements in your enterprise:

- Dashboards—Displays a graphical summary of information about servers, services, management operations (Change Detection, Rule Compliance, and so on), and virtualized components.
- Visualization—Displays the relationships of your networked elements. These relationships can be between any combination of servers, services, applications, and software components.

These views are accessed from the Dashboard link (top right of the UI). The Charts tab displays the dashboards, and the Visualization tab displays the network associations of servers, services, and software components.

Dashboards

CA Configuration Automation includes the following Dashboards:

- Change History
- Clusters
- Communications Relationships
- Compliance
- Servers
- Servers (Managed)
- Servers (Unmanaged)
- Services
- Software Components
- Storage
- Virtualization
- VM Guest Servers
- VM Guest Software Components
- VM Hosting Services

Each Dashboard displays a number of default *charts* (also know as *portlets*) that display detailed information about the Dashboard elements. For example, Servers is the default Dashboard that appears when you click the Dashboard link. The Servers Dashboard displays the following charts by default:

- Servers by OS Family and Operating System
- Servers by Logical CPU Count
- Servers by Server Group
- Servers by Memory Capacity
- Servers by Service
- Servers by Manufacturer

Dashboards, and their corresponding charts, contain options for displaying them, configuring them, removing them, refreshing them, and changing how they display information. This chapter describes those options.

Chart Overview

Consider the following points while working with charts and dashboards:

- Place your cursor over a chart to display informational text
- Summary Table
  - Includes typical table functionality (paging and column sorting)
  - Some charts have a Drill Down header in the first column
    - If you select an object in the column, you drill down into another chart that is based on the information in that column
    - The same drill down action occurs if you click on the corresponding item in the chart in the top half of the frame.
    - Breadcrumb trail takes you back to original chart
    - Configurable options for drill down chart vary from the original chart
    - Any filter that was applied on original chart is propagated to the underlying chart
    - Some charts may have multiple drill downs (that is, multiple nested charts)
    - Some charts have more than one drill down option available. In this case, the column header includes a drop-down list where you can control which chart your mouse click drills into
  - Some rows may include links in one or both columns
    - Links take you into Management pages with context-sensitive filters (breadcrumbs to take you back to chart)

- Chart configuration opens in an independent window
  - You can make edits and click Apply to determine their impact on the chart
  - When this windows is open, you cannot perform drill-down or link actions on the chart (the configuration window will flash)
- When you change the filter criteria, the meaning of the chart may vary dramatically. You may want to change the name of the chart title as it is displayed in the dashboard (by design, this does not change the chart name on the left side)

Display a Dashboard

The currently selected Dashboard appears when you click the Dashboard link. Perform the following procedure to display a different Dashboard.

Follow these steps:

1. Click the Dashboard link.

   The Dashboard panel displays the currently selected Dashboard.

2. Select New from the Configure Dashboard drop-down list.

   The New Dashboard page appears.

3. Click the Dashboard you to display in the Dashboard pane.

   The selected Dashboard appears in the New Dashboard pane.

Configure a Dashboard Portlet

Each Dashboard portlet (also known as a chart) contains the following controls that determine how the portlet is displayed, how the portlet displays the associated data, and what data is displayed in the portlet:


Configure the Portlet—Enables you to configure the portlet as described later in this topic.


Refresh the Portlet—Click this icon to display the current data for the corresponding portlet.


Minimize the Portlet—Click this icon to reduce the corresponding portlet to only display the title bar. Click this icon again to restore the portlet to its original size.


Maximize the Portlet—Click this icon to enlarge the corresponding portlet to fill the web browser. Click this icon again to restore the portlet to its original size.

Close the Portlet—Click this icon to remove the corresponding portlet from the page. You can restore the portlet by selecting New from the Configure Dashboard drop-down list, and selecting the Dashboard of to which the portlet is assigned.

To Configure a Dashboard portlet

1. Click the Dashboard link.

   The Dashboard panel describes the currently selected Dashboard.

2. Click the Configure the Portlet icon in the chart you want to edit.

   The Edit *<chart_title>* dialog appears.

3. Edit the information in the following fields (not all fields are available for all portlets), then click Apply to save the changes:

   Chart Title

   Specifies the display name for the selected chart. For example, if you changed the Server Group field to only include the user-defined group Linux/UNIX, you could change the title to reflect the change.

   Chart Type

   Specifies either Pie or Bar.

   Available Server Groups

   Specifies the server groups that can be displayed in the chart. Double-click a server group to move it to the Selected Server Groups column.

   Selected Server Groups

   Specifies the server groups that will be displayed in the chart. If no server group is specified, all server groups are displayed unless otherwise filtered from the chart.

   Available Services

   Specifies the services that can be displayed in the chart. Double-click a service to move it to the Selected Services column.

   Selected Services

   Specifies the service that will be displayed in the chart. If no service is specified, all services are displayed unless otherwise filtered from the chart.

   Virtualization Filter

   Includes three filter check boxes where you can specify whether Non-virtual Servers, VM Hosting Servers, and VM Guest Servers are displayed in the chart. A check mark indicates that the corresponding option is included in the chart.

Server State

Includes four filter check boxes where you can specify whether servers with the following states: New, Managed, Unmanaged, Imported are displayed in the chart. A check mark indicates that the corresponding option is included in the chart.

Available Operating System

Specifies the operating systems that can be displayed in the chart. Double-click an operating system to move it to the Selected Operating System column.

Selected Operating System

Specifies the operating system that will be displayed in the chart. If no operating system is specified, all operating systems are displayed unless otherwise filtered from the chart.

Note: Some portlets may contain additional fields that enable you to specify or filter other elements from the chart, or to specify a time period to include.

Your changes are saved and the chart is updated.

4. Click Close when you are done making changes.

   The Edit dialog closes.

Create a New Dashboard

Perform the procedure in this section to create a brand new Dashboard. You can also create a new Dashboard using one of the predefined Dashboards as a starting point as described in Create a Custom Dashboard (see page 370).

To create a new Dashboard

1. Click the Dashboard link.

   The Dashboard panel describes the currently selected Dashboard.

2. Select New from the Configure Dashboard drop-down list.

   The following panes appear:

   - Dashboard—Lists the existing Dashboards and Charts.
   - New Dashboard—Blank pane where you can drag-and-drop charts.

3. Click the plus sign (+) to expand the Charts folder, then repeat for any subfolder.

   Note: The charts are contained in element-specific folders (for example, Servers, Services, Applications, and so on) designed to help you locate the appropriate chart, and also in an All Charts folder that contains all charts.

   The charts appear below their folder.

4. Do one or more of the following:
   - Double-click a chart in the Charts folder.
   - Right-click a chart in the Charts folder, then select Add Chart to Dashboard.

   The selected charts appear in the New Dashboard pane.
5. (Optional) Do one or more of the following:
   - Click the title bar of a chart and drag-and-drop it to reposition it.
   - Click either side or the bottom of the chart to resize it.

   The charts are repositioned and resized as appropriate.
6. Select Save from the Configure Dashboard drop-down list

   The Dashboard dialog appears.
7. Enter a name for the custom Dashboard, click the Shared Dashboard check box if you want other users to be able to access the Dashboard, then click OK.

   The New Dashboard pane displays the new name, and the Dashboard appears in the Dashboards folder in the Dashboard pane.
8. (Optional) Configure a chart as described in Configure a Dashboard Portlet (see page 367).

   The new Dashboard is updated with your edits.
9. (Optional) Right-click the new Dashboard and select Set Default.

   The new Dashboard displays by default.

Create a Custom Dashboard

You can create custom Dashboards using one of the predefined Dashboards as a starting point, then adding and removing charts as required.

To create custom Dashboards

1. Click the Dashboard link.

   The Dashboard panel describes the currently selected Dashboard.
2. Select Save As from the Configure Dashboard drop-down list.

   The Dashboard dialog appears.
3. Enter a name for the custom Dashboard, click the Shared Dashboard check box if you want other users to be able to access the Dashboard, then click OK.

   The custom Dashboard appears in the Dashboards folder in the Dashboard pane.
4. Click the custom Dashboard.

   The custom Dashboard appears in the right pane.

5. Do one or more of the following:
   - Click the Remove Portlet icon to remove any unwanted charts.
   - Double-click any chart in the Charts folder in the Dashboard pane to add the chart to the Dashboard.
   - Click the title bar of any chart and drag it to the desired location on the Dashboard.
   - Configure a chart as described in Configure a Dashboard Portlet (see page 367).

   The updated custom Dashboard appears in the right pane.
6. Select Save from the Configure Dashboard drop-down list.

   The custom Dashboard is updated with your edits.
7. (Optional) Right-click the new Dashboard and select Set Default.

   The new Dashboard displays by default.

Export Dashboards

An export utility is provided to export new or custom Dashboards from a CA Configuration Automation implementation so it can be imported into another. The export utility extracts the data for the selected Dashboard, then saves the data to a .jar file.

To export Dashboards

1. Create a new or custom Dashboard as described in either of the following sections:
   - Create a New Dashboard (see page 369)
   - Create a Custom Dashboard (see page 370)
2. Right-click the new or custom Dashboard that you want to export, then select Export Dashboard from the menu that appears.

   A File Download window prompts you for the location where you want to save the export .jar file.

The file is assigned a default name using the following timestamp convention:

Dashboard_Export_*YYYY_MM_DD_HH_MM_SS*.jar

Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

Dashboard_Export_2010_01_06_13_59_57.jar

3. Click Save, specify the location, then click OK.

   Note: If you are using a Windows operating system, Windows may suggest saving the file as a .zip file. Ensure the .jar is included in the file name, and the Save As Type field is set to All Files.

   The Dashboard is saved in the specified location.

Import Dashboards

An import utility is provided to import Dashboards (in the form of .jar files) into your CA Configuration Automation Server.

To import Dashboards

1. Click the Dashboards link.

   The Dashboards panel appears.

2. Select Import Dashboard from the Configure Dashboards drop-down list.

   The Import Dashboards dialog appears.

3. Click Browse and navigate to the .jar file you want to import.

   If the file was exported from a CA Configuration Automation Server, the file was assigned a default name using the following timestamp convention:

   Dashboard_Export_*YYYY_MM_DD_HH_MM_SS*.jar

Where *YYYY* is the year, the first *MM* is the month, *DD* is the day, *HH* is the hour (using a 24-hour clock), the second *MM* is the minutes, and *SS* is the seconds. For example:

Dashboard_Export_2010_01_06_13_59_57.jar

4. Click the Overwrite Existing Dashboards option if you want to update the existing version on your CA Configuration Automation Server, then click one of the following buttons:

   Import All

   Imports all Dashboards in the .jar file.

   Import On Selection

   Displays the Available Imports table which lists the Dashboards in the .jar file, and enables you to specify which of them you want to import.

   The import begins and may take a few moments depending on the size of the file. When the import is complete a confirmation message appears.

Visualization

The Visualization tab contains an expandable tree view with the following main nodes and sub-nodes:

- Graphs
  - Applications
  - Clusters
  - Servers
  - Service Profiler
  - Services
  - Software Components
- Templates
  - Applications
  - Clusters
  - Servers
  - Service Profiler
  - Services
  - Software Components

Each Templates sub-node contains options that determine the elements that are displayed. For example, the Servers sub-node contains the following options:

- All Server Relationships
- Server Change Detection
- Server Communication Relationships
- Server Configuration Relationships
- Server Rule Compliance
- Server Storage Management Relationships
- Server Storage Relationships
- Server Virtualization Relationships

View Predefined Graphs

This section describes the default graphs included in the left pane of the CA Configuration Automation Visualization tab. You can select any of these graphs to view a graphical representation of the components and relationships they include.

Typically, the predefined graphs display a very large amount of information because they are designed to show *all* (that is, all Windows servers, or all communication relationships, for example). To view smaller, more manageable graphical representations, you can select a predefined template (also located on the Visualization tab), create a filter that controls what appears in the view, and save the template as a custom graph as described in To Create a Graph from a Template (see page 379).

The following predefined graphs are available on the Visualization tab:

- Applications
  - All Installed Applications By Server—Displays all installed applications by server.
  - All Installed Applications By Service—Displays all installed applications by service.
  - All Server Application Relationships—Displays all application relationships for all servers.
  - All Servers By Installed Application—Displays all servers by installed application.
  - All Service Application Relationships—Displays all application relationships for all services.

- All Services By Installed Application—Displays all services by installed application.
- Managed Server Application Relationships—Displays all communication relationships for all installed applications on managed servers.
- Managed Servers By Installed Application—Displays all installed applications on managed servers
- MSSQL Service Application Relationships—Displays MSSQL application relationships for all services.
- Oracle Service Application Relationships—Displays Oracle application relationships for all services.
- Unix Server Application Relationships—Displays all communication relationships for all installed applications on Unix servers.
- Unix Servers By Installed Application—Displays all installed applications on Unix servers.
- VM Guest Server Application Relationships—Displays all communication relationships for installed applications on servers that are virtual guest servers.
- VM Guest Servers By Installed Application—Displays all installed applications on virtual guest servers.
- VM Hosting Server Application Relationships—Displays all communication relationships for installed applications on all virtual hosting servers.
- VM Hosting Servers By Installed Application—Displays all installed applications on virtual hosting servers.
- Windows Server Application Relationships—Displays all communication relationships for all installed applications on Windows servers.
- Windows Servers By Installed Application—Displays all installed applications on Windows servers.

- Clusters
  - All Clusters—Displays all known server clusters.
- Servers
  - All Server Change Detection—Displays all change detection results for all servers.
  - All Server Communication Relationships —Displays all communication relationships for all servers.
  - All Server Configuration Relationships—Displays all configuration relationships for all servers.
  - All Server Rule Compliance—Displays all rule compliance results for all servers.

- All Server Storage Management Relationships—Displays all relationships between storage managers and storage devices on all servers.
- All Server Storage Relationships—Displays all storage relationships for all servers.
- All Server Virtualization Relationships—Displays all virtualization relationships for all servers.
- Differences Found Server Change Detection—Displays all change detection results for all servers where differences were found.
- Failures Found Server Rule Compliance—Displays all rule compliance results for all servers where failures were found.
- Managed Server Communication Relationships—Displays all communication relationships for all managed servers.
- Managed Server Configuration Relationships—Displays all configuration relationships for all managed servers.
- Managed Server Relationships—Displays all relationships for all managed servers.
- Managed Server Virtualization Relationships—Displays all virtualization relationships for all managed servers.
- MSSQL Server Communication Relationships—Displays all communication relationships for all MSSQL servers.
- MSSQL Server Relationships—Displays all relationships for all MSSQL servers.
- Oracle Server Communication Relationships—Displays all communication relationships for all Oracle servers.
- Oracle Server Relationships—Displays all relationships for all Oracle servers.
- Unix Server Change Detection—Displays all change detection results for all Unix servers.
- Unix Server Communication Relationships—Displays all communication relationships for all Unix servers.
- Unix Server Configuration Relationships—Displays all configuration relationships for all managed servers.
- Unix Server Relationships—Displays all relationships for all Unix servers.
- Unix Server Rule Compliance—Displays all rule compliance results for all Unix servers.
- Unix Server Storage Relationships—Displays all storage relationships for the Unix servers.
- Unix Server Virtualization—Displays all virtualization relationships for all Unix servers.

- VM Guest Server Change Detection—Displays all change detection results for all virtual guest servers.
- VM Guest Server Communication Relationships—Displays all communication relationships for all virtual guest servers.
- VM Guest Server Configuration Relationships—Displays all configuration relationships for all virtual guest servers.
- VM Guest Server Relationships—Displays all relationships for all virtual guest servers.
- VM Guest Server Rule Compliance—Displays all rule compliance results for all virtual guest servers.
- VM Guest Server Storage Relationships—Displays all storage relationships for the virtual guest servers.
- VM Guest Server Virtualization—Displays all virtualization relationships for all virtual guest servers.
- VM Hosting Server Change Detection—Displays all change detection results for all virtual hosting servers.
- VM Hosting Server Communication Relationships—Displays all communication relationships for all virtual hosting servers.
- VM Hosting Server Configuration Relationships—Displays all configuration relationships for all virtual hosting servers.
- VM Hosting Server Relationships—Displays all relationships for all virtual hosting servers.
- VM Hosting Server Rule Compliance—Displays all rule compliance results for all virtual hosting servers.
- VM Hosting Server Storage Relationships—Displays all storage relationships for all virtual hosting servers.
- VM Hosting Server Virtualization—Displays all virtualization relationships for all virtual hosting servers.
- Windows Server Change Detection—Displays all change detection results for all Windows servers.
- Windows Server Communication Relationships—Displays all communication relationships for all Windows servers.
- Windows Server Configuration Relationships—Displays all configuration relationships for all Windows servers.
- Windows Server Relationships—Displays all relationships for all Windows servers.
- Windows Server Rule Compliance—Displays all rule compliance results for all Windows servers.

- Windows Server Storage Relationships—Displays all storage relationships results for all Windows servers.
- Windows Server Virtualization—Displays all virtualization relationships for all Windows servers.

- Service Profiler (there are no predefined graphs for the Service Profiler)
- Services
  - All Service Change Detection—Displays all change detection results for all services.
  - All Service Communication Relationships—Displays all communication relationships for all services.
  - All Service Rule Compliance—Displays all rule compliance results for all services.
  - Differences Found Service Change Detection—Displays all change detection results for all services where differences were found.
  - Failures Found Service Rule Compliance—Displays all rule compliance results for all services where failures were found.
  - MSSQL Service Communication Relationships—Displays all MSSQL communication relationships for all services.
  - Oracle Service Communication Relationships—Displays all Oracle communication relationships for all services.
- Software Components
  - All Components By Server—Displays all the components belonging to a server or group of servers.
  - All Components By Service—Displays all the components belonging to a service or group of services.
  - All Server Component Relationships—Displays all server component relationships.
  - All Servers By Component—Displays all servers by component.
  - All Service Component Relationships—Displays all service component relationships.
  - All Services by Component—Displays all services by component.
  - Unix Servers By Component—Displays all Unix servers by component.
  - VM Guest Servers By Component—Displays all VM guest servers by component.
  - VM Hosting Servers By Component—Displays all VM hosting servers by component.
  - Windows Servers By Component—Displays all Windows servers by component.

How to Create a Graph from a Template

In general, you need to perform the following steps to create a custom visualization graph from one of the predefined visualization templates:

1. Select one of the templates from the Templates node of the Visualization tree.
2. Create a graph filter that determines the specific elements you want to appear in the custom graph (for example, all Linux servers).
3. Save the now-customized template as a custom graph.

These steps are described in detail in the sections that follow.

Select a Visualization Template

This section describes the predefined visualization templates included with CA Configuration Automation and the procedure for selecting a template to customize and save as a graph.

Follow these steps:

1. Click the Dashboards link, and then click the Visualization tab.

   The Visualization tree appears in the Visualization pane.
2. Click the plus sign next to the Templates node in the tree.

   The Templates node expands to display the following sub-nodes:

   - Applications
   - Clusters
   - Servers
   - Service Profiler
   - Services
   - Software Components
3. Click the plus sign next to the sub-node whose view you want to use for this graph. For example, if you want to create a server-centric view, click the plus sign next to the Servers node.

   The templates associated with the selected sub-node appear.
4. Select one of the following templates to create the graph:

   - Applications
     - Installed Applications by Server—Displays the applications installed on specified servers or server groups.
     - Installed Applications by Service—Displays the applications that are included in specified services.

- Server Application Relationships—Displays servers with specific applications that are associated by a specific communication type (for example rlogin or http_proxy).
- Servers By Installed Application—Displays servers that have a specific application installed.
- Service Application Relationships—Displays services with specific applications that are associated by a specific communication type (for example rlogin or http_proxy).
- Services By Installed Application—Displays services that have a specific application installed.

- Clusters
  - All Clusters—Displays all known server clusters.
- Servers
  - All Server Relationships—Displays servers specified by operating system and communication type.
  - Server Change Detection—Includes servers specified by operating system and whether a Change Detection operation found differences.
  - Server Communication Relationships—Displays servers specified by operating system and communication type.
  - Server Configuration Relationships—Displays the server, application, and database relationships obtained by an NDG discovery. For example, application A on server 1 is configured to use database B on server 2. The relationships are determined using values found in configuration files and database entries.
  - Server Rule Compliance—Displays servers specified by operating system and whether a Rule Compliance operation found compliance failures.
  - Server Storage Management Relationships—Displays the relationship between servers with storage manager software and the storage devices they manage. The relationship type is always Manages.
  - Server Storage Relationships—Displays the relationship between servers and storage devices and servers with storage manager software and the storage devices they manage. The relationship type can be Storage or Manages.
  - Server Virtualization Relationships—Displays any combination of physical, virtualized host, and virtualized guest servers.

- Service Profiler
  - Service Profiler All Server Relationships—Displays the relationships between all known servers.
- Services
  - Service Change Detection—Displays services specified by operating system and whether a Change Detection operation found differences.
  - Service Communication Relationships—Displays services specified by operating system and communication type.
  - Service Rule Compliance—Displays services specified by operating system and whether a Rule Compliance operation found compliance failures.
- Software Components
  - Components by Server—Displays the components that are installed on the specified servers.
  - Components by Service—Displays services that include one or more specified software component.
  - Server Component Relationships—Displays servers that either reference or communicate with other servers.
  - Servers by Components—Displays servers where one or more specified software component is installed.
  - Service Component Relationships—Displays services that either reference or communicate with other services.
  - Services by Components—Displays services where one or more specified software component is installed.

The Graph Filter dialog for the selected template appears. Continue to the next (see page 381) section for instructions about creating a graph filter.

Create a Graph Filter

A graph filter is a filtering mechanism that determines the specific servers, services, relationships, applications, and software components that are displayed in your custom graphs. Creating a graph filter enables you to specify graph views of network objects that are important to you to monitor. For example, if your organization makes assignments based on operation system expertise, you may be assigned to monitor the configuration of servers with Solaris operating systems. You can create a filter whose associated graph only shows the servers with Solaris operating systems.

Note: All filters work using an implied AND when using multiple filter criteria. For example if you create a graph filter that includes a selected server called factotum.ca.com and a server group called My_assigned_servers, and factotum.ca.com was *not* part of the server group, the filter would not display any servers.

To create a graph filter

1. Select a visualization template as described in the previous (see page 379) section.

   The Graph Filter dialog appears.

2. Enter or select the following information as appropriate.

   Note: The tabs and fields described in this step are not available for every visualization template. For example, Change Detection and Rule Compliance options are specific to the templates that include their names, while server or operating system options are available for most templates.

   Applications tab

   Include Patches

   Specifies whether to include application patches in the list of available applications. This option is not available on every template's Applications tab.

   Include Version

   Specifies whether to include application release numbers in the list of available applications. This option is not available on every template's Applications tab.

   Available Applications

   Lists the applications known to CA Configuration Automation that can be included in the graph. Double-click an available application to move it to the Selected Applications field.

   Selected Applications

   Lists the applications that will be included in the graph filter.

   Available Publishers

   Lists the companies or organizations that created the listed applications. Double-click an available publisher to move it to the Selected Publishers field.

   Selected Publishers

   Lists the companies or organizations that will be included in the graph filter.

   Change Detection tab

   Change Detection

   Specifies whether the results of Change Detection operations for the server or service are shown on the graph. The options are any combination of Differences Found, No Differences Found, and No Results Found.

Communication tab

Limit displayed relationships to target servers/applications only

Specifies whether only the relationships of the target servers or applications appear in the graph. If you clear this check box, relationships of target and source servers and applications are shown.

Relationship Type

Lists the available types of relationships for servers within a service. The following options are available on this drop-down list:

Providing Servers—For source servers, the relationship type is "is used by," for destination servers, the relationship types is "uses."

Consuming Servers—For source servers, the relationship type is "uses," for destination servers, the relationship types is "is used by."

Providers and Consumers—Includes all relationships with a relationships type of "uses" and "is used by."

Communicates With—Includes all relationships with a relationship type of "communicates with."

Available Communication Type

Lists the specific communication protocols used by servers in a service. Double-click an available operating system to move it to the Selected Communication field.

Selected Communication Type

Lists the types of communications that will be included in the graph filter.

Date Range

Specifies the time period for which the communication relationship appears in the graph. For example, if you select Today, only the relationships for the current day are shown. If you select Last N Months, the Months field appears and you must select the number of months to display.

Components tab

Available Components

Lists the specific software components and versions known to CA Configuration Automation that can be included in the graph filter. These components have a corresponding predefined or custom Blueprint. Double-click an available component to move it to the Selected Components field.

Selected Components

Lists the components that will be included in the graph filter.

Available Categories

Lists the software categories defined in blueprints (for example, operating systems or relational databases) that can be included in the graph filter. Double-click an available category to move it to the Selected Categories field.

Selected Categories

Lists the categories that will be included in the graph filter.

Configuration tab

Expand Nodes

Specifies if the parent node is expanded to show child nodes in the graph. If you do not select this option, you can manually expand the node in the final graph.

Display Labels

Specifies if the descriptive label is displayed on the relationship lines between nodes in the graph. If you do not select this option, you can mouse-over elements in the final graph to temporarily display pop-up labels. This option is *not* available on every template's Configuration tab.

Operating System tab

Available OS Family

Lists the general operating systems and device types to include in the graph. The WINDOWS option includes both WIN32 and WIN64 families. Double-click an available OS family to move it to the Selected OS Family field.

Selected OS Family

Lists the operating system families that will be included in the graph filter.

Available Operating System

Lists specific operating systems, and in some cases, the specific version to include in the graph. Double-click an available operating system to move it to the Selected Operating System field.

Selected Operating System

Lists the operating systems that will be included in the graph filter.

Relationships tab

Available Relationship Type

Lists the following component relationship types available for the graph: Communicates With and References. Double-click an available relationship type to move it to the Selected Relationship Type field.

Selected Relationship Type

Lists the types of relationships that will be included in the graph filter.

Available Specific Relationship Type

Lists the specific relationship types available for the graph. Options include Database Connectivity, File Reference, Virtualizes, and so on. Double-click an available specific relationship type to move it to the Selected Specific Relationship Type field.

Selected Specific Relationship Type

Lists the types of relationships that will be included in the graph filter.

Rule Compliance tab

Rules Compliance

Specifies whether the results of Rule Compliance operations for the server or service are shown on the graph. The options are to show any combination of Failures Found, No Failures Found, and No Results.

Servers tab

Server Name Mask

Accepts a fully qualified server name or a masked server name with an either an asterisk (*) or percent sign (%) as a wild card.

IP Address Mask

Accepts a fully qualified IPv4 or IPv6 address or a masked IPv4 IPv6 address with an asterisk (*) or percent sign (%) as a wild card. The field does not accept IP Address ranges or CIDR notation.

Available Servers

Lists the servers being managed by CA Configuration Automation. These servers can be discovered by, or imported into CA Configuration Automation. Double-click an available server to move it to the Selected Server field.

Selected Servers

Lists the servers that will be included in the graph filter. Double-click an available server to move it to the Selected Server field.

Available Server Groups

Lists the server groups (including dynamic server groups) being managed by CA Configuration Automation. These server groups can be created by, or imported into CA Configuration Automation. Double-click an available server group to move it to the Selected Server Group field.

Selected Server Groups

Lists the server groups that will be included in the graph filter.

Available Services

Lists the services being managed by CA Configuration Automation. These services can be created in, discovered by, or imported into CA Configuration Automation. Double-click an available service to move it to the Selected Service field.

Selected Services

Lists the services that will be included in the graph filter.

Virtualization Filter

Specifies what combination of Non-virtual Servers, VM Hosting Servers, or VM Guest Servers is included in the graph.

For filtering purposes, VM Hosting Servers includes VM managers, and VM Guest Servers includes actual Virtual Machines running on a VM Host and VM Hosts that have a VM Manager assigned to it.

Server State

Specifies what combination of New, Managed, Unmanaged, and Imported servers is included in the graph. This option is *not* available on every template's Server tab.

Services tab

Available Services

Lists the services that are being managed by CA Configuration Automation. These services can be created in, discovered by, or imported into CA Configuration Automation. Double-click an available service to move it to the Selected Service field.

Selected Services

Lists the services that will be included in the graph filter.

3. Click OK.

   The visualized view resulting from applying the graph filter appears in the main visualization pane (the template name appears at the top of the pane) and the Overview pane. The Legend pane describes elements in the visualized view.

4. (Optional) Click any of the navigation icons to alter the view.

   The icons enable you to zoom in or out, pan, and display the icons and relationship lines in different arrangements. Refer to Graph Display Options (see page 388) for details about using the navigation icons.

Save a Template as a Custom Graph

After creating a custom template view as described in the previous sections, you can save the view as a custom graph so you can reuse it in the future.

To save a template view as a custom graph

1. Display a template view that contains the elements you want to save a a custom graph.
2. Click Save  in the row of buttons above the template view.

   The Save Graph dialog appears.

3. Enter the following information in the corresponding field:

   Name

   Specifies the name of the custom graph. The name you enter appears above the graph in the main visualization pane.

   Description

   Describes the purpose and content of the custom graph. The description appears above the graph in the main visualization pane (to the right of the name).

   Folder

   Specifies the folder in which to save the custom graph. You can enter any of the predefined folder names or enter a new folder name and it will be created.

   Shared Graph

   Specifies whether the custom graph is available for the user who created it (check box cleared), or to all users (check box checked).

4. Click OK.

   The template is saved as a custom graph in the folder you specified.

Working with Graphs

This section describes the visualization graph display options and the procedures that can be performed from the Visualization tab page.

Graph Display Options

This section describes the navigation and display options associated with visualization graphs. These options are activated by series of buttons located above the main visualization pane on the Visualization tab (accessed by clicking the Dashboard link, then the Visualization tab).

Button	Name	Description
	Select	Activates the selection tool after another mode is selected. This tool is active by default.
	Pan	Enables you drag-and-drop the graph view to reposition the view in the main visualization pane. The magnification is not changed.
	Marquee Zoom	Enables you to magnify a specific area of the graph by drawing a rectangle around the area.
	Interactive Zoom	Enables you to increase or decrease the magnification to show greater detail (fewer elements) or more elements (less detail). ■ Click and drag from the outer edge of the graph towards the center to zoom in. ■ Click and drag from the center of the graph towards the outer edge to zoom out.
	Fit	Increases or decreases the magnification to display the entire graph in the main visualization pane.
	Tree Layout	Displays the elements in a graph arranged in a hierarchical tree layout that attempts to depict the flow of the elements (for example, from top to bottom) and arranges similar elements in levels (horizontal rows).
	Symmetric Layout	Displays the elements in a graph arranged using quadrilateral symmetry, that is, elements divided equally on both the horizontal and vertical axes.
	Orthogonal Layout	Displays the elements in a graph arranged orthogonally, that is, at right angles to other elements.

	Circular Layout	Displays the elements in a graph arranged by determining the relationship of the elements, and arranges them as separate circles. The resulting circles are arranged in a radial tree layout fashion.
	Display Filter	Displays the Graph Filter dialog and the filter settings that were used to create the current view. You can edit the settings or save the template as a custom graph.
	Refresh	Redraws the current view.
	Save	Saves the graph using the same name. If you click Save with a template displayed, it is the equivalent of clicking Save As (that is, the Save As dialog appears).
	Save As	Displays the Save As dialog where you can specify the name, format, and location to save the current view as a graphic file.
	Print	Prints the current graph view.
	Print Preview	Displays a preview of the current graph as it will appear when printed.
	Print Settings	Displays the Print Setup dialog where you can configure printer settings.
	Save As Image	Displays the Save Graph dialog that enables you to save the current view as a custom graph. For a description of the fields, see Save a Template as a Custom Graph (see page 387).
	Hide Tree	Hides or displays the tree pane depending on the current setting.
	Display Edge Labels	Displays text labels over the connecting lines in those graph that use them to show relationships.
 	Expand Nodes Collapse Nodes	Expands or collapses all nodes in the graph elements depending on the current setting. ■ If the nodes are currently expanded, the collapse node icon (the minus sign) appears on the button. ■ If the nodes are currently collapsed, the expand node icon (the plus sign) appears on the button. **Note:** You can use the Select tool to expand or collapse individual nodes by clicking the expand node icon (plus sign) or the collapse node icon (the minus sign) next to individual elements in the main visualization pane.

	Set Filter Defaults	Displays the Set Template Filter Defaults dialog with the the default settings displayed. You can create a new filter, or use the default filter settings to filter the graph view.

Display Graph Element Details

Graphs include various graphical elements that depict your servers, services, applications, relationships, and so on. The Visualization page contains a legend in the right pane that describes the elements in the current graph. Additionally, you can expand nodes on certain elements or mouse-over them to display additional details.

The following procedure describes how to view the details about an application element and the server in which it is installed. It also includes an example of the workflow used to filter the view to only the servers with a specific application installed, and to optionally save that view.

To expand an element to display details

1. Click the Dashboard link (top right), then the Visualization tab.

   The Visualization tab page appears.

2. Expand one of the four predefined sub-folders in the Graphs folder, then select a graph. This example uses the Applications sub-folder and the All Servers by Installed Application graph.

   The main visualization pane displays all servers known to CA Configuration Automation that have applications installed (likely all of them).

3. Click the Display Filter  button above the graph.

   The Graph Filter dialog appears.

4. Create a filter that shows a specific application on a smaller set of servers, for example, all Linux servers with Mozilla Firefox:

   a. On the Servers tab on the Graph Filter dialog, click the double right-facing arrow to move all of the servers in the Available Servers column to the Selected Servers column.

   b. Click the Operating System tab, then double-click Linux in the Available Operating System column.

      Linux appears in the Operating System column.

c. Click the Applications tab, then double-click the application you want to include in the graph. In this example, select Mozilla Firefox.

There may be multiple versions of the application (3.6.13, 3.6.14, and so) installed on servers, so select all that you want included. The select versions of the application appear in the Selected Applications column.

d. Click OK.

The graph is filtered to only display servers with the selected application installed:

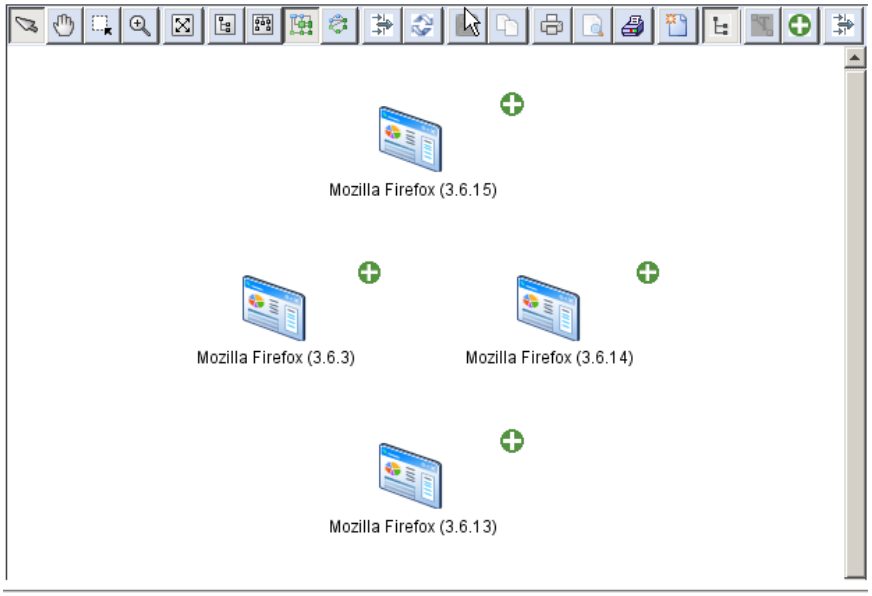

5. Do one or more of the following:

- Mouse-over an application icon to display details about the application.
- Click an application icon to display a navigation icon that you can double-click to display the child element of this application (in this example, the server that hosts it). Right-click the server icon and select Go To Parent to return to the previous view.
- Right-click an application icon to display the filter that created the current graph. You can edit the filter if you want to change the contents of the graph.
- Click the Expand Nodes button above the graph to expand *all* nodes, or double-click an individual node icon in the graph to expand an individual node.

  Server icons representing the servers where the application is installed appear.

You can mouse-over a server icon to display server details (including the assigned Management Profile, status IP address, operating system, and so on) or right-click to display relationship details, or perform server management operations (for example, Create Snapshot, Run Change Detection, and so on).

- Click the Save As button  above the graph to save the current view as a custom graph.

Set a Graph as the Default

You can select a predefined or custom graph to appear by default when you open the Visualization tab page.

To set a graph as the default

1. Click the Dashboard link, then the Visualization tab.

   The Visualization tab page appears.

2. Navigate to the graph you want to designate as the default graph, right-click on it, and then select Set Default.

   A confirmation message similar to the following appears:

   Confirmation: CCA-GR-9806: Default graph set to "All Service Communication Relationships."

   When you return to the Visualization tab page, the selected graph appears by default.

Export a Graph

You can export a predefined or custom graph to JAR file.

To export a graph

1. Click the Dashboard link, then the Visualization tab.

   The Visualization tab page appears.

2. Navigate to the graph you want to export, right-click on it, and then select Export Graph.

   The File Download dialog appears.

3. Click Save, specify a name for the file or accept the default name, and then specify a location to save the file.

   The file is saved in the specified location.

Edit Graph Properties

You can view the properties of a predefined graph, or edit certain properties of graphs you created.

To view or edit the properties of a graph

1. Click the Dashboard link, then the Visualization tab.

   The Visualization tab page appears.

2. Navigate to the graph you want to view or edit, right-click on it, and then select Properties.

   The Properties dialog appears.

3. View or edit one or more of the following fields, then click OK.

   Note: You can only edit the properties of graphs of which you are the owner (that is, where your user name is listed in the Owner field).

   Name

   Specifies the name of the graph. This is a read-only field for all graphs.

   Description

   Describes the function of the graph. The description entered in this field appears after the name of the graph in the main visualization pane, and when you mouse-over the graph name in the tree pane.

   Owner

   Specifies the user name of the person who created the graph. The predefined graphs are owned by system_user.

   Shared Graph

   Specifies whether the graph is viewable by all users (check box checked), or only the owner (check box clear).

   If edits were made, the graph is saved; the dialog closes.

Export Dashboards and Visualization Objects to Tenants

CA Configuration Automation tenant administrators can export dashboards and visualization objects from any CA Configuration Automation instance to a tenant instance on their master instance. You can perform the export from the following places:

- Dashboards tree on the Charts tab
- Graphs tree on the Visualization tab.

Similarly, the Export to Tenants functionality is available for profiles and other objects as described in Export Profiles and Objects to Tenants (see page 97).

Follow these steps:

1. Log in to the CA Configuration Automation Server UI as a tenant administrator user, then click the Dashboard link.

   The Dashboard panel appears.

2. Click the Charts or Visualization tab depending on the objects that you want to export.

   The corresponding tab page appears.

3. Do one of the following:

   - On the Charts tab page, click the plus sign (+) next to the Dashboards folder, right-click a dashboard name, then select Export to Tenant from the menu.
   - On the Visualization tab page, click the plus sign (+) next to the Graphs folder, click the plus sign (+) next a sub-folder, right-click a graph name, then select Export to Tenant from the menu.

   The Import dialog appears.

4. Click the Overwrite Existing Objects check box if you want to overwrite the existing object with the new version on the tenant instance.

5. Double-click one or more tenants in the Available Tenants column.

   The selected tenant appears in the Selected Tenants column.

6. Click OK to the import the selected objects into the selected tenants.

   The Results pane confirms the import was successful, or displays an error description.

Chapter 17: Tasks Panel

Use the Tasks panel to complete the following common tasks:

- Discover Network
- Access Profile and Agent Deployment
- Discover Service
- Run Compliance Job
- Locate and Upgrade Agents

Click a task to open a wizard that contains a detailed description of the task and navigation buttons that link to the subtasks that are required to complete the task.

Chapter 18: Understanding and Creating Rules

Constraint rules are used to place value constraints on particular types of CA Configuration Automation elements, including the following:

- Rules in the Indicators, Verification Rules folder (Component Blueprints only)
- Parameters in the Parameters folder (Services only)
- Files and directories in the Managed, File System Overlay
- Registry keys and values in the Managed, Registry Overlay folder
- Parameters in the Parameter, Rules folder
- Parameters in the Configuration, Structure Classes folder
- Files in the Configuration folder

Constraint rules are always associated one-to-one with an element and can be either created in and inherited from the underlying Component Blueprint or created in and applied directly to a service instance. In addition to the *explicit* constraint rules that you create, there are also built-in CA Configuration Automation *implicit* constraint rules. For example, if you specify a particular value or data type for an element, CA Configuration Automation automatically creates an implicit Check Default or Verify Data Type rule.

Note: When possible, consider creating constraint rules within the Component Blueprint. You create the rule once, and it is automatically inherited by any service that uses the underlying Component Blueprint.

Constraint rules are initiated, viewed, and edited from the Rules field of the selected Component Blueprint.

The number of defined constraint rules or None is displayed in the brackets to the right of Rules. From the drop-down list, you can view a specific constraint rule, view all constraint rules, or add constraint rules.

- If you select Show All Rules, a new page is displayed showing a list of all constraint rules in table format. You can change the sort order of the table by clicking any of the table headings.
- If you select New Rule, the Rule attribute sheet is displayed.
- If you select a specific constraint rule, the name is displayed in the Rules field and the attribute sheet for that constraint rule is displayed. You can edit or delete most rules from the attribute sheet. The exception is that you cannot edit or delete CA Configuration Automation built-in or *implicit* rules, such as Check Default or Verify Data Type. You can only view these implicit rules.

When a default value is specified in a Component Blueprint, an implicit Check Default rule is automatically created. These default rules are informational (show in Rule Compliance as Information, with an icon), and let you know when a value is deviating from the recommended default value.

Explicitly defined constraint rules derive their value from the element itself or from one of the element's attributes. The available attributes and the value type vary for each element within CA Configuration Automation, so constraint rules vary significantly by element type. For example:

Element Type	**Allowable Constraint Rules**
File	■ File Size ■ File Modification Time ■ File Owner ■ File Permissions ■ File Version ■ Product Version
Directory	■ Number of directories ■ Number of files ■ Bytes ■ Depth ■ Directory Modification Time ■ Directory Owner ■ Directory Permissions ■ Directory Must Exist ■ File Must Exist
Rules, Parameters, Groups, Registry Keys and Values	■ Value

Appendix A: Using the Interpreted Cluster

With the new interpretation type named interpreted cluster, you can define the cluster name to form a relationship between cluster nodes and the respective target database instances. The cluster discovery information, by default is available for Oracle 11g Database Instance (UNIX) and Oracle 10g Database Instance (UNIX) blueprints on AIX and RedHat cluster platforms.

Follow these steps:

1. Create a network profile which includes AIX cluster nodes, service hostname or service IP with the correct credential vault details associated to it.

   Credential Vault of target AIX cluster nodes is not required as SNMP can get the required cluster relationship information.

2. Run the network profile.

   You can get the Cluster nodes and Service Hostname added in the Servers section and Cluster information is available in the Clusters section.

3. Create management profile by selecting the required blueprint pair:

   - Oracle Database 11g (UNIX), Oracle 11g Database Instance (UNIX)
   - Oracle 10g Database (UNIX), Oracle 10g Database Instance (UNIX)

4. Assign this management profile to both the Cluster nodes which are added in Servers section.

   The discovered Oracle 10g or 11g components displays the configuration relationships with the new cluster name added in the Cluster Name column. You can export this relationship to CA CMDB using the catalyst server.

Chapter 19: Understanding and Creating Directives

Directives are used to either extract values from elements managed within a service, or to retrieve values from managed servers using the CA Configuration Automation Agent. The following list introduces the four directive types:

- Verification directives—Eliminate initially discovered components that are partially installed, of the wrong version, or not of interest to a particular service.
- Parameter directives—Define parameters that are critical for locating and identifying the component, such as the file system or registry root, component version, vendor, and database connection information.
- Executable directives—Define directives to extract and interpret configuration information from a server.
- Macro step directives—Help diagnose problems specific to the servers containing the data being managed by its Component Blueprint and provide additional information about a server or service, such as viewing system information, memory statistics, or disk volume statistics.

The sections that follow describe the directives in detail.

Verification Directives

To initiate verification-related value directives in the Component Blueprint, select the Indicators, Verification Directives element and click Add Directive. View and edit existing verification directives from the Verification Directive tree that appears when you select a verification directive element.

The fields that are displayed vary by the Directive Type that you select.

Examples:

- The Oracle 8i Database (UNIX) v8.* Component Blueprint defines a Remote Execution verification directive that uses SQL Plus to retrieve the version and verifies that the discovered component is version 8.
- The Apache HTTP Server (UNIX) v1.3.* Component Blueprint defines a Constant verification directive that retrieves the version information and verifies that the discovered component is version 1.3.*.

  Note: The example uses a wildcard (*) when specifying the regular expression on which to match. Discovery finds all versions that start with 1.3.

Parameter Directives

Parameter-related value directives are initiated in the Component Blueprint by selecting the Parameters, Directives element and clicking Add Directive. Existing parameter directives are viewed and edited from the Parameter tree that appears when you select a parameter directive element.

The fields displayed vary by the Directive Type selected.

Examples:

- The WIN32 v*.* Component Blueprint defines a Constant parameter directive that exposes the product name and associated Service Pack version in the parameters of a discovered Windows operating system component in a service.
- The same WIN32 v*.* Component Blueprint defines a Constant parameter directive that exposes the Vendor in the parameters of a discovered Windows operating system component in a service.

Configuration Executable Directives

Configuration executable-related value directives are initiated in the Component Blueprint by selecting the Configuration, Executables element and clicking Add Directive. Existing executable directives are viewed and edited from the Executable Directive tree that appears when you select a configuration executable directive element.

The fields displayed vary by the Directive Type selected.

Examples:

- The Active Directory Service v*.* Component Blueprint defines a Get LDAP configuration executables directive that retrieves the FSMO Roles.
- The BIG-IP Load Balancer v*.* Component Blueprint defines a Get SNMP configuration executables directive that retrieves values at the specified MIB address.

Macro Step Directives

Macro step-related value directives are initiated in the Component Blueprint by selecting a macro element under Diagnostics, Macros or Utilities, or Macros, and then clicking Add Step. Existing macro steps are viewed and edited from Macro Step tree that appears when you select a macro step element.

The fields displayed vary by the Directive Type selected.

Examples

- The IBM WebSphere 6 Server Instance (UNIX) v6.* Component Blueprint defines a Remote Execution macro step directive that starts WebSphere Server with the help option.
- The BIG-IP Load Balancer v*.* Component Blueprint defines a Get SNMP macro step directive that retrieves the total uptime values at the specified MIB address.

Appendix B: Configuring sudo for UNIX and Linux Softagent Discovery

When using the NDG Softagent to discover UNIX and Linux servers, NDG attempts to establish an SSH connection to the UNIX and Linux hosts using the set of credentials provided in the credential vault. Depending how your UNIX/Linux security is configured, it is possible that some commands issued by the NDG Softagent cannot be authorized for the non-root user, resulting in less data being discovered for the server.

You have the following options to avoid having a non-root user issue discovery-related commands:

- Provide root user credentials in the credential vault so all discovery-related commands are issued as root. This ensures the commands are authorized.
- Use the sudo command to enable a non-root user to issue discovery-related commands under the authority of root, without having to supply root's credentials.

You also have to define a path for the userid that is associated with the sudo user that includes all the locations for the following commands and utilities that NDG discovery uses:

- /bin
- /sbin
- /usr/sbin
- /opt/xensource/bin

To configure the /etc/sudoers file to use sudo to authorize non-root users

1. Edit the /etc/sudoers file using the visudoers command.
2. Create the following entry for the user ndguser to issue all NDG Softagent commands using sudo without prompting for root credentials:

```
# simple entry for ndg discovery if client does not need granularity
# ndguser ALL=NOPASSWD: ALL
# detailed entry for ndg discovery permitting only those commands used by discovery

ndguser ALL = NOPASSWD: /bin/uname, /bin/echo, /bin/cat, \
                        /bin/domainname, /bin/hostname, \
                        /bin/netstat, /bin/df, /bin/ps, /bin/rpm, \
                        /bin/ls, /sbin/ifconfig, /sbin/ip, \
                        /sbin/mii-tool, /sbin/chkconfig, \
                        /sbin/sfdisk, /usr/sbin/dmidecode, \
                        /usr/bin/cdrecord, \
                        /opt/xensource/bin/xe, /bin/lshmc
```

   Note: You can modify this entry to authorize an existing user instead of creating ndguser. If your system already has a user configured in the /etc/sudoers file to issue all commands without password prompting, or a granular list that contains all of the commands shown, that user can be used without any modifications by adding this user to your credential vault.

3. Save and close the sudoers file.
4. Click the Enable use of sudo check box on the Network Scan Policy page as described in Create a Network Scan Policy.

Define the path for the sudo user

1. Edit the shell configuration file for your UNIX or Linux system's shell (typically, .bashrc in the user's $HOME directory), and add the following lines to the user's PATH definition:

```
PATH=$PATH:/bin:/sbin:/usr/sbin:/opt/xensource/bin
export PATH
```

2. Save and close the file.

Appendix C: Configuring Telnet Access Mode for Component Discovery

Telnet discovery lets you discover components in the CA Configuration Automation network. Telnet discovery has the following requirements:

- The ftpd daemon and Telnet must be running on the Grid Node server.
- An FTP client and Telnet must be available on the managed or target server that the FTP daemon uses.

When you initiate Telnet discovery, CA Configuration Automation runs the Discovery Job on the Grid Node server. The Grid Node server starts a sequence of Telnet communications with the managed server. Following a Telnet negotiation, a Telnet command uses FTP to copy Discovery scripts to the managed server. The Discovery commands use Telnet to run on the managed server. CA Configuration Automation captures the output in a file and copies it to the Grid Node server.

Follow these steps:

1. On the Servers tab under Management, click the Access Profile link.

   Note: As a best practice, make a copy of a predefined access profile before you modify it.

2. On the Access Profile page, click the Telnet link.
3. On the Details for Profile: Telnet page, select Access Mode and define values in the required fields.

   Notes:

   - The Connection Timeout value is important to the Telnet access mode and CA Configuration Automation sets it to a high value by default. CA Configuration Automation uses the Connection Timeout value to stop idle connections.
   - The Prompts field values locate and identify the prompts that CA Configuration Automation displays while establishing a Telnet session with the targeted server.

This section contains the following topics:

Modify CA Configuration Automation Configuration Properties for Telnet

You can modify the Telnet configuration properties.

Follow these steps:

1. Click the Administration link.
2. On the Properties page of the Configuration tab, click > to display the next properties page.
3. Edit the following cca group properties in the Value column:

   telnet.connection.retries

   Defines the number of times to retry the Telnet connection process if it fails during discovery.

   telnet.read.timeout_secs

   Defines the maximum time to wait to gather the results after issuing a command.

   telnet.read.byte_to_byte_delay_secs

   Defines the maximum time to wait for the next byte while reading the results. CA Configuration Automation uses this value only if Look for Prompts is not selected in the access profile.

Configure Grid Node Server Properties for FTP

You can configure Grid Node server properties for FTP.

Follow these steps:

1. Click the Administration link.
2. On the Configuration tab, click the Properties link.

3. On the Properties page of the Configuration tab, click > to display the next properties page.
4. Edit the following grid group properties in the Value column:

   ftp.account

   Defines the account name with which to connect to the FTP Server.

   ftp.password

   Defines the password that is associated with the specified ftp.account value.

   ftp.port

   Defines the port number on which the FTP Server listens.

   ftp.root

   Defines the root path of the FTP Server (that is, the FTP site directory).

Appendix D: Mapping CA Configuration Automation Tasks to CA EEM Permissions

This appendix maps the CA Configuration Automation UI tasks to the CA EEM permissions required to perform that task. It describes CA Configuration Automation tasks by the UI location and which task to select from the Table Actions drop-down Select Actions drop-down list.

The CA EEM permissions are described in terms of the Policy and the corresponding Action.

Notes:

- All users have view access to the entire CA Configuration Automation UI
- An error message appears when the authorization fails for an operation
- CA EEM performs the policy evaluation
- You can assign policies to users directly, or to user groups

This section contains the following topics:

Service Options

Service Options	CA EEM Policy and Action
Delete Services	Service Management, delete
Take Snapshot	Service Management, create
Run Change Detection	Service Management, run_change_detection
Run Compare	Service Management, run_compare
Run Rule Compliance	Service Management, run_rule_compliance
Refresh Services	Service Management, refresh
Run Discovery	Service Management, run_discovery
Stop Discovery	Service Management, stop_discovery
Run Management Profile	Service Management, run_management_profile
Assign Management Profile	Service Management, update
Export Services	Service Management, export
View all Services	View permissions are granted to all users
Create Service	Service Management, create
Update Service	Service Management, update
Import Service	Service Management, import

Service Snapshot Options

Service Snapshot Options	CA EEM Policy and Action
View all Snapshots	View permissions are granted to all users
Delete Snapshots	Service Management, delete
Set as Gold Standard	Service Management, update
Set as Silver Standard	Service Management, update
Set as Bronze Standard	Service Management, update
Set as Baseline	Service Management, update
Remove Gold Standard	Service Management, update
Remove Silver Standard	Service Management, update

Remove Bronze Standard	Service Management, update
Remove Baseline Designation	Service Management, update
Export Snapshots	Service Management, export

Service Component Options

Service Component Options	CA EEM Policy and Action
Delete Components	Server Management, delete
Refresh Components	Server Management, update
View Components	View permissions are granted to all users

Server Options

Server Options	CA EEM Policy and Action
Delete Server	Server Management, delete
Manage Servers	Server Management, update
Reject Servers	Server Management, update
Test Servers	Server Management, test_servers
Take Snapshot	Server Management, create
Run Change Detection	Server Management, run_change_detection
Run Compare	Server Management, run_compare
Run Rule Compliance	Server Management, run_rule_compliance
Refresh Servers	Server Management, refresh
Run Discovery	Server Management, run_discovery
Stop Discovery	Server Management, stop_discovery
Run Management Profile	Server Management, run_management_profile
Assign Profiles	Server Management, update
Secure Agents	CCA Admin Access, update and Server Management, install_uninstall_agent
Install Agents	Server Management, install_uninstall_agent

Uninstall Agents	Server Management, install_uninstall_agent
View all Servers	View permissions are granted to all users
Create Server	Server Management, create
Add Server from File	Server Management, create
Update Server	Server Management, update
Import Server	Server Management, create

Server Snapshot Options

Server Snapshots Options	CA EEM Policy and Action
Delete Snapshots	Server Management, delete
Set as Gold Standard	Server Management, update
Set as Silver Standard	Server Management, update
Set as Bronze Standard	Server Management, update
Set as Baseline	Server Management, update
Remove Gold Standard	Server Management, update
Remove Silver Standard	Server Management, update
Remove Bronze Standard	Server Management, update
Remove Baseline Designation	Server Management, update
Export Snapshots	Server Management, export
Import Snapshots	Server Management, import
View all Snapshots	View permissions are granted to all users

Server Component Options

Server Component Options	CA EEM Policy and Action
Delete Components	Server Management, update
Refresh Components	Server Management, update

View Components	View permissions are granted to all users

Server Group Options

Server Group Options	CA EEM Policy and Action
Create Server Groups	Server Management, create
Update Server Groups	Server Management, update
View Server Groups	View permissions are granted to all users

Management Profile Options

Management Profile Options	CA EEM Policy and Action
Set As Default Profile	Management Profile Management, update
Enable Profile	Management Profile Management, update
Disable Profile	Management Profile Management, update
Delete Profile	Management Profile Management, delete
Create Profile*	Management Profile Management, create
Update Profile	Management Profile Management, update
Export Profile	Management Profile Management, export
Import Profile	Management Profile Management, import
Run Profile	Server Management, run_management_profile and Service Management, run_management_profile
View Profiles	View permissions are granted to all users

Network Profile Options

Network Profile Options	CA EEM Policy and Action
Set As Default Profile	Network Management, update
Enable Profile	Network Management, update

Disable Profile	Network Management, update
Delete Profile	Network Management, delete
Create Profile	Network Management, create
Update Profile	Network Management, update
View Profiles	View permissions are granted to all users

Network Scan Policy Options

Network Scan Policy Options	CA EEM Policy and Action
Delete Network Scan Policy	Network Management, delete
Create Network Scan Policy	Network Management, create
Import Network Scan Policy	Network Management, import
Export Network Scan Policy	Network Management, export
Update Network Scan Policy	Network Management, update
View Network Scan Policy	View permissions are granted to all users

Access Profile Options

Access Profile Options	CA EEM Policy and Action
Delete Access Profile	Access Profile Management, delete
Create Access Profile	Access Profile Management, create
Import Access Profile	Access Profile Management, import
Export Access Profile	Access Profile Management, export
Update Access Profile	Access Profile Management, update
View Access Profile	View permissions are granted to all users

Credential Vault Profile Options

Credential Vault Profile Options	CA EEM Policy and Action
Set As Default Profile	Network Management, update
Delete Credential Vault Profile	Network Management, delete
Create Credential Vault Profile	Network Management, create
Update Credential Vault Profile	Network Management, update
View Credential Vault Profile	View permissions are granted to all users

Notification Profile Options

Notification Profile Options	CA EEM Policy and Action
Set As Default Profile	Notification Profile Management, update
Delete Notification Profile	Notification Profile Management, delete
Create Notification Profile	Notification Profile Management, create
Update Notification Profile	Notification Profile Management, update
View Notification Profile	View permissions are granted to all users

Blueprints Options

Blueprints Options	CA EEM Policy and Action
Copy Blueprint	BlueprintsManagement, copy
Delete Blueprint	BlueprintsManagement, delete
Enable Discovery	BlueprintsManagement, update
Disable Discovery	BlueprintsManagement, update
Export Blueprint	BlueprintsManagement, export

Import Blueprint	BlueprintsManagement, import
Create Blueprint	BlueprintsManagement, create
Update Blueprint	BlueprintsManagement, update
View Blueprint	View permissions are granted to all users

Structure Class Options

Structure Class Options	CA EEM Policy and Action
Copy Structure Class	BlueprintsManagement, create
Delete Structure Class	BlueprintsManagement, delete
Create Structure Class	BlueprintsManagement, create
Import Structure Class	BlueprintsManagement, import
Export Structure Class	BlueprintsManagement, export
Update Structure Class	BlueprintsManagement, update
View Structure Class	View permissions are granted to all users

Global Variable Options

Global Variable Options	CA EEM Policy and Action
Delete Global Variables	BlueprintsManagement, delete
Create Global Variables	BlueprintsManagement, create
Import Global Variables	BlueprintsManagement, import
Export to Excel	BlueprintsManagement, export
Update Global Variables	BlueprintsManagement, update
View Global Variables	View permissions are granted to all users

Compliance Management Options

Compliance Management Options	CA EEM Policy and Action
Create Compliance Profile	Compliance Management, create
Delete Compliance Profile	Compliance Management, delete
Update Compliance Profile	Compliance Management, update
Run Job	Compliance Management, run_job

Dashboard Options

Dashboard Options	CA EEM Policy and Action
Create Dashboards	Dashboard Management, create
Import Dashboards	Dashboard Management, import
Export Dashboards	Dashboard Management, export
Update Dashboards	Dashboard Management, update
Delete Dashboards	Dashboard Management, delete
View Dashboards	View permissions are granted to all users

Remediation Options

Remediation Options	CA EEM Policy and Action
Allow Remediation	Remediation Management, allow

Report Options

Report Options	CA EEM Policy and Action
Run Reports	Report Management, run
Save Reports	Report Management, create

Update Reports	Report Management, update
Delete Reports	Report Management, delete
Schedule Reports	Report Management, run
View Saved Reports View Report Templates	View permissions are granted to all users

Administration Options

Administration Options	**CA EEM Policy and Action**
Access Management	CCA Admin Access, update
Configuration, Security Certificates	CCA Admin Access, update

Appendix E: Using the Command-line Interface

CA Configuration Automation includes a command-line interface (CLI) that enables you to incorporate server- and service-related operations into scripts and other administrative processes within your infrastructure.

The CLI is installed automatically when you install a CA Configuration Automation Agent, but you can also install the CLI on any server for integration into existing scripts and processes. Follow the installation instructions provided in the *CA Configuration Automation Implementation Guide* to ensure access to these commands.

The CLI runs on all CA Configuration Automation Agent-supported platforms. See the *CA Configuration Automation Release Notes* for platform and version support details.

Note: Before using the CLI, you must successfully log into the CA Configuration Automation UI at least once to verify user authentication.

ccautil

The ccautil executable is the basis for all CLI operations. The arguments you supply with the executable control which operations are performed. The ccautil executable (ccautil.bat for Windows and ccautil.sh for UNIX and Linux) is located on the CA Configuration Automation Server in the *<CA Configuration Automation Server_install_directory>*\bin directory.

You can use ccautil to perform the following operations:

- Query the CA Configuration Automation Server to obtain service information, such as service names and service UUIDs, and to get server information such as server names and server UUIDs, without having to use the browser-based UI
- Create and store a hashed CA Configuration Automation password in a named password file
- Initiate a server or service Refresh operation with additional options to do the following:
  - Create a Snapshot
  - Run Change Detection between current data to the last Snapshot created
  - Run Rule Compliance against a current or Snapshot data
  - Send an email notification of job completion, job failure, and Change Detection or Rule Compliance failures

- Run customized reports
- Import a server data file
- Create and install a new SSL certificate for securing CA Configuration Automation Agent communications
- Assign or update:
  - Access Profile to Server
  - Management Profile to Server
  - Management Profile to Service

Usage Notes

Consider the following general points when using the CLI:

- Most ccautil commands require a -task argument, the CA Configuration Automation Server identification, and user name and password arguments.
- Optional arguments are enclosed in square brackets [ ].
- The ccautil arguments can be ordered in any sequence.
- The CLI generates error messages when you use incorrect or incomplete syntax.
- If you specify argument text that contains special characters or embedded blanks, you must enclose the entry in quotation marks.

  For example, if the location of a password file is in the Program Files directory or the name of a service is San Francisco Datacenter, you need to enclose these entries in quotation marks (for example, "Program Files" and "San Francisco Datacenter").

Execute ccautilTasks

This section provides the syntax and argument descriptions for the following ccautil tasks:

- List ccautil help
- List Services
- List Servers
- Hash Password
- Refresh Service or Server

- Run Report
- Import Server File
- Secure Agent
- Assign Profile

Follow these steps:

1. Open a command prompt.

   The command window appears.

2. Change to the *<CA Configuration Automation Server_install_directory>*\bin directory directory and execute the ccautil command as follows:

   `ccautil <task>`

   The options to perform each task are described in the sections that follow.

Display ccautil Help

The following command generates the list of all ccautil options:
`ccautil -h`

-h

Displays a list of ccautil arguments and a brief description of what each one does.

List Services Option

The listservices command produces a list of services and their internal identifiers, which can be used to uniquely reference services for other ccautil commands:

```
ccautil -task listservices
-s <server:port_number>
-u <user_name>
-p <password>
-pwfile </path/filename>
```

-s <*server:port_number*> or -s https://<*server:port_number*>

Specifies the HTTP or HTTPS CA Configuration Automation Server name and port number from which you want to obtain the list of services.

-u <*user_name*>

Specifies a valid CA Configuration Automation Server user name.

-p <*password*>

Specifies the user's password. The password shows as clear text on the command line. Alternatively, you can create and use a scrambled password file for authentication. For more information, see Hash Password Option (see page 425).

-pwfile <*/path/filename*> (on UNIX servers) or -pwfile <*pathfilename*> (on Windows servers)**

Specifies the user's password file. You can use this argument instead of -p *password*.

Usage Examples

- The following example uses the -p argument to specify the password on the command line:

  ```
  ccautil -task listservices -S qaserver1:8080 -u psmith -p user250
  ```

- The following example uses the -pwfile argument to specify a password file on a UNIX server:

  ```
  ccautil -task listservices -S qaserver1:8080 -u psmith -pwfile
  /home/psmith/ccapwd
  ```

- The following example uses the -pwfile argument to specify a password file on a Windows server:

  ```
  ccautil -task listservices -S qaserver1:8080 -u psmith -pwfile "\Program
  Files\CA\ccapwd"
  ```

  Notice the path and file name for the password file is enclosed in quotes because the directory Program Files contains a space.

List Servers Option

The listservers command produces a list of servers and their internal identifiers, which can be used to uniquely reference servers for other ccautil commands:

```
ccautil -task listservers
-s <server:port_number>
-u <user_name>
-p <password>
-pwfile </path/filename>
```

-s <*server:port_number*>

Specifies the CA Configuration Automation Server name and port number from which you want to obtain the list of services.

-u <*user_name*>

Specifies a valid CA Configuration Automation Server user name.

-p <*password*>

Specifies the user's password. The password shows as clear text on the command line. Make sure to clear your screen after using the CLI directly at the command line prompt. Alternatively, you can create and use a scrambled password file for authentication. For more information, see the Hash Password Option.

-pwfile <*/path/filename*> (on UNIX servers) or -pwfile <*pathfilename*> (on Windows servers)**

Specifies the user's password file. You can use this argument instead of -p *password*.

Usage Examples

- The following example uses the -p argument to specify the password on the command line:

  ```
  ccautil -task listservers -S qaserver1:8080 -u psmith -p user250
  ```

- The following example uses the -pwfile argument to specify a password file on a UNIX server:

  ```
  ccautil -task listservers -S qaserver1:8080 -u psmith -pwfile /home/psmith/ccapwd
  ```

- The following example uses the -pwfile argument to specify a password file on a Windows server:

  ```
  ccautil -task listservers -S qaserver1:8080 -u psmith -pwfile "\Program Files\CA\ccapwd"
  ```

  Notice the path and file name for the password file is enclosed in quotes because the directory Program Files contains a space.

Hash Password Option

The hash password option lets you create and store a scrambled or hashed CA Configuration Automation Server password in a named password file. Stored passwords can be used in subsequent invocations of ccautil. This option lets you execute commands without the password appearing in clear text on the command line.

The following command hashes the supplied password into the specified password file:

```
ccautil -task hashpw -p <password> -newpwfile [set the File Name variable]
```

-p *<password>*

Specifies the user's password. Your password shows as clear text on the command line. Use the-task hashpw command to create and use a scrambled password file for authentication in other ccautil commands.

-newpwfile *</path/filename>* (on UNIX servers) or
-newpwfile *<\path\filename>* (on Windows servers)

Creates a password file containing the hashed password created from the required -p argument, at the specified location with the specified name. This is a required argument.

Note: Make sure that you have write permission to the path you specify for the password file.

Usage Example

The following example creates a password file containing a scrambled version of the user's user250 password:

```
ccautil -task   hashpw -p    user250 -newpwfile    /home/bsmith/ccapwd
```

Refresh Service or Server Option

The following command refreshes the identified service or server and directs ccautil to perform other optional tasks (listed in square brackets):

```
ccautil -task refresh
-s <server:port_number>
-u <user_name>
-p <password> or -pwfile </path/filename>
-service
-server
-id <service/server_UUID>
-servicename <service_name>
-servername <server_name>
[-job <refresh_jobname>]
[-snap]
[-cd]
[-rc]
[-date <mm/dd/yyyy>]
[-time <hh:mm>]

[-delay <secs>]
[-n error or -n completion]
[-email <email_account@your_company.com>]
[-emailsubject <email_subject>]
```

-s <*server:port_number*>

Specifies the CA Configuration Automation Server name and port number for which you want to refresh a service.

-u <*user_name*>

Specifies a valid CA Configuration Automation Server user login name.

-p <*password*>

Specifies the user's password. The password shows as clear text on the command line. Alternatively, you can create and use a scrambled password file for authentication. For more information, see Hash Password Option (see page 425).

-pwfile <*/path/filename*> (on UNIX servers) or -pwfile <*pathfilename*> (on Windows servers)**

Specifies the user's password file. You can use this argument instead of -p *password*.

-id <*service_uuid*> or <*server_uuid*>

Specifies the service or server ID number returned from the -task listservices query operation of the service or server you want to refresh. Use this argument *and* the -servicename or -servername arguments to uniquely identify the service or server you want to refresh.

-servicename <*service_name*>

Specifies the name of the service you want to refresh.

-servername <*server_name*>

Specifies the name of the server you want to refresh.

[-job <*refresh_jobname*>]

Specifies a job name for the Refresh operation. If no name is specified, the CLI uses refresh_jobname, ccautil.

[-snap]

Creates a Snapshot of the current data.

[-cd]

Runs Change Detection between the current data and the most recent Snapshot.

[-rc]

Runs Rule Compliance to verify service and Blueprint rules on the current data.

[-date *mm/dd/yyyy*]

Specifies the date the Refresh operation is to run. If no date is specified, the CLI uses today's date.

[-time *hh:mm*]

Specifies the time (24-hour clock) the Refresh operation is to run. If no date is specified, the CLI uses the current time.

[-delay *secs*]

Specifies the number of seconds to delay the start of the Refresh operation. You can specify any integer.

Note: This feature is unique to the CLI and is not available in the CA Configuration Automation Server GUI.

[-n]

Specifies if a Refresh job completion or failure notification is to be sent, depending on how you append the argument. If notification is to be sent, you must also specify the -email argument with a valid email account. For more information, see the details for each of these arguments.

[-n error]

Sends email notification only if remediations or rule violations are detected.

[-n completion]

Sends email each time a Refresh operation completes.

[-email <*email_name@company.com*>]

Specifies that an email notification be sent and specifies the email account to send the notification to. Specify this argument with the -n argument.

[-emailsubject <*email_subject*>]

Specifies the subject line of any sent email notifications.

Usage Example

- The following example uses the -servicename argument to uniquely identify the service. The other arguments direct the utility to Refresh the service and take a Snapshot at 11:00 PM on the specified date. Upon job completion, an email notification will be sent to cca@ca.com:

  ```
  ccautil -task refresh -service -servicename "service1" -snap -date 06/15/2009 -time 23:00 -n completion -email cca@ca.com -S ccaservername:8080 -u username –p password
  ```

- The following example uses the -id argument to specify the service. The other arguments direct the utility to Refresh the service 60 seconds from now, and take a Snapshot, run Change Detection, and run Rule Compliance. Upon job completion, an email notification will be sent to cca@ca.com:

  ```
  ccautil -task refresh -service -id 5dfe7ce5-fbed-44b1-80ef-1a49f5780236 -snap –cd -rc -delay 60 -n completion -email cca@ca.com -S ccaservername:8080 -u username –p password
  ```

- The following example uses the -servername argument to uniquely identify the server. The other arguments direct the utility to Refresh the server and take a Snapshot at 11:00 PM on the specified date. Upon job completion, an email notification will be sent to cca@ca.com:

  ```
  ccautil -task refresh -server -servername "server1" –snap -date 06/15/2009 -time 23:00 -n completion -email cca@ca.com -S ccaservername:8080 -u username –p password
  ```

- The following example uses the -id argument to specify the server. The other arguments direct the utility to Refresh the server 60 seconds from now, take a Snapshot, run Change Detection, and then run Rule Compliance. Upon job completion, an email notification will be sent to cca@ca.com:

  ```
  ccautil -task refresh –server -id 5dfe7ce5-fbed-44b1-80ef-1a49f5780236 -snap –cd -rc -delay 60 -n completion -email cca@ca.com -S ccaservername:8080 -u username –p password
  ```

Run Report Option

The following command lets you run an existing customized report:

```
ccautil -task report
-s <server:port_number>
-u <user_name>
-p <password> or -pwfile </path/filename>
-reportname <report_name>
[-dir <output_directory>]
```

-s <*server:port_number*>

Specifies the CA Configuration Automation Server name and port number from which you want to run a report.

-u <*user_name*>

Specifies a valid CA Configuration Automation Server user name.

-p <*password*>

Specifies the user's password. The password shows as clear text on the command line. Alternatively, you can create and use a scrambled password file for authentication. For more information, see Hash Password Option (see page 425).

-pwfile </*path/filename*> (on UNIX servers) or -pwfile <*pathfilename*> (on Windows servers)**

Specifies the user's password file. You can use this argument instead of -p <*password*>.

-reportname <*report_name*>

Specifies the name of the customized report that was created using the CA Configuration Automation Server GUI. For more information, see Report Management (see page 359).

[-dir <*output_directory*>]

(Optional) Specifies the directory path where the report is saved.

Default: Current directory

Usage Example

The following example generates an existing report called SF Datacenter Refresh on qaserver1 for user psmith with a I<3Artaud password:

```
ccautil -task report -reportname "SF Datacenter Refresh" -S qaserver1:8080
-u psmith -p I<3Artaud
```

Notice the report name is enclosed in quotes because "SF Datacenter Refresh" contains spaces.

Import File Option

The following command lets you import a server data file, instead of having to enter or maintain servers manually in the CA Configuration Automation Server user interface:

```
ccautil -task import
-s <server:port_number>
-u <login_name>
-p <password> or -pwfile </path/filename>
-file <import_filename>
```

-s <*server:port_number*>

Specifies the CA Configuration Automation Server name and port number that you want to import a server data file into.

-u <user_*name*>

Specifies a valid CA Configuration Automation Server user login name.

-p <*password*>

Specifies the user's password. The password shows as clear text on the command line. Alternatively, you can create and use a scrambled password file for authentication. For more information, see the Hash Password Option (see page 425).

-pwfile *</path/filename>* (on UNIX servers) or -pwfile *<\path\filename>* (on Windows servers)

Specifies the user's password file. You can use this argument instead of -p *<password>*.

-file *<import_filename>*

Specifies the name of the server data file in XML or CSV format. Sample server data files, called servers.xml and servers.csv, are included in the following location on the CA Configuration Automation Server host:

- (UNIX) */<install_dir>*/CA/CCAServer/doc/samples
- (Windows) *\<install_dir>*\CA\CA Configuration Automation Server\doc\samples

The sample files describe each field in detail and include examples.

Usage Example

- The following example imports a file called servers.csv to the CA Configuration Automation Server called cmdbtest:

```
ccautil -task import -file servers.csv -s cmdbtest:8080 -u user -p password
```

- The following example imports a file called servers.xml to the CA Configuration Automation Server called cmdbtest:

```
ccautil -task import -file servers.xml -s cmdbtest:8080 -u user -p password
```

Secure Agent Option

The following command creates a new SSL certificate for the agent, installs this new certificate in the agent installation directory, configures the agent to only accept secure connections, and restarts the agent with the new configuration:

```
ccautil -task secureagent
-s <server:port_number>
-u <login_name>
-p <password> or -pwfile </path/filename>
-server <agent_servername>
-acp <agent_certificate_password>
-cap <certificate_authority_password>
```

-s *<server:port_number>*

Specifies the HTTP or HTTPS CA Configuration Automation Server name and port number that contains the agent you want to secure.

-u *<user_name>*

Specifies a valid CA Configuration Automation Server user login name.

-p *<password>*

Specifies the user's password. The password shows as clear text on the command line. Alternatively, you can create and use a scrambled password file for authentication. For more information, see Hash Password Option (see page 425).

-pwfile*< /path/filename>* (on UNIX servers) or -pwfile <*pathfilename*> (on Windows servers)**

Specifies the user's password file. You can use this argument in place of -p *password*.

-server *<agent_name>*

Specifies the server for which you want to secure CA Configuration Automation Agent communications.

-acp *<agent_certificate_password>*

Provides a certificate password for the new agent certificate.

-cap *<certificate_authority_password>*

Provides the CA Configuration Automation Server certificate authority password to create the new agent certificate. The agent certificate you are creating must be issued for the proper server.

For more information, about SSL certification and the passwords you must enter to create a new SSL certificate, see Creating and Managing Security Certificates.

Usage Example

The following example creates and installs a new SSL certificate for the agent installed on a server named sqldbserver and enables agent security:

```
ccautil -task secureagent -S qaserver1:8080 -u bsmith -p user250
-server sqldbserver -acp fzr1000 -cap hdfxd1200
```

Assign Profile Option

The following command allows the CA Configuration Automation administrator to automate the administration process of updating the Server Management Profile, Server Access Profile and Service Management Profile:

```
ccautil -task assignprofile
-s <server:port of the CCA Server>
-u (user name)
-p <password> or -pwfile <file containing the password>
-service
-server
-id <service/server uuid>
-servicename <service name>
-servername <server name>
-apname <access profile name>
-mpname <management profile name>
```

-s *<server:port_number>*

Specifies the HTTP or HTTPS CA Configuration Automation Server name and port number that contains the agent you want to secure.

-u *<user_name>*

Specifies a valid CA Configuration Automation Server user login name.

-p *<password>*

Specifies the user's password. The password shows as clear text on the command line. Alternatively, you can create and use a scrambled password file for authentication. For more information, see Hash Password Option (see page 425).

-pwfile*< /path/filename>* (on UNIX servers) or -pwfile *<\path\filename> (*on Windows servers)

Specifies the user's password file. You can use this argument in place of -p *password*.

-service

Specifies the assign or update profile to Service.

-server

Specifies the assign or update profile to Server.

-id <service/server uuid>

Specifies the service or server user id.

-servicename *<service_name>*

Specifies the name of the service you want to refresh.

-servername <*server_name*>

Specifies the name of the server you want to refresh.

-apname <access profile name>

Specifies the name for the access profile.

-mpname <management profile name>

Specifies the name for the management profile.

Usage Example

The following example uses the -servername argument to uniquely identify the server, -apname argument to uniquely identify the access profile name.

```
ccautil -task assignprofile -server -servername "server1" -S ccaservername:8080 -u username –p password -apname "AProfile1"
```

The following example uses the -id argument to specify the server, -mpname argument to uniquely identify the mangement profile name.

```
ccautil -task refresh –server -id 5dfe7ce5-fbed-44b1-80ef-1a49f5780236 -S ccaservername:8080 -u username –p password -mpname "MProfile1"
```

The following example uses the -servicename argument to uniquely identify the service, -mpname argument to uniquely identify the mangement profile name.

```
ccautil -task assignprofile -service -servicename "service1" -S ccaservername:8080 -u username –p password -mpname "MProfile2"
```

The following example uses the -id argument to specify the service, -mpname argument to uniquely identify the mangement profile name.

```
ccautil -task assignprofile -service -id 5dfe7ce5-fbed-44b1-80ef-1a49f5780236  -S ccaservername:8080 -u username –p password -mpname "MProfile2"
```

How to Run ccautil on an HTTPS-enabled CA Configuration Automation Server

You must perform the following procedures to use the command-line interface on an HTTPS-enabled CA Configuration Automation Server:

- Configure the CA Configuration Automation Server in HTTPS mode as described in Create a Certificate Authority, Server Certificate, and HTTPS Certificate.
- Import the CCA Certificate as described in Import the Certificate Authority to the JRE Key Store (see page 435).

Import the Certificate Authority to the JRE Key Store

After configuring the CA Configuration Automation Server to use HTTPS, you must copy the ccaca.cert file from the CA Configuration Automation Server host to the computer from which you want to run ccautil, then import it into the JRE key store.

Follow these steps:

1. Navigate to one of the following directories on the client computer (the computer from which you want to run ccautil):
   - *<JAVA_HOME>*\jre\lib\security
   - *<JRE_HOME>*\jre\lib\security

   Ensure that the cacert certificate file is in this directory.
2. Copy the ccaca.cer file from the CA Configuration Automation Server's security directory to the ...\jre\lib\security directory on the client computer.

   If you installed the CA Configuration Automation Server in the default location, the ccaca.cer file is located in one of the following locations:
   - C:\Program Files\CA\CCA Server\security
   - opt/ca/CCAServer/jre/lib/security
3. Open a command prompt on the agent host computer, change (cd) to the ...jre\lib\security directory, and then run one of the following commands to import the Certificate Authority into the JRE key store:

   ```
   %JRE_HOME%\bin\keytool -import -trustcacerts -keystore cacerts -storepass changeit -noprompt -alias acmca –file <path>\ccaca.cer
   ```

   ```
   %JRE_HOME%\bin\keytool -importcert -trustcacerts -file <path>\ccaca.cer -keystore cacerts -storepass changeit -noprompt -alias ccaca
   ```

   The keytool import command copies the file, and ccautil can be accessed using HTTPS as follows:

   ```
   ccautil -task listservers -s https://ccaservername:8080 -u username –p password.
   ```

   Notes:
   - The CA Configuration Automation Server name must be specified using the HTTPS protocol on HTTPS-enabled CA Configuration Automation Servers for all command-line tasks.
   - Do not copy and paste the commands from this document—they may not work as expected. For example, the hyphen (-) cannot be copied properly, and will produce an error similar to the following:

     ```
     keytool error: java.lang.RuntimeException: Usage error, ûfile is not a legal command
     ```

Appendix F: Using the CA Configuration Automation SDK

CA Configuration Automation includes a set of integration capabilities called the CCA Software Development Kit (SDK). The CCA SDK consists of a Web Services front end to the CCA Server, and a set of Java classes wrapping the Web Services interface that allows access to the SDK classes.

This section contains the following topics:

SDK Web Service

The CA Configuration Automation Server is configured during installation to export a web service called SDKService. The service is defined in the services.xml file as follows:

```
<serviceGroup>
    <service name="SDKService">
        <messageReceivers>
            <messageReceiver mep="http://www.w3.org/ns/wsdl/in-out"class="com.ca.
cca.sdk.websvc.SDKServiceMessageReceiverInOut"/>
        </messageReceivers>
        <parameter
name="ServiceClass">com.ca.cca.sdk.websvc.SDKServiceImpl</parameter>
        --------------------------
        //more parameters and operations can be defined here.
        --------------------------
</serviceGroup>
```

The services.xml file is located on the CA Configuration Automation Server host in the following location:

<CA Configuration Automation Server_Install_Directory>\tomcat\webapps\cca\WEB-INF\services\SDKService\META-INF\

The SDKService can be enabled and disabled by editing the sdk.enabled property in the Properties table on the Configuration tab page in the CA Configuration Automation Server UI. The property is enabled by default.

The SDKService provides services for authentication and subsequent access to CA Configuration Automation data and management functions. The data and functions available include the following:

- Access to services, servers, their attributes, custom attributes, and parameters
- Refresh and Snapshot operations
- Running Discovery and Management Profiles
- Change Detection, Comparison, and Rule Compliance operations
- Execution of customized reports
- Assignment of custom attributes

All methods available through the SDKService are subject to access control policies as defined by the CCA Administrator.

The SDK distribution media includes a WSDL definition for the service called sdk.wsdl that can be viewed using the CA Configuration Automation Server user interface at the following URL: http://*<CA Configuration Automation Server>*:*port* /services/SDKService?wsdl

The WSDL interface implements an XML level exchange, where input parameters other than primitive types (for example, strings and integers) are formatted as XML, and values returned from the web service are similarly formatted as XML.

For example, the getServerByName method is defined by the following WSDL:

```
<wsdl: message name="getServerByNameRequest">
    <wsdl:part name="credential" type="xsd:string" />
    <wsdl:part name="name" type="xsd:string" />
    <wsdl:part name="getComponents" type="xsd:boolean" />
</wsdl:message>

<wsdl:message name="getServerByNameResponse">
    <wsdl:part name="getServerByNameReturn" type="xsd:string" />
</wsdl:message>
```

To request the data for a server by name, the web service method takes a string credential, the string name of the server, and a Boolean specifying whether software component details should be returned. The return value is an XML formatted string, similar to the following example:

```
<Server>
    <uuid>930ce95b-9c4a-4358-bb14-0777be7dc866</uuid>
    <type>cca_srvr</type>
    <name>mv0090.ca.com</name>
    <snapshot>false</snapshot>
    <attributes>
        <entry>
            <string>agent_protocol</string>
            <string></string>
        </entry>
        <entry>
            <string>os_ver</string>
            <string>5.2 (Build 3790)</string>
        </entry>
    ...
</Server>
```

The XML format of returned data is simple and self-documenting. The SDK, as described in the next section, takes these XML representations of objects and de-serializes them into actual Java objects, allowing programmatic access to data and functions. The web service is intentionally data- and XML-oriented, allowing language neutral utilization of the web service. XML is the standard of interchange, so that any program or scripting language that is capable of SOAP and XML processing (which should be most, including PERL and Python) can immediately gain access to the web service interface by compilation of the WSDL into client stubs, and processing of the resulting XML.

SDK Client API

The SDK is a Java wrapper that uses the SDKService web service to access CA Configuration Automation data and present it in a logical, object-oriented manner. The SDK is shipped in the cca-api.jar JAR file. The SDK uses an auxiliary file (cca-aux.jar) that contains open-source utilities. Applications external to CA Configuration Automation can use the classes that the SDK defines to access the SDKService and to program against objects that represent CA Configuration Automation-managed data.

You can access Javadoc pages for classes in the cca-api.jar file in the following locations:

- On the CA Configuration Automation Server host in the *<CA Configuration Automation Server_Install_Directory>*\sdk\doc\javadoc directory
- On the SDK distribution media

Note: If bulk SDK client API methods fail, they display the following message:

```
java.lang.StackOverflowError.
```

To avoid this issue, use the following JVM option to increase the native stack size to 2 MB:

```
>java -Xss2m
```

Gradually increase the native stack size until the error message does not appear if the issue persists.

com.ca.acm.sdk.net

The com.ca.acm.sdk.net package contains classes that establish the connectivity with, and provide session credentials to the CA Configuration Automation Server. As the following example shows, com.ca.acm.sdk.net.ACMSDKService is the primary class that CA Configuration Automation uses for this setup:

```
import com.ca.acm.sdk.net.ACMSDKService;

if (ACMSDKService.locateService(http://<yourserver>:port/services/SDKService) )
{
    if (ACMSDKService.beginSession(username, password) )
    {
        //  do some work
        Server[] server = Server.getAllServers();
        ...
     }
        ACMSDKService.endsession();
 }
```

The product uses ACMSDKService static methods to set up and tear down ThreadLocal connectivity and session credentials. After you use the locateServer() and beginSession() methods to set up a session, the product manages the connectivity and session parameters in the ACMSDKService class. The Subsequent calls to SDK methods use the ThreadLocal connectivity and session parameters without a need to view or handle session credentials directly.

Note: Connectivity and session setup are ThreadLocal. Therefore, in multithreaded applications, verify that SDK calls are made from the thread where the connectivity was originally established.

The ACMSDKService.endsession() method closes established sessions and clears the connectivity that the locateServer() call set up.

Establishing CA Configuration Automation Server Connectivity

The com.ca.acm.sdk.net package contains classes that establish the connectivity with, and provide session credentials to the CA Configuration Automation Server. As the following example shows, com.ca.acm.sdk.net.ACMSDKService is the primary class that CA Configuration Automation uses for this setup:

```
import com.ca.acm.sdk.net.ACMSDKService;

if (ACMSDKService.locateService(http://<yourserver>:port/services/SDKService) )
{
    if (ACMSDKService.beginSession(username, password) )
    {
        //  do some work
        Server[] server = Server.getAllServers();
        ...
     }
        ACMSDKService.endsession();
 }
```

The product uses ACMSDKService static methods to set up and tear down ThreadLocal connectivity and session credentials. After you use the locateServer() and beginSession() methods to set up a session, the product manages the connectivity and session parameters in the ACMSDKService class. The Subsequent calls to SDK methods use the ThreadLocal connectivity and session parameters without a need to view or handle session credentials directly.

Note: Connectivity and session setup are ThreadLocal. Therefore, in multithreaded applications, verify that SDK calls are made from the thread where the connectivity was originally established.

The ACMSDKService.endsession() method closes established sessions and clears the connectivity that the locateServer() call set up.

SDK Support for HTTPS-enabled CA Configuration Automation Server

This section describes how to configure SDK support for an HTTPS-enabled CA Configuration Automation Server.

In general, there are two steps:

- Configure the server in HTTPS mode
- Import CA Configuration Automation Certificate Authority to JRE key store

These steps are described in detail in the sections that follow.

To configure the CA Configuration Automation Server to use HTTPS

1. Log in to the CA Configuration Automation Server you want to use in HTTPS mode.
2. Click the Administration link, the Configuration tab, and then the Security Certificates link.

   The Security Certificates page appears.
3. Select Create Certificate Authority from the Table Actions drop-down list.

   The Create Certificate Authority dialog appears.
4. Enter and confirm the three required passwords, click the Set Up HTTPS check box, and then click OK.

   The ccaca.cer file is created in *<CA Configuration Automation Server_home>*\security directory, and the Security Summary area of the Security Certificates page shows that a Certificate Authority was created.
5. Restart the CA Configuration Automation Server.

   The CA Configuration Automation Server is configured to run in HTTPS mode.

To import the CA Configuration Automation Server Certificate Authority into the JRE key store

1. Navigate to the jre\lib\security directory of the JRE that is being used at the client, and ensure that a certificate file called cacerts exists.
2. Copy the ccaca.cer file from *<CA Configuration Automation Server_home>*\security to jre\lib\security.
3. Open a command prompt, then change directory (cd) to jre\lib\security.
4. Run either of the following commands to import the ccaca.cer into JRE key store:

   ```
   >%JRE_HOME%\bin\keytool -import -trustcacerts -keystore cacerts -storepass changeit -noprompt -alias ccaca –file ccaca.cer
   ```

   ```
   >%JRE_HOME%\bin\keytool -importcert -trustcacerts -file ccaca.cer -keystore cacerts -storepass changeit -noprompt -alias ccaca
   ```

After importing ccaca.cer into JRE using this command, if you still cannot access your CA Configuration Automation Server with the SDK client using HTTPS protocol, perform the following step:

5. (Optional) Do one of the following:

   - If the jssecacerts keystore file exists in the JRE Security folder, run the following command to import the ccaca.cer into jssecacerts keystore:

     ```
     >%JRE_HOME%\bin\keytool -importcert -trustcacerts -file ccaca.cer –keystore jssecacerts -storepass changeit -noprompt -alias ccaca
     ```

   - If the jssecacerts file does not exist in the JRE Security folder, copy the cacerts file, and rename it to jssecacerts, then run the import command.

   Note: Do not copy and paste commands from this document, the hyphen (-) symbol may not be copied properly and may cause the following error:

   ```
   keytool error: java.lang.RuntimeException: Usage error, ûfile is not a legal command
   ```

   The SDK can now be accessed using HTTPS.

SDK Support for Client Authentication using X.509 Certificates

This section describes how to use the SDK to configure CA Configuration Automation to support client authentication with X.509 certificates. Use the following SDK API method to use a client certificate to establish CA Configuration Automation Server connectivity.

```
import com.ca.acm.sdk.net.ACMSDKService;
if (ACMSDKService.locateService(http://<yourserver>:<port>/services/SDKService) )
{
     if (ACMSDKService.beginSessionWithCertificate(certificateFileName,
certificatePassphrase) )
     {
          // do some work
          ...
     }
     ACMSDKService.endsession();
}
```

Note: For more information about this method, see the SDK javadoc that is available in the CA Configuration Automation Server installation directory.

In the beginSessionWithCertificate() method, the code runs on the client to set the client certificate to SSL context before making a server call with an empty user name and password.

Use an SDK Web Service Call to Establish Connectivity

To use client certificate authentication to establish CA Configuration Automation Server connectivity, run the following code before making a call to the server:

```
//This is java code (write equivalent code in your language (C, C++, .NET etc...) before
making call to server.
String certificateFileName = “C:\\certs\\client.p12”
String certificatePassphrase = “password”
//create a trust manager that does not validate certificate chains
TrustManager[] trustAllCerts = new TrustManager[] { new X509TrustManager() {
     public java.security.cert.X509Certificate[] getAcceptedIssuers() {
          return null;
     }
     public void checkClientTrusted(
          java.security.cert.X509Certificate[] certs, String authType) {
     }
     public void checkServerTrusted(
          java.security.cert.X509Certificate[] certs, String authType) {
     }
} };
char[] passphraseChar = certificatePassphrase.toCharArray();
try {
     KeyManagerFactory kmf = KeyManagerFactory.getInstance(“SunX509”);
     KeyStore ks = KeyStore.getInstance(“PKCS12”);
     ks.load(new FileInputStream(certificateFileName), passphraseChar);
     kmf.init(ks, passphraseChar);
     SSLContext sc = SSLContext.getInstance(“SSL”);
     sc.init(kmf.getKeyManagers(), trustAllCerts, new
java.security.SecureRandom());
     SSLContext.setDefault(sc);
} catch (Exception e) {
   //handle exceptions here.
 }
```

Use the GetSessionCredential web service call with an empty user name and password as input to establish connectivity and to create a session with the CA Configuration Automation Server. The GetSessionCredential call returns the credential string to use in subsequent CA Configuration Automation web service calls.

Appendix G: Opening the CA Configuration Automation Server UI in Context

Context launch lets you launch the CCA Server UI directly using a URL. The context launch provides a mechanism to launch some CA Configuration Automation operations or dialogs using a URL. This URL can act as an integration tool and can be used in any report or application UI.

URL Parameters

The common parameters are as follows:

eem=eem-artifact

Specifies credentials to log in to CA EEM.

type

Specifies the type of the operation or dialog.

Parameters to run a Change Detection operation on a server are as follows:

type=cd

Specifies the run Change Detection operation on a server.

server=server-name OR ipaddress=ip-address

Specifies the server name or the IP address.

source=(defaults to 1)

Specifies the source snapshot to run Change Detection. The values are as follows:

- 1 — Current data
- 2 — Most recent snapshot
- 3 — Second most recent snapshot
- 4 — Baseline
- 5 — Gold Standard
- 6 — Silver Standard
- 7 — Bronze Standard
- Snapshot ID for Change Detection and compare operations

target=(defaults to 2)

Specifies the target snapshot. The values are as follows:

- 1 — Current data
- 2 — Most recent snapshot
- 3 — Second most recent snapshot
- 4 — Baseline
- 5 — Gold Standard
- 6 — Silver Standard
- 7 — Bronze Standard
- Snapshot ID for change detection and compare operations

Example:

http://servername:portnumber/CCAUI.html?type=cd&eem=eem-artifact&server=servername.ca.com&source=1&target=2

Parameters to view Change Detection dialog for a server are as follows:

type=cdd

Specifies the Change Detection operation for a server.

server=server-name OR ipaddress=ip-address

Specifies the server name or IP address.

Example:

http://
servername:portnumber/CCAUI.html?type=cdd&eem=token&server=servername.ca.com

Parameters to run Compare for two servers are as follows:

type=compare

Specifies to run a Compare operation for two servers.

srcserver=server-name or srcip=ip-address

Specifies the source server name or the IP address.

tgtserver=server-name or tgtip=ip-address

Specifies the target server or IP address.

source=(defaults to 1)

Specifies the source snapshot to run compare operations. The values are as follows:

- 1 — Current data
- 2 — Most recent snapshot
- 3 — Second most recent snapshot
- 4 — Baseline
- 5 — Gold Standard
- 6 — Silver Standard
- 7 — Bronze Standard
- Snapshot ID for change detection and compare operations

target=(defaults to 1)

Specifies the target snapshot. The values are as follows:

- 1 — Current data
- 2 — Most recent snapshot
- 3 — Second most recent snapshot
- 4 — Baseline
- 5 — Gold Standard
- 6 — Silver Standard
- 7 — Bronze Standard
- Snapshot ID for change detection and compare operations

For componentOnly parameter=(defaults to 0)

Specifies software components that are compared on the selected servers.

- 0 — All
- 1— Component only

Note: If you select component only (1), provide the following parameters:

- sourceComponentUUID = the source component uuid
- targetComponentUUIDs= the target component uuid.

Use comma to separate multiple target component UUIDS.

Examples:

- http://
  servername:portnumber/CCAUI.html?type=compare&srcserver=servername.ca
  .com&tgtserver=servername.ca.com&source=2&target=1
- http://servername:portnumber/CCAUI.html?type=compare&srcserver=servern
  ame.ca.com&tgtserver=servername.ca.com&source=2&target=1&componentO
  nly=1& sourceComponentUUID=xyzasdsdsadhkh1&
  targetComponentUUIDs=asdkhkhsd,sdshdkhak

Parameters to show Compare dialog for two servers:

type=compared

Displays the Compare Servers dialog.

srcserver=server-name or srcip=ip-address

Specifies the source server name and the IP address.

tgtserver=server-name or tgtip=ip-address

Specifies the target server name or the IP address.

Example:

http://servername:portnumber/CCAUI.html?type=compared&srcserver=servernam
e.ca.com&tgtserver=servername.ca.com

Parameters to view component for a selected component on server:

type=cv

Specifies view compare operations for a selected component on a server.

server=server-name or ipaddress=ip-address

Specifies the server name or IP address.

comp=component-name

Specifies the component name.

Example:

http://servername:portnumber/CCAUI.html?type=cv&server=servername.ca.com&
comp=CCA%20Server

Parameters to view Compliance Job results:

type=cjresult

Specifies the Compliance job operations.

resultuuid= result uuid of the history record

Specifies the UUID of the history entries.

Example:

http://servername:portnumber/CCAUI.html?type=cjresult&eem=eem-artifact&resultuuid="xyz"

Parameters for Access Profile details:

type=ap

Specifies the Access profile operations.

profile=profile-name

Specifies the profile name.

Example:

http://servername:portnumber/CCAUI.html?type=ap&profile=xyz

Parameters for Management Profile details:

type=mp

Specifies the Management Profile operations.

profile=profile-name

Specifies the profile name.

Example:

http://servername:portnumber/CCAUI.html?type=mp&profile=xyz

Parameters for Remediation Profile details:

type=rp

Specifies the Remediation Profile operations.

profile=profile-name

Specifies the profile name.

Example:

http://servername:portNumber/CCAUI.html?type=rp&profile=xyz

Parameters for Blueprint tab:

type=bp

Specifies the Blueprint.

Example

http://servername:portnumber/CCAUI.html?type=bp

Parameters for Server Log:

type=sl

Specifies the server log operations.

server=server-name OR ipaddress=ip-address

Specifies the server name or the IP address.

Example:

http:// servername:portnumber/CCAUI.html?type=sl&server=servername.ca.com

Parameters for Server Component List:

type=sc

Specifies the Server Component list operations.

server=server-name OR ipaddress=ip-address

Specifies the server name or the IP address.

Example

http:// servername:portnumber/CCAUI.html?type=sc&server=servername.ca.com

Parameters to run Rule Compliance for a server or service:

type=rc

Specifies the Rule Compliance operations.

server=server-name OR ipaddress=ip-address OR service=service_name

Specifies the server name, IP address, or service name.

source=(defaults to 1)

Specifies the source snapshot to run Rule compliance operations.

- 1 – Current data
- 2 – Most recent snapshot
- 3 – Second most recent snapshot
- 4 – Baseline

- 5 – Gold Standard
- 6 – Silver Standard
- 7 – Bronze Standard

Example

http://servername:portnumber/CCAUI.html?type=rc&eem=eem-artifact&server=servername.ca.com&source=1

Parameters to view Rule Compliance dialog for a server or service:

type=rcd

Specifies the Rule Compliance operations.

server=server-name or ipaddress=ip-address or service_name=service

Specifies the server name, IP address, or service name.

Example:

http://
servername:portnumber/CCAUI.html?type=rcd&eem=eem-artifact&server=servername.ca.com

Parameters to view Charts:

type=chart

Specifies the chart type.

name=chart-name

Specifies the chart name.

Example:

http://servername:portnumber/CCAUI.html?type=chart&name=Servers by OS Family

The chart names are as follows:

- Servers by OS Family and Operating System
- Servers by OS Family
- Servers by Operating System
- Servers by OS Version
- Servers by OS Detail
- Servers by Service
- Servers by Server Group
- Server Inventory over Time

- Servers by Manufacturer
- Servers by Memory Capacity
- Servers by Logical CPU Count
- Servers by Provider Count
- Servers by Consumer Count
- Servers by Packet Count
- Communication Types by Server Count
- Communication Types by Packet Count
- Servers by Packet Count Over Time
- Communication Types by Packet Count Over Time
- Software Components by Service
- Software Components
- Software Components by Category
- VM Management
- VM Hosting Servers by VM Guest Count and Virtual Environment
- VM Hosting Servers by VM Guest Count
- VM Hosting Servers by CPU Count
- VM Hosting Servers by Memory Capacity
- Installed Applications
- Servers with Rule Violations over Time
- Services with Rule Violations over Time
- Software Components with Rule Violations over Time
- Rule Violations over Time
- Total Rule Violations over Time
- Software Component Differences Over Time
- Software Component Differences by Server Over Time
- Software Component Differences by Service Over Time
- Management Profile Rule Violations Over Time
- Compliance Profile Rule Violations Over Time
- Average Grid Processing Time

- Average Grid Wait Time
- Grid Jobs Submitted

Parameters required for graphs:

type=graph

Specifies the graph type.

name=localized name

Specifies the name of the localized graph.

id=non-localized name

Specifies the name of the predefined graphs installed with CA Configuration Automation.

Example:

http://servername:portnumber/CCAUI.jsp?type=graph&name=All Server Relationships.

Note: Localized and non-localized graph names are same for English CCA Server. The ID parameter is applicable only for non-English CCA Server to access predefined graphs with English names.

Parameters required for CA CMDB:

type=ci&objtype=&name=&component=&eem=

type=ci&objtype=&name=&component=&eem=path= (Component only) path = installed directory or default directory

The following table shows the URL changes required for CA CMDB to support the context launch supported by CA Configuration Automation.

Object type	URL type argument (objtype=xxx)	URL name argument (name=xxx)	URL component argument (component=xxx)	EEM artifact	Example federated asset ID
Server (including virtual and lpar)	server	server name	n/a	optional	objtype=server&name=my_server_name&component=&eem=artifact

Object type	URL type argument (objtype=xxx)	URL name argument (name=xxx)	URL component argument (component=xxx)	EEM artifact	Example federated asset ID
Service	service	service name	n/a	optional	objtype=service&name=my service name&component=&eem=artifact
Component	component	server name	component name	optional	objtype=component&name= my_server_name &component=my component name&eem=artifact
File system (not supported)					
NIC	nic	server name	nic ip address	optional	objtype=nic&name=my_server_name &component=my nic ip address&eem=artifact
Hard drive (not supported)					
Cluster (not supported)					
Blueprint (not supported)					

Appendix H: Blueprint Wizard UI Reference

This section contains the following topics:

Blueprint Page: Component Blueprint Fields

The Component Blueprint pane on the Blueprint wizard Blueprint page contains the following fields:

Component Blueprint Name

Defines a unique name for the blueprint.

Limitations: The blueprint name cannot contain the following characters:
< > ; : " ' * + = \ / | ?

Component Version

Defines the software component version (release number).

- When a single blueprint supports all versions of a specific component, set this field to *.*
- Otherwise, define the specific supported version or version series (for example, 1.2 or 2.*)

Default: *.*

Blueprint Version

Defines the blueprint version, in your preferred format.

Default: 1.0.0

Discovery

Specifies whether discovery is enabled or disabled for the blueprint.

Default: Enabled

Description

Describes the blueprint item. The blueprint Details page displays the description.

Operating System

Specifies the operating system on which the software component runs. The operating system that you select determines whether the blueprint:

- Supports a specific operating system (for example, if the file system structure varies significantly across platforms or relies heavily on the Windows registry).
- Works across multiple operating systems (if the file system paths and configuration parameters can be normalized).
- Specifically supports an F5 BIG-IP load balancer.
- Specifically supports Cisco routers and network switches that use IOS.

  You can always add or remove operating system-specific files. However, choosing the target operating system lets you start with a more representative base blueprint structure. For example:

  - If you select Windows, the wizard automatically creates registry- and registry overlay-related components.
  - If you select Any UNIX or another UNIX-based operating system, the wizard does not create the registry- and registry overlay-related components.

Default: Any

Category

Specifies the component category under which the Component Blueprints page lists the new blueprint. **Note:** To separate your custom blueprints from predefined blueprints, use the Custom Components category. A separate category lets you find the blueprints easily and avoid accidentally editing the predefined blueprints.

Default: Custom Components

Discovery Methods Page: Search Options Fields

The Add New Search Options pane on the Blueprint wizard Discovery Methods page contains the following Search Options fields:

Search From

Defines where the search starts. If you do not set this value, the search begins at the top of the component file system.

For Windows, the top of the file system (/) includes all physical drives (C:, D:, and so on).

Default: / (root)

Search Depth

Defines the number of folders or directories the search includes.

Default: 10

Component Always Found

Specifies whether the process always considers the component as discovered.

Yes

The product does not search for the indicators. If you define a $(Root) parameter as a constant, the product manages files under the root. Select Yes to define and manage components in services that the product cannot or does not need to discover.

No

The product searches for the indicators.

Default: No

Discovery Methods Page: File Indicators Fields

The Add New File pane on the Blueprint wizard Discovery Methods page contains the following fields:

Name (posix)

Defines the directory or file to discover. You can use wildcard characters (*) to define the name with the POSIX pattern-matching syntax (for example, Agent_*).

Type

Specifies whether the product adds a File or a Directory to the blueprint.

Default: File

Note: Specify either the Path From Root field or the Depth From Root value.

Path From Root

Defines the partial path on which to locate the component root relative to the file indicator path. Define only the directories or folders between the file indicator and the component root (not the entire component path).

For example, the file indicator mx.ini locates the configuration file C:\Program Files\mx\setup\conf\mx.ini. To define the path C:\Program Files\mx as the component root, define the path from the root as setup\conf.

- The search does not support matching the full expression. Use wildcard characters in the partial path for intermediate directories with unknown variable names. For example, *\conf
- You can use the * or ? wildcard characters in the directory names. For example: *conf, \conf*, /*bin, /bin*, /b*n, and /bi?

Depth From Root

Defines the number of folders or directories the discovery includes, beginning from the search root.

Discovery Methods Page: Registry Indicators Search Options Fields

The \HKEY_LOCAL_MACHINE pane on the Blueprint wizard Discovery Methods page for Registry Indicators contains the following Search Options fields:

Search From

Defines where the search starts. If you do not set this value, the search begins at the top of the component file system.

For Windows, the top of the file system (/) includes all physical drives (C:, D:, and so on).

Default: \HKEY_LOCAL_MACHINE

Search Depth

Defines the number of folders or directories the search includes.

Default: 10

Component Always Found

Specifies whether the process always considers the component as discovered.

Yes

The product does not search for the indicators. If you define a $(Root) parameter as a constant, the product manages files under the root. Select Yes to define and manage components in services that the product cannot or does not need to discover.

No

The product searches for the indicators.

Default: No

Discovery Methods Page: Registry Indicators Fields

The Add New Registry Value/Key pane on the Blueprint wizard Discovery Methods page for Registry Indicators contains the following Registry Indicators fields:

Name (posix)

Defines the directory or file to discover. You can use wildcard characters (*) to define the name with the POSIX pattern-matching syntax (for example, Agent_*).

Type

Specifies whether the product adds a File or a Directory to the blueprint.

Default: File

Note: Specify either the Path From Root field or the Depth From Root value.

Path From Root

Defines the path from the search root.

Depth From Root

Defines the number of folders or directories the discovery includes, beginning from the search root.

Discovery Methods Page: Network Probe Fields

The Add New Network Probe pane on the Blueprint wizard Discovery Methods page for Network Probes contains the following Network Probe fields:

Name

Defines a name for the network probe.

Description

Describes the blueprint item.

Primary Ports

Defines a comma-delimited list of the first port numbers from which to collect responses.

Alternate Ports

Defines a comma-delimited list of port numbers from which to collect responses if the primary ports fail.

Probe

Describes an optional probe that is sent from CCA to the remote service after the port communications have been established. The probe is to verify that the service or component exists.

Regex Match

Defines the attribute that the product compares to the regular expression directive value. If the value does not match, the directive fails.

Version Regex

Defines the regular expression version.

Protocol

Specifies whether the probe uses Transmission Control Protocol (TCP) or User Datagram Protocol (UDP).

Default: TCP

Discovery Verification Rules Page: Discovery Verification Rule Fields

The Add New Discovery Verification Rule pane on the Blueprint wizard Discovery Verification Rules page contains the following Discovery Verification Rule fields:

Name

Defines a name for the discovery verification rule.

Description

Describes the blueprint item.

Directive Type

Defines which of the following directive types the product uses to:

- Extract the values from elements that are managed in a service
- Retrieve the values from managed servers

The Directive Type drop-down list contains the following options:

Constant

Defines a fixed value that the product uses to substitute variables or to construct complex strings. This directive type has the following common uses:

- Define the fixed Root and the RegistryRoot parameters when the component installation location is not ambiguous.
- Combine data from multiple sources by assembling variables and fixed text in a complex string.
- Expose component constants (for example, the vendor name) as parameters.

Database Query

Queries a database on a managed server, and (optionally) extracts a value from the data retrieved. This directive type has the following common uses:

- Run a query that extracts a column value from a database row.
- Run a stored procedure that changes the managed server or extracts a value.
- Run a query that extracts a result set and displays it in tabular form in a macro step directive.

Get File

Retrieves the contents of a specific file (when the file location is known). The product then filters the contents to define a directive value. This directive type has the following common uses:

- Extract the parameters from unstructured files by fetching the file contents and filtering them with a regular expression.
- Retrieve log files for display, storage, and import/export in macro step directives.

Get SNMP

Retrieves the value at a management information base (MIB) address from an SNMP agent, and (optionally) sets the directive value from the results.

Match File Name

Lists all files and directories at a designated location in a managed server file system, and (optionally) extracts the directive value from the list.

Match Registry Name

(Windows Only) Lists registry key and registry value names at a location in the registry tree, and (optionally) extracts the directive value from the list.

Match Registry Name Get Data

(Windows Only) Lists all registry key and registry value names at a location in the registry tree.

(Optional) Retrieves the data value of the selected names as the directive value.

Network Probe

Defines the parameters of a TCP socket opened to a remote service server and port. You can (optionally) send a probe to the port, and collect a response from the socket as the directive value.

Registry

(Windows only) Retrieves the data value that is associated with a registry key or value, and (optionally) defines the directive value from the filtered results.

Remote Execution

Runs a command or script on a managed server. The product captures the command output and returns it as the directive value. The regular expression typically filters the value to extract a concise value from verbose command output. This directive type has the following common uses:

- Access the configuration information that is only available from output (for example, operating system configurations).
- Access transient data such as memory, network, or CPU statistics.
- Access custom scripts and tools and import the output to the blueprint.
- Run utilities and scripts to update a managed server.

Web Service Call

Retrieves the application configuration data that is exposed as web services. The directive queries the web services and parses the returned data to configuration variables. The directive supports the configuration executables that can retrieve and store the configuration data.

Provide the WSDL URL where the web service is running. Web Services Description Language (WSDL) is an XML-based language that provides a model for describing web services.

Default: Constant

Stage

Specifies when verification directives run during discovery, relative to when the component parameter directives run:

Late

The verification directive requires a Component parameter value for the variable substitution. The Component parameter value is the parameter name that is defined in the Component Parameters and Variables tab. The Component parameter value ensures that the substituted variables have the values that are required for the verification.

Early

The verification directive does not depend on the value of any Component parameters.

Default: Late

Value

Defines a fixed string or a string that contains one or more variable substitutions.

- If the directive resolves to a string with a length greater than 0, the string becomes the directive value.
- If the string length is 0 or undefined, the product considers the directive to be without a value and uses any specified default value instead.

You can define a constant value as blank (one or more spaces). The blank value appears as no value on the user interface, but the product considers it a valid value.

Examples:

- Variable substitution with a fixed default:

  `$(VariableName)`

- Multiple substitutions in fixed text:

  `C:\$(V)\$(V2)\file.xml`

You can also use the following variable syntax:

$(*parameter_name*)

Defines normal component variable substitution.

$(#*parameter_name*)

References a parameter in the parent component.

$(@*global_variable_name*)

References a global variable.

Regex

Defines a regular expression that describes a set of strings. If you use the variable substitution to define a constant value, you can use a regular expression to filter the value. You can also use the regular expression like a conditional expression to test the value and return it in one of the following ways:

- Return paren(0) if the value matches the expression.
- Return no value if the value and the expression do not match.

Paren

Defines which subexpression the product returns on match if a regular expression includes a parenthesized subexpression.

Default: 0 (if you specify no Paren value).

Regex Match

Defines the attribute that the product compares to the regular expression directive value. If the value does not match, the directive fails.

Must Equal

Defines the attribute to which the product compares the directive value. If the values are not equal, the directive fails.

Value Case

Specifies whether the product considers text case for matches.

Default: Case Sensitive

Note: The default (Case Sensitive) option does *not* consider Agent_conf and agent_conf to be a match. The Case Insensitive option considers Agent_conf and agent_conf to be a match.

Translate

Defines a unique name that the translation applies to the directive output. The product precedes the specified Translate value with $CCTranslation$__ at run time so the value is identifiable in the database as a translation. If no translation matches, the product retains the original directive result.

The directive values that the product extracts from the configuration files and the registry is sometimes cryptic strings or integers that belong to an enumeration. Each value corresponds to a configuration state, but the value does not interpret the state. The product can translate these values to meaningful strings so they are clear when it displays them in the blueprint.

Map values in the product Database to refer to the translation name and trigger the translation of *from values* to *to values*. The following example maps the translation $CCTranslation$__IIS SERVER STATE:

```
<BlueprintTranslation name="IIS server state"
coh_name="IIS_SERVER_STATE" coh_id="23501"
created_by="system_user">
   <BlueprintTranslationEntry translate_from="2"
   translate_to="Running" />
   <BlueprintTranslationEntry translate_from="4"
   translate_to="Stopped" />
   <BlueprintTranslationEntry translate_from="6"
   translate_to="Paused" />
</BlueprintTranslation>
```

The product does not associate translations with a specific blueprint. To avoid conflict with other translations, ensure that each translation has a unique name.

Note: The product user interface does not currently support defining translation tables. Create the files in CA Configuration Automation data load format (as the example shows), and then load them to the database with the data loader utility.

Transform

Defines the XSL transform with which to filter a returned XML-formatted directive value. The transformed value replaces the original returned value and can be either a new XML value or any other text that the transformation generates.

To run an XSL transform, specify the location of an XSL file to apply to the directive value. The transform location can be a URL (file, http, or other) that is visible from the server, or a file name in the server CLASSPATH. The product retrieves the transform and applies it to the directive value to produce a new value for the directive.

Insert

Defines the value of a directive that the product inserts in another string to produce the appropriate result for display or for a subsequent operation.

The insertion location can be any string that contains $(VALUE) (all uppercase letters). The directive value replaces $(VALUE) to produce the filtered result. For example, if the initial directive value is 3 and the insertion string is 1.$(VALUE).export, the filtered directive value is 1.3.export.

Modifier

Modifies the discovery results. Select either the Does Host Use Agent or Is Alterpoint Device modifiers.

Modifier Parameters

Defines the parameters that the product applies to the specified modifier.

Management Page: File Management Options Fields

The $(Root) pane on the Blueprint wizard Management page contains the following File Management Options fields:

Managed Depth

Defines how many directory levels the discovery operation verifies below the file system root.

Retrieve configured maximum files (50,000)

Specifies how many files the discovery operation retrieves.

Selected: The discovery operation retrieves up to 50,000 files.

Cleared: The discovery operation retrieves up to the number of files that the Maximum Files field specifies.

Maximum Files

Defines the maximum number of files that the discovery operation retrieves.

Management Page: Directory Fields

The Add New Directory pane on the Blueprint wizard Management page contains the following Directory fields:

Name (posix)

Defines a name for the element to add. You can use wildcard characters (*) to define the name with the POSIX pattern-matching syntax (for example, Agent_*).

Path From Root

Defines the partial path on which to locate the element. Define only the directories or folders between the file indicator and the component root (not the entire component path).

Management Page: Directory Fields

The Add New File pane on the Blueprint wizard Management page contains the following File fields:

Name (posix)

Defines a name for the element to add. You can use wildcard characters (*) to define the name with the POSIX pattern-matching syntax (for example, Agent_*).

Path From Root

Defines the partial path on which to locate the element. Define only the directories or folders between the file indicator and the component root (not the entire component path).

Filters and Attributes Page Rules Tab Fields

The Rules tab on the Component Blueprint wizard File Filters and Attributes page contains the following fields:

Name

Defines a name for the rule.

Description

Describes the blueprint item.

Documentation URL

Specifies a URL where the rule documentation is located.

Constraint Type

Specifies either the File or Directory constraint type to which the rule applies.

File

- File Modification Date
- File Owner
- File Permissions
- File Size
- File Version
- Product Version

Directory

- Bytes
- Depth
- Directory Modification Date
- Directory Must Exist
- Directory Owner
- Directory Permissions
- File Must Exist
- Number of Directories
- Number of Files

Operation

Specifies the operation that the rule for the selected Constraint Type uses.

- = (equals)
- != (not equal)
- = (ignore case)
- != (do not ignore case)
- > (greater than)
- >= (greater than equal)
- < (less than)
- =< (less than equal)
- Integer Range (inclusive)
- Must Match (regex)
- Must Match (regex, ignore case)
- Must Not Match (regex).

Value

Specifies the value that is taken as the reference for the Operation. The constraint type, operation, and value define a match criteria for the rule.

Disable

Specifies whether the rule is enabled (selected) or disabled (cleared).

On Failure

Defines the string that the product writes to the Rule Compliance results when a compliance operation that uses this blueprint fails the rule.

Severity

Specifies which error level determines that the rule failed.

- Information
- Warning
- Error
- Critical

Registry Management Fields

The $(RegistryRoot) pane on the Blueprint wizard Management page contains the following Registry Management Options fields:

Managed Depth

Defines how many directory levels the discovery operation verifies below the file system root.

Retrieve configured maximum elements (50,000)

Specifies how many elements the discovery operation retrieves.

Selected: The discovery operation retrieves up to 50,000 elements.

Cleared: The discovery operation retrieves up to the number of elements that the Maximum Elements field specifies.

Maximum Elements

Defines the maximum number of elements that the discovery operation retrieves.

Registry Filters and Attributes Page Add Key Fields

The Component Blueprint wizard Registry Filters and Attributes page contains the following Add Key fields:

Default

Defines the default registry key value.

Interpret As

Specifies which of the following entities the product interprets the key as:

- Database Name
- Database Table
- Host Name and Port
- Host Name or IP Address
- JDBC URL
- TCP Port Number
- URL
- Web Service URL

Relationship Key

Specifies whether the key is a relationship key.

Relationship Type

Specifies one of the following relationship types:

Communicates With

Establishes the relationship between an application and a database server.

Manages

Establishes the relationship between a server that manages another server. For example, select Manages to view the relationship between a VMware vCenter Server that manages a VMware ESX host.

Hosts

Establishes the relationship between a server that hosts another server. For example, select Hosts to view the relationship between a VMware ESX host that hosts a VMware virtual machine.

Uses

Establishes the relationship between servers.

Visibility

Specifies whether the discovered registry key is shown or hidden. Also, specifies whether the default value defined in the Blueprint UI is shown or hidden. The value is masked and encrypted when you select the Visibility value as Hide Value.

Show Value

Displays the default value, and discovered registry key.

Hide Value

Hides the default value in the Blueprint UI, and discovered registry key. Displays ******* instead of the values. Hide the value if:

- The value is confidential (for example, a password).
- The value is binary.
- The value is too long to display.

The product encrypts the hidden registry key value in the database. They are not displayed or viewable in the UI.

Note: To ensure the security of this field, reenter the value if you later decide to show it.

Interpreted Server

Defines the target server information that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Target Instance

Defines the target database instance or application instance that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Application

Specifies the target application information available in any other parameters. Use variable substitution methods to define the parameter.

Registry Filters and Attributes Page Value Details Fields

The Component Blueprint wizard Registry Filters and Attributes page contains the following Value Details fields:

Default

Defines the default registry key value.

Interpret As

Specifies which of the following entities the product interprets the key as:

- Database Name
- Database Table
- Host Name and Port
- Host Name or IP Address
- JDBC URL
- TCP Port Number
- URL
- Web Service URL

Relationship Key

Specifies whether the key is a relationship key.

Relationship Type

Specifies one of the following relationship types:

Communicates With

Establishes the relationship between an application and a database server.

Manages

Establishes the relationship between a server that manages another server. For example, select Manages to view the relationship between a VMware vCenter Server that manages a VMware ESX host.

Hosts

Establishes the relationship between a server that hosts another server. For example, select Hosts to view the relationship between a VMware ESX host that hosts a VMware virtual machine.

Uses

Establishes the relationship between servers.

Visibility

Specifies whether the discovered registry value is shown or hidden. Also, specifies whether the default value defined in the Blueprint UI is shown or hidden. The value is masked and encrypted when you select the Visibility value as Hide Value.

Show Value

Displays the default value, and discovered registry value.

Hide Value

Hides the default value in the Blueprint UI, and discovered registry value. Displays ******* instead of the values. Hide the value if:

- The value is confidential (for example, a password).
- The value is binary.
- The value is too long to display.

The product encrypts the hidden registry key value in the database. They are not displayed or viewable in the UI.

Note: To ensure the security of this field, reenter the value if you later decide to show it.

Interpreted Server

Defines the target server information that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Target Instance

Defines the target database instance or application instance that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Application

Specifies the target application information available in any other parameters. Use variable substitution methods to define the parameter.

Interpreted Cluster

Defines the target Cluster name that is available in any other component parameter. To define the parameter, use variable substitution.

Database Page Fields

The Component Blueprint wizard Database page contains the following fields:

Name

Defines a name for the database.

Description

Describes the blueprint item.

Database

Defines the database access description.

User

Defines the account name of a user that can access the database.

Password

Defines the password that is associated with the specified user account.

Server

Defines the server where the database is installed.

Port

Defines the listening port of the database host server.

DB Type

Specifies the type of database in use, from the following options:

- DEFAULT_CONTEXT
- SQL_SERVER_CONTEXT
- ORACLE_CONTEXT
- INFORMIX_CONTEXT
- DB2_CONTEXT
- MYSQL_CONTEXT
- POSTGRES_CONTEXT
- HSQLDB_CONTEXT
- ORACLE9_CONTEXT
- ODBC
- CLOUDSCAPE
- ORACLE10_CONTEXT
- SYBASE11_CONTEXT
- INGRES_CONTEXT
- SQL_SERVER2K5_CONTEXT
- JAVA_DB_CONTEXT
- SYBASE15_CONTEXT
- ORACLE11_CONTEXT
- SQL_SERVER2K8_CONTEXT

Environment

Defines the environment variables that the agent sets before it runs the command.

The product runs the remote command in the CA Configuration Automation Agent environment. If the agent does not have the environment setup that is required to run a command, define the necessary environment variables as a comma-separated list of name/value pairs in the form name=value. For example:

<var_name>=<value>,<var_name2>=<value2>

The agent supports the variable substitution method.

Component Parameters and Variables Page Fields

The Component Blueprint wizard Component Parameters and Variables page contains the following fields:

Name

Defines a parameter substitution expression in the following form:

`$(`*`VariableName`*`)`

VariableName

Defines the name of a discovery parameter that is defined in the component. By default, the Name value is case-sensitive, so it must match the parameter name exactly.

- You can embed the parameter expressions in string literals or you can use them on their own.
- You can define recursive multiple substitutions in the same expression. For example, a substitution value can be a string in the parameter expression. If it is, the blueprint evaluates it recursively.

Example:

If the Discovery parameters have the following values:

```
User=info
Domain=ca.com
v1=$(User)
v2=$(Domain)
```

the following parameter substitution expression:

```
$(User)@$(Domain) [$(v1) at $(v2)]
```

returns the following results after evaluation:

```
info@ca.com [info at ca.com]
```

Description

Describes the blueprint item.

Folder

Defines the item location.

Selected Categories

Defines which of the following categories the product uses to organize the elements. When you double-click an item in the Available Categories column, it moves to the Selected Categories column.

Administration

Defines administrative settings relative to general component availability and management. Examples of administrative settings include:

- How to back up elements.
- When to delete a cache.
- How many times to retry an action.

Configuration

Defines component configuration settings that are not related to the Administration, Log and Debug, Network, or Performance settings. Examples of configuration settings include a component alias or default web page.

Documentation

Defines the elements that document the component behavior or provide information to users. For example, guides, readme files, FAQs, or online help pages.

Log And Debug

Defines the elements that let the user set log locations, log levels, debug output, or diagnostic variable types.

Network

Defines the elements that represent a network-related component setting (for example, the SNMP port settings). If an element can be categorized as both Network and Security (for example, enabling LDAP Authentication), set the category to Security.

Other

CA internal use only.

Performance

Defines the performance settings, which are generally a specific subset of configuration parameters. For example, number of threads or number of concurrent users.

Product Info

Defines general (static) product information (for example, licensing, installation location, vendor, or module name).

Resources

Defines component resources (for example, storage, memory and cache allocation or size, or CPU).

Note: The static Resource category is different from the Transient category, which defines real-time information.

Security

Defines the elements that represent security-related settings, which are generally a specific subset of configuration parameters. For example, authentication types, enabling authentication, encryption settings, directory browsing, SSL, or HTTPS.

Transient

Defines the elements that regularly change. For example, server states (up, down, running, or stopped), current number of connected clients, current number of threads, or current disk utilization.

Versioning and Patches

Defines the elements that indicate the product version or patch levels.

Selected Filters

Defines which of the following elements the product filters from the Compare operations. When you double-click an item in the Available Filters column, it moves to the Selected Filters column.

Component Specific

Filters the elements that are specific to a single component instance (for example, installation root and Service Server Name). Excludes component-specific elements that are already known to be different from Change Detection operations and results.

Service Specific

Filters the elements that are specific to a service (for example, server names and installation roots). Excludes service-specific elements that are already known to be different from Change Detection operations and results.

Server Specific

Filters the elements that are specific to a single server (for example, the Server name and IP address). Excludes server-specific elements that are already known to be different from Change Detection operations and results.

Never Run Change Detection

Identifies the elements to exclude permanently from Change Detection operations and results. Examples include temporary directories, log files in the managed folder, or elements that are known to be transient.

Never Run Rule Compliance

Identifies the elements to exclude permanently from Rule Compliance operations and results because they are inconsequential or variable. Examples include temporary directories, known old configuration files, templates, or example files.

Time Variant

Filters the elements that are known to change over time, but not necessarily across servers or across services. Examples include log files, process start times, and registry event counters. Excludes time-variant parameters that are already known to be different from Change Detection operations and results.

File Size Variant

Filters the files based on the file size in the Change Detection Operations only when the Time Variant filter is not selected.

Weight

Specifies the relative importance of an element:

- Low
- Medium
- High

Default: Medium (elements for which you assign no weight)

Directive Type

Defines which of the following directive types the product uses to:

- Extract the values from elements that are managed in a service
- Retrieve the values from managed servers

The Directive Type drop-down list contains the following options:

Constant

Defines a fixed value that the product uses to substitute variables or to construct complex strings. This directive type has the following common uses:

- Define the fixed Root and the RegistryRoot parameters when the component installation location is not ambiguous.
- Combine data from multiple sources by assembling variables and fixed text in a complex string.
- Expose component constants (for example, the vendor name) as parameters.

Configuration

Retrieves the value of a configuration parameter from a parsed configuration file or from a configuration program file.

Database Query

Queries a database on a managed server, and (optionally) extracts a value from the data retrieved. This directive type has the following common uses:

- Run a query that extracts a column value from a database row.
- Run a stored procedure that changes the managed server or extracts a value.
- Run a query that extracts a result set and displays it in tabular form in a macro step directive.

Get File

Retrieves the contents of a specific file (when the file location is known). The product then filters the contents to define a directive value. This directive type has the following common uses:

- Extract the parameters from unstructured files by fetching the file contents and filtering them with a regular expression.
- Retrieve log files for display, storage, and import/export in macro step directives.

Get LDAP

Retrieves named data sets from a Directory Server, and (optionally) extracts the directive value from the output.

Get SNMP

Retrieves the value at a management information base (MIB) address from an SNMP agent, and (optionally) sets the directive value from the results.

Match File Name

Lists all files and directories at a designated location in a managed server file system, and (optionally) extracts the directive value from the list.

Match Registry Name

(Windows Only) Lists registry key and registry value names at a location in the registry tree, and (optionally) extracts the directive value from the list.

Match Registry Name Get Data

(Windows Only) Lists all registry key and registry value names at a location in the registry tree. You can (optionally) set the directive value from the data value of the selected names.

Network Probe

Defines the parameters of a TCP socket opened to a remote service server and port. You can (optionally) send a probe to the port, and collect a response from the socket as the directive value.

Registry

(Windows only) Retrieves the data value that is associated with a registry key or value, and (optionally) defines the directive value from the filtered results.

Remote Execution

Runs a command or script on a managed server. The product captures the command output and returns it as the directive value. The regular expression typically filters the value to extract a concise value from verbose command output. This directive type has the following common uses:

- Access the configuration information that is only available from output (for example, operating system configurations).
- Access transient data such as memory, network, or CPU statistics.
- Access custom scripts and tools and import the output to the blueprint.

Run utilities and scripts to update a managed server.

Default

Defines a fixed string or a string that contains one or more variable substitutions. If the directive runs and the product cannot determine a value, the product uses the default value instead of an undefined or zero-length value.

Persistence

Specifies whether a parameter persists in a service.

Always Persist

The product saves the parameter and displays it in the service tree view.

Never Persist

The product does not save the parameter only uses it for discovery.

Persist If Not Null

The product saves the parameter and (if the parameter has a value) displays it in the service tree view.

Default: Always Persist

Visibility

Specifies whether the component shows or hides the component parameter and its value. Also, specifies whether the default value defined in the Blueprint UI is shown or hidden. If you hide the value, the value is encrypted.

Show Value

Displays the default value, and component parameter and its value.

Hide Value

Hides the default Value in Blueprint UI and the component parameters value in the component viewer UI and displays ******** instead of the respective values. Hide the component parameter value and default value if:

- The value is confidential (for example, a password).
- The value is binary.
- The value is too long to display.

The product encrypts the hidden component parameters value in the database. They are not displayed or viewable in the UI.

Note: To ensure the security of this field, reenter the value if you later decide to show it.

Hide Element

Hides the parameter. Select Hide Element so the value is available for variable substitution, but is not displayed in the service tree view.

Default: Show Value

Value Case

Specifies whether the product considers text case for matches.

Default: Case Sensitive

Note: The default (Case Sensitive) option does *not* consider Agent_conf and agent_conf to be a match. The Case Insensitive option considers Agent_conf and agent_conf to be a match.

Interpret As

Defines a hint about the string format of a configuration parameter and how the associated component uses it. The product uses context-sensitive parsers to inspect interpreted parameter values. The parsers let the product extract multiple subvalues from complex parameter strings.

Default: blank (no interpretation)

Relationship Key

Specifies whether the product determines and assigns relationships according to the Interpret As value.

To establish relationships, set the Interpret As value and then set the Relationship Key field to **Yes**. Not all interpretations can define relationships.

Default: No

Relationship Type

Specifies one of the following relationship types:

Communicates With

Establishes the relationship between an application and a database server.

Manages

Establishes the relationship between a server that manages another server. For example, select Manages to view the relationship between a VMware vCenter Server that manages a VMware ESX host.

Hosts

Establishes the relationship between a server that hosts another server. For example, select Hosts to view the relationship between a VMware ESX host that hosts a VMware virtual machine.

Uses

Establishes the relationship between servers.

Interpreted Server

Defines the target server information that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Target Instance

Defines the target database instance or application instance that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Application

Specifies the target application information available in any other parameters. Use variable substitution methods to define the parameter.

Translate

Defines a unique name that the translation applies to the directive output. The product precedes the specified Translate value with $CCTranslation$__ at run time so the value is identifiable in the database as a translation. If no translation matches, the product retains the original directive result.

The directive values that the product extracts from the configuration files and the registry is sometimes cryptic strings or integers that belong to an enumeration. Each value corresponds to a configuration state, but the value does not interpret the state. The product can translate these values to meaningful strings so they are clear when it displays them in the blueprint.

Map values in the product Database to refer to the translation name and trigger the translation of *from values* to *to values*. The following example maps the translation $CCTranslation$__IIS SERVER STATE:

```
<BlueprintTranslation name="IIS server state"
coh_name="IIS_SERVER_STATE" coh_id="23501"
created_by="system_user">
   <BlueprintTranslationEntry translate_from="2"
   translate_to="Running" />
   <BlueprintTranslationEntry translate_from="4"
   translate_to="Stopped" />
   <BlueprintTranslationEntry translate_from="6"
   translate_to="Paused" />
</BlueprintTranslation>
```

The product does not associate translations with a specific blueprint. To avoid conflict with other translations, ensure that each translation has a unique name.

Note: The product user interface does not currently support defining translation tables. Create the files in CA Configuration Automation data load format (as the example shows), and then load them to the database with the data loader utility.

Transform

Defines the XSL transform with which to filter a returned XML-formatted directive value. The transformed value replaces the original returned value and can be either a new XML value or any other text that the transformation generates.

To run an XSL transform, specify the location of an XSL file to apply to the directive value. The transform location can be a URL (file, http, or other) that is visible from the server, or a file name in the server CLASSPATH. The product retrieves the transform and applies it to the directive value to produce a new value for the directive.

Insert

Defines the value of a directive that the product inserts in another string to produce the appropriate result for display or for a subsequent operation.

The insertion location can be any string that contains $(VALUE) (all uppercase letters). The directive value replaces $(VALUE) to produce the filtered result. For example, if the initial directive value is 3 and the insertion string is 1.$(VALUE).export, the filtered directive value is 1.3.export.

Modifier

Modifies the discovery results. Select either the Does Host Use Agent or Is Alterpoint Device modifiers.

Modifier Parameters

Defines the parameters that the product applies to the specified modifier.

Interpreted Cluster

Defines the target Cluster name that is available in any other component parameter. To define the parameter, use variable substitution.

Configuration - File Parsing Page

The Configuration - File Parsing page contains the following fields:

Name (regex)

Defines the name of the configuration file that the product locates and parses during discovery. You can use wildcard characters (for example, ^ or $) to define the name as per regex naming conventions.

Display Name

Defines the name of the configuration file that the product displays after discovery.

Description

Describes the blueprint item.

File Type

Specifies the type of configuration file (for example, text, binary, or log)

Configuration Executables Page

The Configuration Executables page contains the following fields:

Name

Defines a name for the configuration executable. This name appears under Executables in the Configuration folder. The name must be unique in the Executables folder.

Description

(Optional) Describes the blueprint item. The product displays the description as a tool-tip over the directive name in the blueprint and over the parameter name in the tree views.

Directive Type

Specifies the type of directive the product uses to obtain the parameter value. All directives return a value and a Boolean result. The value appears next to the parameter name when the discovered parameter appears in the service tree view.

Note: The Configuration directive type extracts a value from a managed configuration file.

Default: Constant

File

Defines the name of a configuration file that is in either the Configuration, Files folder or the Configuration, Executables folder. You can specify a specific file name or a regular expression. If you specify a regular expression, the product selects all matching files in the Configuration folder. If multiple files match, the product selects the file closest to the root of the managed file system. The product allows variable substitution.

Path

Defines the path to the configuration file. The product does not support wildcard characters or pattern matching, but it allows variable substitution. The path can either be absolute (for example, $(Root)/conf) or relative (for example, /conf) to the root of the managed file system. If the configuration file is at the configuration file system root, you can define Path as $(Root) or you can leave it undefined.

Parameter

Defines the name of the configuration parameter in the file. The product does not support wildcard characters or pattern matching, but it allows variable substitution. You can leave the Parameter attribute undefined if the value belongs to a Group file block. In that case, specify only the Group Path attribute.

Regex

Specifies whether to filter the named value with a regular expression if the product finds it.

Paren

Defines which subexpression the product returns on match if a regular expression includes a parenthesized subexpression.

Default: 0 (if you specify no Paren value).

Group Path

Specifies hierarchically organized configuration files. Use the Group Path value within the file to the name Parameter and that is specified like a partial file system path (for example, a/b/c).

Because a single file often contains the same name many times, use a unique Group Path to differentiate a specific value from the matching names. For example, an XML configuration file can have the following format, where the tag server and the buildDate and type attributes occur many times:

```
<configuration>
   <server name= “wxp123”>
      <setup buildDate=”10/30/2003” type=”webserver”/>
   </server>
   <server name=”wxp123” auxiliary=”true”>
      <setup buildDate=”09/02/2002” type=”webserver”/>
   </server>
      <server name=”wxp124”>
      <setup buildDate=”10/29/2003” type=”dataserver”/>
   </server>
</configuration>
```

To identify children that must match for the product to follow the path, use a comma-separated list of Group Path names (or names and values). For example, to extract the buildDate value for server wxp124, define the following values:

Parameter

buildDate

Group Path

configuration/server,name=wxp124/setup

Extracting the buildDate value from the server wxp123 auxiliary data only requires the name because it is the only server tag with an auxiliary attribute:

Parameter

buildDate

Group Path

configuration/server,auxiliary/setup

You can qualify any or all of the elements in a Group Path and you can specify multiple qualifiers for each element. If you specify multiple qualifiers, each qualifier must match to select a path. The product allows variable substitution.

If a structure class for the configuration file defines group blocks qualifiers, the product allows an alternate syntax. For group blocks, you can surround the qualifier value with parentheses and append it to the block name. For example, to extract the value of buildDate from server wxp124 where the attribute name is a Qualifier Child, specify the following values:

Parameter

buildDate

Group Path

configuration/server(wxp124)/setup

The value in parentheses matches the qualifier of the named group block. Either a Qualifier Child value or a constant Qualifier in the structure class defines the qualifier, or the qualifier is the group value.

If you are not sure about the parameter location in the file, you can specify multiple paths. Separate two or more specified paths with a pipe (|). The product tries the paths one at a time, from left to right, until a match it finds a match. For example:

Parameter

buildDate

GroupPath

configuration/server(abc)/setup|configuration/server(xyz)/setup

Empty Is Null

Specifies whether to override the default product behavior when the named configuration parameter in the file has no value. The default behavior is to set the value to an empty, zero-length value. The Empty Is Null field lets you distinguish between a directive that returns an empty value and a directive that returns no value.

Yes

Uses the value in the Default field.

No

Does not use the value in the Default field.

Default: No

Default

Defines a fixed string or a string that contains one or more variable substitutions. If the product cannot determine a value when it runs the directive, the product uses this value instead of an undefined or zero-length value.

In some cases (for example, if a Component Blueprint plug-in can set a directive), you do not want the Default value to replace the directive. To keep the Default value from replacing the plug-in value, set it to **Cohesion.PLACEHOLDER**.

Value Case

Specifies whether the product considers text case for matches.

Default: Case Sensitive

Note: The default (Case Sensitive) option does *not* consider Agent_conf and agent_conf to be a match. The Case Insensitive option considers Agent_conf and agent_conf to be a match.

Translate

Defines a unique name that the translation applies to the directive output. The product precedes the specified Translate value with $CCTranslation$__ at run time so the value is identifiable in the database as a translation. If no translation matches, the product retains the original directive result.

The directive values that the product extracts from the configuration files and the registry is sometimes cryptic strings or integers that belong to an enumeration. Each value corresponds to a configuration state, but the value does not interpret the state. The product can translate these values to meaningful strings so they are clear when it displays them in the blueprint.

Map values in the product Database to refer to the translation name and trigger the translation of *from values* to *to values*. The following example maps the translation $CCTranslation$__IIS SERVER STATE:

```
<BlueprintTranslation name="IIS server state"
coh_name="IIS_SERVER_STATE" coh_id="23501"
created_by="system_user">
   <BlueprintTranslationEntry translate_from="2"
   translate_to="Running" />
   <BlueprintTranslationEntry translate_from="4"
   translate_to="Stopped" />
   <BlueprintTranslationEntry translate_from="6"
   translate_to="Paused" />
</BlueprintTranslation>
```

The product does not associate translations with a specific blueprint. To avoid conflict with other translations, ensure that each translation has a unique name.

Note: The product user interface does not currently support defining translation tables. Create the files in CA Configuration Automation data load format (as the example shows), and then load them to the database with the data loader utility.

Transform

Defines the XSL transform with which to filter a returned XML-formatted directive value. The transformed value replaces the original returned value and can be either a new XML value or any other text that the transformation generates.

To run an XSL transform, specify the location of an XSL file to apply to the directive value. The transform location can be a URL (file, http, or other) that is visible from the server, or a file name in the server CLASSPATH. The product retrieves the transform and applies it to the directive value to produce a new value for the directive.

Insert

Defines the value of a directive that the product inserts in another string to produce the appropriate result for display or for a subsequent operation.

The insertion location can be any string that contains $(VALUE) (all uppercase letters). The directive value replaces $(VALUE) to produce the filtered result. For example, if the initial directive value is 3 and the insertion string is 1.$(VALUE).export, the filtered directive value is 1.3.export.

Modifier

Modifies the discovery results. Select either the Does Host Use Agent or Is Alterpoint Device modifiers.

Modifier Parameters

Defines the parameters that the product applies to the specified modifier.

Parser Details

Specifies either Structure Class or Parser. Select an option from the corresponding drop-down list.

Add Query Pane

The Component Blueprint wizard Add Query pane contains the following fields:

Name

Defines which directive queries a database on a managed server, and (optionally) extracts a value from the output data. This directive is commonly used to:

- Run a query that extracts a column value from a database row.
- Run a stored procedure that modifies the managed server or to extract a value.
- Run a query that extracts a result set and displays it in tabular form in a directive macro step.

The specified name must be unique in the executables folder of a blueprint.

Description

(Optional) Describes the blueprint item. The product displays the description as a tool-tip over the directive name in the blueprint and over the parameter name in the tree views.

Query Type

Specifies one of the following query types that direct the product how to run the query and whether the query returns data:

Select

The query is an SQL select statement that returns data.

Insert

The query is an SQL insert statement that returns no data.

Delete

The query is an SQL delete statement that returns no data.

Update

The query is an SQL update statement that returns no data.

Others

The query is an SQL alter table or other DDL statement that returns no data.

Stored Procedure

The query string invokes a stored procedure that can return data.

Query

Defines an SQL query or stored procedure name. Query syntax can vary depending on the database type, and variable substitution is allowed. The product returns all rows that the query or stored procedure finds. To filter the result set, use the column and regular expression attributes.

Max Rows

Defines the maximum number of rows that a query returns.

Default: 5000

Primary Columns

Defines the name of the column or an aliased column in the returned result set. The product takes the value in the named column as the directive value. If the query returns more than one row, the product uses the named column in the first row.

Structure Class

Defines the hierarchical collection of groups and parameters that the product uses to map metadata to the name/value pairs that the parser retrieves.

Note (Edit Mode only)

Defines a note about the selected element. The product shows the number of notes on the element or None in brackets. From the drop-down list, you can view all notes, add a note, or view, update, or delete a note.

- If you select Show All Notes, a new page lists all notes. To change the table sort order, click any table heading.
- If you select New Note, the following fields appear:

  Name

  Defines a name for the note.

  Text

  Defines the note text.

If you select an existing note, the Notes field shows the note name with the note text below the name. You can update, delete, or cancel the selected note.

File Structure Class Tab

The File Structure Class tab contains the following fields:

Name

Defines a name for the structure class. The Blueprint list and (unless you specify a Display Name) the tree view show the specified structure class name.

Version

Defines the structure class version of the structure class. The Blueprint list and (unless you specify a Display Name) the tree view show the specified version.

Display Name

Defines the class name that the Component Tree view shows if it differs from the Name and Version fields. For example, you can specify a Display Name value that is more descriptive of the class function.

Description

(Optional) Describes the blueprint item. The product displays this description as a tool-tip over the class name in the Blueprint tree view.

Allow Remediation Jobs

Specifies whether the product can modify the changeable blueprint elements as specified in a remediation job.

Default: Yes

Parser

Defines the parser to use for files of the defined type. The product includes a library of predefined lexical analyzers (lexers) and parsers that let the product analyze most common configuration file formats.

File Structure Class Group and Parameter Fields

The Add Group and Add Parameter panes on the Precedence tab contain the following fields:

Name

Defines name for the precedence group.

Description

Describes the blueprint item.

Available Categories

Specifies a category for the group. Double-click one or more categories to move them to the Selected Categories column.

Available Filters

Specifies a filter for the group. Double-click one or more filters to move them to the Selected Filters column.

Weight

Specifies the relative importance of an element:

- Low
- Medium
- High

Default: Medium (elements for which you assign no weight)

Data Type

Specifies the elements data type (for example, string, Boolean, or integer).

Valid Values

Defines the valid parameter or group values for a data type range (for example, integer enumeration, integer range, or string enumeration).

Default Value

Defines a fixed string or a string that contains one or more variable substitutions. If it cannot determine a value at run time, the product uses the Default Value instead of an undefined or zero-length value.

If a Component Blueprint plug-in can set the directive value, you can set Default Value to Cohesion.PLACEHOLDER to keep the specified Default Value from replacing the plug-in value.

Visibility

Specifies whether the component shows or hides the configuration parameter or group value. Also, specifies whether the default value defined in the Blueprint UI is shown or hidden. If you hide the value, the value is encrypted.

Show Value

Displays the default value, and configuration parameter or group value.

Hide Value

Hides the default value in Blueprint UI and configuration parameter or group value in the component viewer UI.

- The value is confidential (for example, a password).
- The value is binary.
- The value is too long to display.

The product encrypts the hidden configuration parameters or group value in the database. They are not displayed or viewable in the UI.

Note: To ensure the security of this field, reenter the value if you later decide to show it.

Default: Show Value

Value Case

Specifies whether the product considers text case for matches.

Default: Case Sensitive

Note: The default (Case Sensitive) option does *not* consider Agent_conf and agent_conf to be a match. The Case Insensitive option considers Agent_conf and agent_conf to be a match.

Interpret As

Defines a hint about the string format of a configuration parameter and how the associated component uses it. The product uses context-sensitive parsers to inspect interpreted parameter values. The parsers let the product extract multiple subvalues from complex parameter strings.

Default: blank (no interpretation)

Relationship Key

Specifies whether the product determines and assigns relationships according to the Interpret As value.

To establish relationships, set the Interpret As value and then set the Relationship Key field to **Yes**. Not all interpretations can define relationships.

Default: No

Relationship Type

Specifies one of the following relationship types:

Communicates With

Establishes the relationship between an application and a database server.

Manages

Establishes the relationship between a server that manages another server. For example, select Manages to view the relationship between a VMware vCenter Server that manages a VMware ESX host.

Hosts

Establishes the relationship between a server that hosts another server. For example, select Hosts to view the relationship between a VMware ESX host that hosts a VMware virtual machine.

Uses

Establishes the relationship between servers.

Interpreted Server

Defines the target server information that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Source Instance

Defines the source database instance or application instance.

Interpreted Target Instance

Defines the target database instance or application instance that is available in any other component parameter. To define the parameter, use variable substitution.

Interpreted Application

Specifies the target application information available in any other parameters. Use variable substitution methods to define the parameter.

Interpreted Cluster

Defines the target Cluster name that is available in any other component parameter. To define the parameter, use variable substitution.

Expected Value

Defines the expected value of the parameter or group.

Value

Defines a unique value for the group in the file.

Qualifier Child (Add Group only)

Defines the parameter that the product uses as a qualifier under the group.

Macros Page

The Macros page contains in the following fields:

Name

Defines a name for the macro.

Description

Describes the blueprint item.

Folder

Defines the item location.

Diagnostic

Specifies whether the product uses the macro to diagnose issues.

Default: No

Read Only

Specifies whether the macro can modify the target system.

Note: To run a read-only macro, the user must have Server View or Service View permissions.

Default: No (the macro cannot modify the target system).

Finish Page

The wizard page contains the following fields:

Allow this blueprint to contain other components

Specifies whether the blueprint can contain other components.

Allow this blueprint to nest in other components

Specifies whether other components can contain the blueprint.

Note: By default, the Allow this blueprint to contain other components, and Allow this blueprint to nest in other components checkboxes are enabled.

Allow this blueprint to contain only named components

Specifies whether the blueprint can contain only named components.

This blueprint will not use nesting

Specifies whether the blueprint can use nesting.

Refresh Order

Specifies the position that this component refreshes relative to all others on a specific server. The product refreshes components on each server sequentially and refreshes components on other servers in parallel.

Refresh order is important when components depend on each other for variable substitution. For example, if one component uses a parameter value from another component for the variable substitution, the product must refresh the other component first. If you do not specify Refresh Order in this case, the dependent component can pick up a value that the current refresh has not updated.

Refresh Order has the following values. No specific order is guaranteed within a level. :

- **Don't Care:** The product refreshes the component as if it has the Middle ordinal.
- **Initialization:** The product refreshes components that are set to Initialization first.
- **First**
- **Early**
- **Middle**
- **Late**
- **Last**
- **Cleanup:** The product refreshes components that are set to Cleanup last.

Modifier

Specifies one of the following options: to modify the discovered results:

- Device Authority
- IIS Nesting
- Replace All
- Sun Application Server Nesting
- TIBCO
- TIBCO Nesting

Modifier Parameters

Defines the parameters for the selected Modifier.

Nearest Neighbor

Specifies whether the product selects the component with the closest file system root (in terms of depth) if a component can nest in multiple components.

Default: Yes

Nest By

Specifies the file system relationship that determines whether the component nests in other components.

Note: If you select This Blueprint Will Not Use Nesting, Nest By is not required:

Child

Defines that the component should nest in any component whose file system root is higher in the file system directory tree.

Child or Sibling

Defines that the component should nest in any component whose file system root is equal to or higher in the file system directory tree.

Direct Child

Defines that the component can nest only in components whose file system root is one level higher in the file system directory tree.

Sibling

Defines that the component should nest in any component with the same file system root.

Nest Only In

Defines the blueprint that uses the nesting defined on this page. Double-click a blueprint in the Available Component Blueprints column to move it to the Selected Component Blueprints column.

Path from Root

Specifies the component only nests in another component whose file system root is located above it, with the exact relative path between them. Use this option to select a specific nesting match or to restrict nesting to an exact location in the file system directory tree.

Depth from Root

Specifies the component only nests in another component whose file system root is located above it, with the exact depth between them. Use this option to select a specific nesting match or to restrict nesting to an exact depth in the file system directory tree.

Index

F

G

I

J

L

M

N

P

R

S

T

U

V