

CA Service Operations Insight

Service Modeling Best Practices Guide
r3.2



This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time. This Documentation is proprietary information of CA and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA.

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

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

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

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

The manufacturer of this Documentation is CA.

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

Copyright © 2013 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Technologies Product References

This document references the following CA Technologies products:

- CA Application Performance Management
- CA Business Intelligence
- CA Clarity™ Project and Portfolio Manager
- CA CMDB
- CA Configuration Automation (formerly CA Application Configuration Manager)
- CA eHealth® Performance Manager (CA eHealth)
- CA Embedded Entitlements Manager (CA EEM)
- CA Event Integration
- CA Insight™ Database Performance Manager
- CA NSM
- CA Process Automation
- CA Service Desk
- CA Server Automation (formerly CA Spectrum® Automation Manager)
- CA SiteMinder®
- CA Spectrum®
- CA Systems Performance for Infrastructure Managers
- CA SystemEDGE
- CA Virtual Assurance

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: About This Guide	9
Intended Audience	9
Related Publications	9
Local Documentation and Online Bookshelf	11
Chapter 2: Introduction	13
Service	13
Service Concepts	14
Relationships	14
Propagation Types	17
Severity	24
Significance	26
Impact	27
Health, Quality, and Risk	28
Granularity	29
Service Model Types	30
Federated Modeling	31
Chapter 3: Planning Service Models	33
Planning Service Models	33
Service Identification	33
Existing Services	34
Resource Collection	35
Resource Importance	36
Best Practices	36
Chapter 4: Building Service Models	39
Service Models	39
How to Build a Service Model	39
Create the Service	41
Add CIs to the Service	41
Create Subservices	44
Create Groups	45
Set the Granularity Level	46
Assign Propagation Types and Relationships	47

Create a Propagation Policy with Operative Propagation	49
Create a Propagation Policy with Custom Propagation	50
Assign the CI Significance in the Service	53
Create and Assign a Service-Level Agreement	54
Create and Assign an Escalation Policy	55
Validate and Save the Service	55
How Service Validation Works	57

Chapter 5: Customizing Service Model Display 59

How to Customize Service Model Display	59
Control Zoom Level	60
Adjust Model Layout	61
Control Line Layout and Appearance	62
Adjust Background Type	64
Add a Background Image	64

Chapter 6: Editing and Managing Services 67

Modify a Service Model	67
Perform Copy, Cut, Paste, Delete, or Remove Operations on a Service in the Console	68
View Service Information	70
Set Service Priority	72
Set the Service Location for Google Earth	72
Control User Group Service Access	73
Change Multiple CI Relationships Simultaneously	74

Chapter 7: Importing Services 75

How to Import Services	75
Import Services Automatically	75
Import Services Manually	76
Processing Updates for an Imported Service	77
Launch the Source of Imported Services	77

Chapter 8: Using Service Discovery 79

Service Discovery	79
How to Create Dynamic Service Policies	80
Launch the Wizard	82
Define the Service	82
Confirm the Dynamic Service Creation	85
How to Create Automatic Relationship Policies	85

Launch the Wizard	87
Define the Relationship and Source CI Criteria	88
Define the Target CI Criteria	89
Define the Match Criteria	90
Define the Relationship Scope	91
Confirm the Automatic Relationship Policies Creation	91
How to Create Unmanaged Relationship Policies	92
Launch the Wizard	93
Define the Relationship Criteria	94
Define the Source CI Criteria	95
Define the Target CI Criteria	96
Define the Match Criteria	97
Define the Relationship Scope	98
Confirm the Unmanaged Relationship Policies Creation	98
Topology Warnings	99
How to Manage Service Discovery Policies	100
View Service Discovery Policy Details	102
View Service Discovery Policy Set Revision Information	102
Modify Service Discovery Policies	103
Add a Service Discovery Policy	103
Enable or Disable Service Discovery Policies	104
Delete Service Discovery Policies	105
Service Discovery Connector Configuration	105
How to Use Command Line Service Discovery Operations	107
Create a Service Discovery Policy Set File	108
Upload a Service Discovery Policy Set File	108
Undo the Latest Version of a Service Discovery Policy Set	109
Redo an Undo Operation	110

Chapter 9: Creating and Working with Service-Level Agreements 111

How to Create and Work with Service-Level Agreements	111
Define a Service-Level Agreement	113
View SLAs on the Dashboard	116
View Service SLA Summary Report	117
View SLAs Status on the Operations Console Data	117
Edit or Terminate an SLA	118
Disassociate Business Hours in a Service-Level Agreement	118
SLA Example	119

Chapter 10: Working with Customers 121

How to Create and Manage Customers	121
--	-----

Identify Customers, Sub-Customers, and Priorities	123
Identify Services	124
Create Customers.....	124
Create Sub-Customers	126
Configure Customer Priorities and Labels.....	127
Create Escalation Policies and Alert Queues.....	128
View Customer and Sub-Customer Details	129
View Customer Metrics.....	130
View Customers Associated with a Service.....	131
Example: Creating and Working with Customers	132
 Appendix A: Service Modeling Examples and Scenarios	 145
Service Modeling Example 1 - Finance Service	145
Finance Service Resources	145
Finance Service Model	146
Example Finance Service Escalation Flow	147
Service Modeling Example 2- Shopping Cart Service	149
Shopping Cart Service Resources	149
Shopping Cart Service Model	150
Example Shopping Cart Service Escalation Flow	152
Service Modeling Example 3 - Dynamic Service Policy.....	153
Policy - All Running Hardware with Specific Starting IP Address	154
Policy - All Virtual Systems by a Specific Vendor.....	155
Service Modeling Example 4 - Automatic Relationships Policy.....	157
Policy - All Computer Systems and Running Software	157
Policy - All Database Instances and Tablespace	161
Service Modeling Example 5 - Unmanaged Relationships Policy	163
 Glossary	 169

Chapter 1: About This Guide

The *Service Modeling Best Practices Guide* contains information about planning, building, and managing service models in CA SOI. This guide introduces the key service concepts and provides detailed procedures, processes, examples, and scenarios about service modeling. The guide also includes details about creating policies that automatically create and maintain services and relationships between source and target CIs according to specified criteria. Additionally, you can find information about how to create, monitor, and manage service-level agreements.

This section contains the following topics:

[Intended Audience](#) (see page 9)

[Related Publications](#) (see page 9)

[Local Documentation and Online Bookshelf](#) (see page 11)

Intended Audience

This guide is intended for product administrators who are responsible for modeling and managing services in CA SOI.

Related Publications

The following publications, provided on the installation media and the CA SOI online bookshelf, provide complete information about CA SOI:

Administration Guide

Provides information about administering and maintaining the product after installation.

Connector Guide

Provides general information about connectors, the CA Catalyst infrastructure, and writing custom connectors.

Event and Alert Management Best Practices Guide

Provides concepts, procedures, and best practices for managing the event and alert stream that CA SOI receives from connectors.

Implementation Guide

Provides information about installing and implementing the product.

Online Help

Provides information about performing tasks in CA SOI user interfaces.

Readme

Provides information about known issues and information that is discovered after the guides were finalized. A CA SOI release may not have a Readme.

Release Notes

Provides information about operating system support, system requirements, database requirements, web browser support, and international support.

Troubleshooting Guide

Provides information and procedures to diagnose and resolve problems with CA SOI.

User Guide

Provides information for nonadministrative users about using the product, such as responding to alerts and viewing reports.

Web Services Reference Guide

Provides information about the CA SOI web services for interacting with resources such as CIs, services, alerts, relationships, and escalation policy.

The following publications provide information about the CA Catalyst infrastructure:

<Product Name> Connector Guide and Readme

Provides information about a specific CA Catalyst connector, including prerequisites, installation, configuration, data mapping, and known issues. The documentation for each CA Catalyst connector is included with its downloadable package.

Local Documentation and Online Bookshelf

CA SOI provides access to the documentation locally and online.

Local Documentation

The local documentation is installed in the SOI_HOME\Documentation folder and includes the PDFs for all guides. The online help is also installed with CA SOI and accessed through the Dashboard (PC and Mobile) and USM Web View. The local documentation is updated with specific releases only.

Online Bookshelf

The online bookshelf is on support.ca.com and provides the most current documentation set, which can be updated between releases. The online bookshelf also provides the documentation for the latest supported versions of CA Business Intelligence, CA EEM, and CA Process Automation. For a list of Bookshelf updates, click the Update History link on the Bookshelf.

CA SOI provides access to the online bookshelf in the following locations:

- The Dashboard provides a Bookshelf link.
- The Operations Console provides a menu link under Help, Bookshelf.

Note: If you are unable to access the online bookshelf, contact your system administrator to provide the documentation set PDFs.

Chapter 2: Introduction

This section contains the following topics:

[Service](#) (see page 13)

[Service Concepts](#) (see page 14)

[Service Model Types](#) (see page 30)

[Federated Modeling](#) (see page 31)

Service

The concept of a service model, or service, is central to CA SOI. A *service* typically consists of several CIs, which you can group to represent, for example, web server farms or clusters. Services can also contain *subservices* (see page 44), which are subordinate service models. Subservices are previously created services that are reused as building blocks of another service. Service models typically represent high-level abstract entities like a web-based retail transaction service, an application server service, a printing service, or a routing service. You can define any type of service with CA SOI as long as one of the integrated domain managers monitors the service components.

CA SOI provides a comprehensive understanding of how a fault condition, which CA SOI represents as an infrastructure alert, impacts the business. Consider a managed resource such as a router. You can accurately, but narrowly, define it as a device that forwards data from one network to another. From a service perspective, however, a router is an indispensable component among other cooperating components that support interconnected business activities.

When router performance is compromised, the activities that depend on that router are likely compromised also. A router can be associated with other network devices such as switches or servers, which are associated with the applications or databases that they host. These relationships and dependencies comprise the logical and physical topology of the service. CA SOI lets you incorporate these relationships in the service model. These relationships help you capture how one CI relates to another and how they collectively deliver the service logic.

Service models contain policy that determines how alert conditions on one CI can impact related items and the service itself. You can modify and extend this policy to refine the model and capture the collective behavior of all associated entities.

You can reuse a service model any number of times. You can also combine it with other configuration items (CIs) and services to build higher-level service models. For example, the DNS service can be critical to several higher-level services such as Microsoft Exchange and SAP. Similarly, Exchange can itself form part of still higher-level services such as email or Blackberry.

You can define new service models, import them from domain managers, or define policies that automatically discover and create services according to specified criteria. For example, an operator working with a service owner can select configuration items that are discovered through integration with the domain managers. The operator can then create relationships among those configuration items. Similarly, if a service model is defined in a domain manager, you can import that service model and all its topographic information directly into CA SOI. You can extend or combine imported service models in the same manner as the service models defined in CA SOI. This ability provides a powerful mechanism to leverage your existing investment.

Service Concepts

The concepts in this section help you understand how CA SOI monitors services and calculates service impact. The following entities play a role in impact calculation:

- Alert and CI conditions from domain managers display as [severity](#) (see page 24) in the appropriate CI and alert.
- CI severity affects related CIs in service models according to [propagation type and policy](#) (see page 17) settings.
- If propagation settings cause severities to propagate, the multiplied value of severity and [significance](#) (see page 26) determines the related CI impact and, ultimately, service [impact](#) (see page 27).

Relationships

Relationships in a service model show how CIs are linked to form the service topology. A relationship between linked objects has a semantic (name or type, for example 'HasAccessTo') and a propagation type (for example, 'Custom'). CA SOI relationships correspond to instances of USM BinaryRelationship type. If you import a service from a connector, its relationships in the domain manager are mapped to USM relationships.

You can assign relationships to every link between objects in a service model. To assign relationships, you select the appropriate propagation type and then select a relationship from the list of relationships that map to that propagation type.

The available USM BinaryRelationship semantics are as follows:

Note: For details about each type or information about relationship updates, see the USM schema documentation. For information about how to access the schema documentation, see the *Connector Guide*.

Has Access To

Specifies that a CI accesses another CI's functionality. Use this relationship to indicate that a resource can access another resource, or that software communicates with or accesses a specific entity, such as a database. This relationship differs from Has Requirement For, which indicates a mandatory presence of the target CI for the source CI to function.

Has Contact

Specifies that a CI plays a specific contact role for another CI, such as an owner or assignee.

Has Detail

Specifies that a CI provides additional information for another CI. For example, an Asset CI can relate to another CI that provides more detail about the asset.

Has Member

Specifies that a child CI is a member of another CI. For example, several User CIs can be members of a user group. This relationship is the default assignment.

Has Requirement For

Specifies that a CI requires the existence of another CI, its operation, or both. For example, an Application CI can require the operation of an Application Server CI.

Is Affected By

Specifies that a CI impacts another CI and that custom policy defines the impact. Whereas the Has Requirement For relationship causes impact when the target CI is in a critical or down state, Is Affected By lets you configure the scenarios that impact the source CI. For example, a web server farm group is impacted if 40 percent of its ComputerSystem CIs are down.

Note: CIs in maintenance mode are excluded from custom policy calculations that are related to average or percentage.

Is Bound To

Specifies a symmetric relationship where two CIs are intrinsically linked, so that one cannot function without the other.

Is Cause Of

Specifies that a ChangeOrder is the cause of a Request, Incident, or Problem being opened.

Is Clone Of

Indicates that the source CI is a clone of the Target CI and that the two elements are synchronized.

Is Composed Of

Specifies a compositional relationship where a source CI is the aggregate of several target CIs.

Is Connected To

Indicates a network connection carrying data between the source and target CIs, such as between physical ports or application components.

Is Discovered By

Indicates that the specified ManagementAgent target CI discovers and manages the source CI.

Is Evolution Of

Indicates that the source CI is evolved from the target. Examples include next generations of hardware or software.

Is Hosted By

Indicates that the target CI hosts the source CI. This relationship is the inverse of Is Host For.

Is Host For

Indicates that a source CI is hosting a target CI. For example, a ComputerSystem CI can host a VirtualSystem or RunningSoftware CI.

Is Impacted By

Indicates that another CI impacts or affects the source CI. For example, a Memory child CI or Port child CIs affect the parent ComputerSystem CI.

Is Instance Of

Specifies that a source CI is an occurrence of a target CI. For example, a ProvisionedSoftware CI can run an instance of a RunningSoftware CI.

Is Location For

Specifies that a source CI defines the target CI location.

Is Manager For

Specifies that a source CI controls or manages a target CI. For example, a DatabaseInstance CI can manage a Database CI.

Is Request For

Specifies that a source CI has been created to create, provision, or otherwise handle the target CI.

Is Resolved By

Specifies that the specified ChangeOrder target CI resolves or corrects a Request, Incident, or Problem source CI.

Is Result Of

Specifies that a source CI is created as a result of an automated workflow or manual processing of a target CI.

Relationships display as color-coded arrows between two objects in the Topology view. The Topology view is available on the Operations Console and the Service Modeler. Propagation types of relationships define CA SOI derives CI and service impact.

More Information:

[Propagation Types](#) (see page 17)

Propagation Types

Propagation type defines how CA SOI derives the impact when conditions change in related objects. Every relationship in a service model has a propagation type, and each USM relationship type maps to a specific propagation type in CA SOI. The propagation types are as follows:

- [Aggregate](#) (see page 17)
- [Bound](#) (see page 19)
- [Custom](#) (see page 19)
- [Operative](#) (see page 20)

Relationships are depicted as arrows between two CIs in the Topology view. The color of the arrow reflects the propagation type of the underlying relationship. The Topology view is available on the Operations Console and the Service Modeler. The first letter of the propagation type is a label on the arrows (A, B, C, or O for Aggregate, Bound, Custom, and Operative, respectively).

Note: Sometimes the arrows are animated red dashes to indicate the root cause of a situation.

Aggregate Propagation

Aggregate propagation indicates a general-purpose relationship that propagates impact from one CI to another. Typically, an association between CIs that are part of a higher-level CI uses this propagation type. CA SOI uses the highest impact condition of all aggregate child CIs when calculating the impact on a parent CI.

Usage

Use aggregate propagation in the following situations:

- When child CIs individually affect a parent CI
- When a parent CI contains multiple child CIs that could potentially impact its condition

Examples

Examples of appropriate situations to use aggregate propagation include the following:

- A CPU or Memory CI may have aggregate propagation to the Operating System on their server.
- Services or components that make up an application may have aggregate propagation to their parent Application CI.
- The components of a network (Router, Bridge, Network Interface Card, and so on) may have aggregate propagation to a high-level Network service, CI, or group.

USM Relationships

The following USM relationships map to aggregate propagation by default in CA SOI:

- Is Impacted By
- Is Composed Of
- Is Host For
- Is Location For
- Is Manager For
- Has Member
- Has Access To

You must select one of these relationships to effect the aggregate propagation behavior. For example, you could assign Is Host For to a relationship between a ComputerSystem CI and its hosted RunningSoftware and VirtualSystem CIs. Or you could assign Is Composed Of to a ComputerSystem CI and its compositional parts. When you add services and resources to a service model, the default propagation type is aggregate with a relationship of Has Member until you change it to something else.

Bound Propagation

Bound propagation indicates a bidirectional relationship between two CIs. If one CI is bound to another CI and either CI has a severity change, the change results in the same impact on both CIs.

Usage

Use bound propagation in the following situations:

- When a fault condition that affects one CI equally affects the other CI
- When two CIs are on the same hierarchical level and cannot function without each other

Examples

Examples of appropriate situations to use bound propagation include the following:

- Mirrored disks may require bound propagation with one another.
- A Database CI connected to an Application CI may have bound propagation with a replication database CI that belongs to a related replication server.

USM Relationships

The following USM relationships map to bound propagation by default in CA SOI:

- Is Bound To
- Is Clone Of
- Is Connected To
- Is Result Of

You must select one of these relationships when you assign bound propagation. For example, you could assign Is Bound To to a set of mirrored disks. When bound propagation is automatically assigned to a relationship, the default relationship assignment is Is Bound To.

Custom Propagation

Custom propagation indicates that one CI may depend on several other CIs for some behavior or function. Custom propagation lets you specify policy that defines when and how to change the severity value of a parent item in a dependent relationship. If the policy is not met, there is no impact to the CI or service.

Usage

Use Custom propagation in the following situations:

- When a group of CIs work together to deliver an aspect of the service
- When a parent CI is not directly affected by a fault condition in an individual CI and is instead affected by the combined performance of all child CIs

Examples

Examples of appropriate situations to use custom propagation include the following:

- A Group CI for a web server farm may require custom propagation with its servers to indicate that the farm can lose 30% of its servers and still function.
- A clustered server configuration with failover capacity may require custom propagation to indicate that a certain number of failures are permitted as long as the failover servers are functioning.

USM Relationships

The following USM relationships map to custom propagation by default in CA SOI:

- Is Affected By
- Is Evolution Of

You must select one of these relationships to effect the custom propagation behavior. The relationships define that the impact from one CI to another is derived through configurable policy.

Custom propagation requires you to [create a propagation policy](#) (see page 50) that defines how to calculate the impact from a group of related CIs to a parent CI. You can configure CA SOI to automatically assign custom propagation and the associated propagation policy when a CI is added to a group that already uses custom propagation instead of the default aggregate propagation.

When you assign custom propagation policy to a CI and the threshold is reached, the target CI attains the specified severity value and an infrastructure alert with the appropriate severity. However, because the alert applies to the CI as it relates to the parent service, the CI icon does not change color as expected, and the severity propagates to a separate impact calculation on the service, which is reflected in the service icon color.

Operative Propagation

Operative propagation indicates that the related item is affected only if the impact value of a CI exceeds a defined threshold. Operative propagation lets you specify policy that defines the impact value that the CI must exceed for impact to propagate to the parent CI. If the policy is not met, there is no impact to the CI or service.

Usage

Use operative propagation in the following situations:

- When a CI is only needed to be available, even if it is not performing well
- When the service is not affected unless the state of the CI is serious
- When you are not concerned with minor state changes as long as the CI is still running

Examples

Examples of appropriate situations to use operative propagation include the following:

- An application server may have operative propagation with a database server; the application server requires the database server to provide data storage and retrieval functions. As long as the database server is up, it is assumed that the application server is being served adequately to perform its function.
- An operating system may have operative propagation with the CI for its memory. In this case, operative propagation can prevent impact from the memory CI from propagating to the operating system unless the memory usage exceeds a defined threshold.

USM Relationships

The following USM relationships map to operative propagation by default in CA SOI:

- Has Requirement For
- Has Contact
- Has Detail
- Is Instance Of
- Is Request For
- Is Cause Of
- Is Discovered By
- Is Hosted By
- Is Resolved By

You must select one of these relationships to effect the operative propagation behavior. For example, you could assign Has Requirement For to an Application CI that requires an ApplicationServer CI to be running.

By default, operative propagation requires an impact value of 20 to propagate impact, and the default relationship is Has Requirement For. You can [create propagation policy](#) (see page 49) to define a different impact threshold.

Configure Default Propagation Policies

Each relationship type has a predetermined [propagation type](#) (see page 17). You can view the relationship mappings and change the parameters of default propagation policies for each relationship that maps to custom or operative propagation.

Important! When you change the default propagation policy parameters for a relationship type, any instance of that relationship in a service model that has not been manually overridden acquires that policy's parameters by default.

The changes will affect all new relationships/propagations from that moment on, but will not cause the reevaluation of the policies currently in SA Manager's memory. To effect that, you have to restart the CA SOI Application Server.

You must close the Service Modeler to change default propagation policy.

Note: A propagation policy determines how [impact](#) (see page 27) is derived and propagated. For more information about each policy type, see [Define Operative Propagation Policy](#) (see page 49) and [Define Custom Propagation Policy](#) (see page 50).

Follow these steps:

1. Select Tools, Propagation Policies from the Operations Console.

The Default Propagation Policy dialog opens.

Note: If the Service Modeler is currently active (in any user session), a warning dialog opens stating that the Default Propagation Policy dialog is read-only. Close all Service Modeler sessions and reopen the dialog to make changes.

2. Click the set link next to the value in the Propagation Rule column for the Is Affected By or Is Evolution Of relationship row, which have custom propagation defined.

The Default Thresholds for Custom dialog opens.

3. (Optional) Select the Automatically Maintained check box to automatically change the relationship to Is Affected By or Is Evolution Of and apply the custom propagation policy when new CIs are added to a group that uses custom propagation. Auto-maintenance applies to relationships added by both manually editing a service and through automation by service discovery or service import.

Note: You can set Automatic Policy Maintenance for the Modeler operation to on by default in the Set Preferences dialog.

4. Select one of the following from the Policy Type drop-down list:

Average

Sets the impact of the parent item based on the average impact values of CIs associated with the policy.

Percentage

Sets the impact of the parent item based on a percentage of CIs that have the impact specified in the rule.

Any

Sets the impact of the parent item when any CIs associated with the policy have the impact specified in the rule.

All

Sets the impact of the parent item when all CIs have the impact specified in the rule.

Default: Any

The text of the rules on the dialog changes to reflect the selected policy type.

5. Complete the following fields and settings for each rule that you want to create:

% of items

(Percentage type only) Defines the percentage of items that must exceed the impact threshold to meet the rule criteria.

Threshold

Defines the impact threshold. The appropriate amount of CIs (as defined by the policy type) must meet or exceed the impact threshold to meet the rule criteria. Define a number between 0 and 40. The impact numbers translate to the following impacts:

- 0: Operational
- 1-10: Slightly Degraded
- 11-20: Moderately Degraded

- 21-30: Severely Degraded
- 31-40: Down

Set Severity

Defines the severity to assign to the parent item if the rule criteria are met.

The text for each rule changes to reflect the new settings. Define as many of the four available rules as the policy requires.

6. Click OK.

The rule changes in the Propagation Rule column for the Is Affected By or Is Evolution Of relationship. You can assign different default policy parameters for each relationship type. The new default custom policy is assigned to all new custom propagation assignments and existing custom propagations that use the associated relationship, including those that previously used a default custom policy. Customized default propagations are not recalculated; they take effect from the point they are customized.

Note: CIs in maintenance mode are excluded from custom policy calculations related to average or percentage.

7. Click set next to any value in the Propagation Rule column for relationship rows that map to Operative propagation.

The Default Threshold for Operative dialog opens.

8. Set the default threshold value at which impact starts to propagate to parent CIs, and click OK.

The rule changes in the Propagation Rule column for the relationship, and the new default operative policy is assigned to all new and existing operative propagation assignments of the specific relationship, including those that previously used a default operative policy.

9. Click Save.

The propagation changes are saved.

Changes take effect for ensuing alerts by default. To recalculate the impact of existing alerts, restart the CA SOI Manager and CA SOI User Interface.

Severity

Severity indicates the condition of a CI as reported from the domain manager to CA SOI through alerts. If multiple domain managers send alerts for the same CI, the highest severity is used. CI severity helps determine the service impact by propagating the impact of the condition to related CIs in the service model according to propagation settings.

The following table describes each severity:

Severity	Color	Description
Normal	Green	Operational
Minor	Yellow	A nominal displacement of CI function that can require an inspection
Major	Orange	A serious causal change typically leading to degradation of function
Critical	Red	High probability of imminent failure and severe degradation of service
Down	Burgundy	The CI is incapable of providing function or service

Color-coded icons on the Operations Console indicate CI severity (the color-coded icons for services indicate the service impact). Alerts in the Contents pane have the color corresponding to their severity. The Navigation pane also represents severity in columns next to services and CIs. The following graphic shows that each column lists the number of items with the corresponding severity (represented by the colors in the previous table).

The screenshot shows the 'Navigation' pane with the 'Services' tab selected. It displays a tree view of services with columns for severity counts. The columns are color-coded: Green (Normal), Yellow (Minor), Orange (Major), Red (Critical), and Burgundy (Down). The data is as follows:

Name	Green	Yellow	Orange	Red	Burgundy
Services (LODL...	87	6...	62	39	
Appended Se...	11	7	2	3	
Active Direct...		4	4		
Asia Banking ...					
Applied Kineti...					
Canada Ware...					

Note: If no alerts are raised for a CI, its severity is green even if the device contains child CIs with different severities. Also, groups are simply containers and would not usually have alerts. You can expand the tree and can follow the numbers to the row that lists the item whose severity you are looking for.

Significance

Significance indicates the importance of a CI. In a service, many CIs can affect another CI or service. Although each CI affects the health of the service, some CIs can be more important than others. For example, it is important that a print server is available for an online order service. However, the print server does not affect order processing as much as the inventory database does. In this case, the print server has a lower significance than the inventory database.

Significance is a property of a CI in the context of a service. Each CI type has a default significance, but the actual significance is stored in the relationship, thus providing significance with a service scope. When you change significance for an individual CI in a service, the CI's significance value only changes within that service. If the CI belongs to other services, it retains its previous significance value in those services. This behavior ensures that you can maintain different significance values for the same CI if it is more or less important to other services.

CIs can be in as many services as needed and can affect each service differently, based on the significance setting. Suppose that the Payroll department also uses the print server to print checks. Therefore, the print server has much higher significance.

Significance is a value from 1 through 10, where 1 is the least significant and 10 is the most significant. The numeric value denotes how important a child object (antecedent) is to the functioning of its parent object (dependent). Each CI is assigned a significance when it is imported from a domain manager.

CA SOI assigns a default global significance value to every CI type available in the product. For instance, all servers of a Computer System type added to a service model have a significance of 5, while switches and routers have a value of 9. The types Service, Network, and Operating System are considered to be the most important and have a significance of 10.

You can set significance globally or for individual CIs and relationships.

More Information:

[Assign the CI Significance in the Service](#) (see page 53)

Impact

Impact indicates how much a CI affects a service and related CIs. The following factors determine impact:

- CI severity
- CI significance
- Propagation type and policy

Impact provides IT personnel with a good understanding of what fault conditions really mean to the services that CIs support.

CA SOI calculates the impact by multiplying [severity](#) (see page 24) by [significance](#) (see page 26). Significance is a number from 1 (lowest) to 10 (highest), and severity ranges from 0 (Normal) to 4 (Down). Therefore, the highest possible impact is 40. Consider an application with a severity of 4 and a significance of 4. The resulting impact is 16.

Note: When a custom propagation policy connects CIs, you can define complex rules for changing the severity value. For more information, see [Define Custom Propagation Policy](#) (see page 50).


Consider the following items:


- If you change the significance of an item, the impact changes after a new alert is received.
- [Priority](#) (see page 72) is used in the calculation of impact instead of significance in the following situations:
 - The service is a top-level service
 - The significance of the parent relationship is zero

The following table defines the impact ranges:

Impact	Color	Description
0	Green	Operational
1-10	Yellow	Slightly Degraded
11-20	Orange	Moderately Degraded
21-30	Red	Severely Degraded
31-40	Burgundy	Down

The Topology tab of the Contents pane shows the impact color as a small box above and to the right of a service, CI, or group. Mouse over the box to display the impact value.



The impact number shows on the small box when you click Chart Display Complexity Level  and switch to the Advanced view.

The color of the service icons in the Operations Console indicates the quality impact of the service. *Quality impact* is the inverse value of service health, which displays on the Dashboard and the Component Detail pane of the Operations Console.

Health, Quality, and Risk

Health, quality, and risk are the primary metrics exposed to Dashboard and external interfaces for monitoring service status. They categorize service impact values to reflect the type of outage or impact according to alert categories.

Alerts impacting a service belong to one of the following categories:

Quality

Indicates the level of excellence that consumers of an IT service experience, whether they are other IT services, customers, or end users. The quality levels are Operational, Slightly Degraded, Moderately Degraded, Severely Degraded, Down, and Unknown. The highest propagated [impact](#) (see page 27) of an associated quality alert determines the service quality value.

Risk

Indicates the likelihood of delivering the quality of service required to support the overall business objectives. The risk levels are Down, Severe, Moderate, Slight, None, and Unknown. The highest propagated [impact](#) (see page 27) of an associated risk alert determines the service risk value. If an alert has no defined type, it is a risk alert by default.

Service health is the highest impact held by quality or risk. The following table shows the available Health, Quality, and Risk values:

Health	Quality	Risk
Normal	Operational	None
Minor	Slightly Degraded	Slight
Major	Moderately Degraded	Moderate
Critical	Severely Degraded	Severe
Down	Down	Down

For example, a slightly degraded service with a severe risk of degradation would have a service health of Critical.

Granularity

CA SOI supports granularity at two levels: Normal and Low. Normal granularity mode represents an explicit modeling principle where alerts are presented in CA SOI only if all the impacted CIs are included in the model. For example, a computer system CI that is running many service-supporting resources. When the service granularity is Normal, you include all the resources in the model to show their alerts as impacting the service.

Low granularity mode represents a mixed modeling principle where you manage the service granularity as follows:

- You include only the parent CIs and no associated child CIs in the service. In this case, parent CIs act as aggregators for all alerts that affect them directly or indirectly through their commonly related resources. For example, consider the same scenario of a computer system running various service-supporting resources. In this case, you only include the computer system, and any alerts affecting any of the resources hosted on the computer system are aggregated to the computer system.
- You include the parent CIs and specific child CIs (not all child CIs) in your service. In this case, any alerts that are directly associated with the included child CIs are aggregated to those CIs. The alerts that are associated with the nonmodeled child CIs are aggregated to the corresponding parent CIs. For example, you include a computer system CI and its child CPU CI in your service model. If any CPU-related alert comes in, it is associated directly to the CPU. Any other alert not directly associated with the CPU gets attached to the parent computer system CI. Therefore, including granular child CIs does not change the low granularity of the service model for excluded child CIs, enabling mixed-mode low granularity modeling.

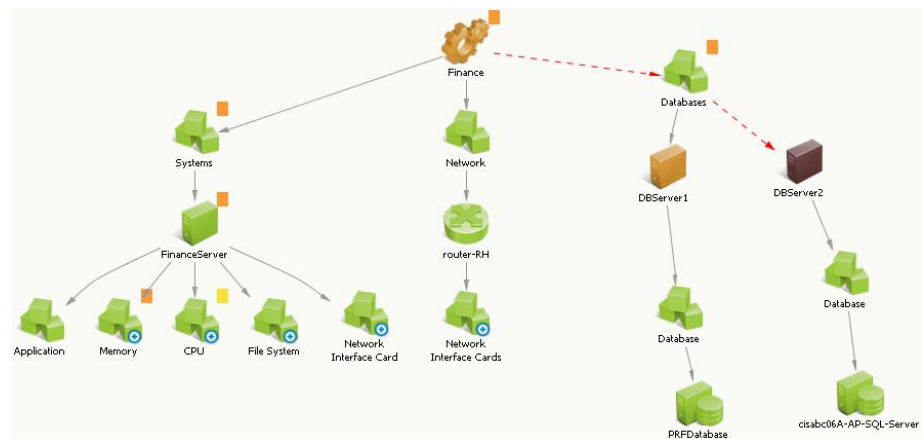
Service Model Types

The most common types of service models are as follows:

Bottom-Up

Bottom-up service models are constructed from a domain-oriented perspective (network, systems, applications, and so on). These models take less time to define and are easier to understand. They focus on impact analysis, not root cause, and are the most common starting point. Most service models that are imported from domain managers are bottom-up, because they define a service from the perspective of their managed domain.

The following graphic shows an example of a bottom-up service model:




Notice that this bottom-up example is modeled based on groups that define distinct domains: Databases, Network, and Systems. Separate products can manage these domains, and the model combines the intelligence of these domain managers to create a service representation for a Finance department. The model uses groups to aggregate domains and specific objects within those domains (such as the network cards in a router). Note how the model is able to pinpoint the root cause of Finance service degradation as the database server DBServer2.

Note: For a detailed scenario that builds on this example, see Service Modeling Scenarios and Examples.

Top-Down

Top-down service models are organized based on a logical business service topology, not by domain. These models are more advanced. They require a better understanding of the logical structure of your enterprise and of the capabilities of service modeling in CA SOI. They ultimately provide better root cause information than bottom-up models.



This example uses subservices to provide a solid representation of the business-oriented topology of the service. The items in the specific subservices may or may not be specific to domains. Impact still propagates from subservices to the top-level service.

Federated Modeling

A *federated modeling* approach indicates that intelligence is distributed across integrated products. The network management tool has intelligence around the network model, the application management tool has intelligence around applications and transaction models, and so on. CA SOI combines the collective intelligence of all integrated products to derive the overall impact of any CI condition on modeled services.

The benefits of federated modeling are as follows:

- Distributed models are scalable.
- Models in each environment are optimized to suit their role within that environment.
- Leveraging existing models lowers the cost and effort of implementation.
- Models are flexible and extensible to meet the specific detail level that each service requires.

Consider the following example of federated modeling in practice:

- A service in CA SOI models an application and its dependency to a host.
- CA Spectrum models the same host and its dependency to the network and systems infrastructure.
- When a change to the operational state of the infrastructure impacts the host, CA SOI is notified through its CA Spectrum connector. CA SOI can identify the root cause of the problem using the intelligence that CA Spectrum provides without having to model the underlying infrastructure and topology in the service.

Chapter 3: Planning Service Models

This section includes information about how you can plan your service models.

This section contains the following topics:

[Planning Service Models](#) (see page 33)

Planning Service Models

Service models are the basis of all administrative and management functions in CA SOI. Managing your enterprise from the perspective of business services moves management away from domain-specific management to consolidated, holistic, and practical views of all IT resources. To accomplish this comprehensive business service management, first gain an accurate picture of your enterprise and what resources make up discrete business services.

This section describes the information you must collect to model accurate, end-to-end business services using the Service Modeler.

Service Identification

Any set of resources that depend on each other to provide a useful business function can compose a service. Identifying services requires an understanding of the resources that make up your enterprise and the typical contents of a service. A service can take many forms, including the following:

- A low-level IT service such as DNS or DHCP
- A set of resources that you want to manage collectively, such as the servers or printers contained in a specific location
- A network service such as VPN
- A systems service such as Active Directory
- A database service such as a Microsoft SQL Server cluster
- A high-level business service such as Payroll, Email, and so on
- External services such as an Internet Service Provider or Cloud

Most services are a combination of all of the above, with high-level business services containing various low-level services. For example, a Blackberry business service may contain subservices that represent key IT services that support Blackberry communication, such as Active Directory, Exchange, DHCP, and so on.

Most modern, complex service models do not, by definition, follow a hierarchical tree-like structure. They follow a more dynamic, multi-layered structure that requires you to collect comprehensive information about what makes up your services. As you begin identifying the services in your enterprise that require modeling, you can consult the following sources for the required information:

- Department leads and line of business owners
- Your IT department
- Existing management products
- Old design documents, statements of work, and so on
- Service Catalog, defined workflows
- Application configuration management tools
- Existing help desk or CMDB products

Other sources may also be useful. In a large enterprise, you rely on all available sources of information to gain an accurate picture of your services.

Existing Services

You may have other management products that manage services or a concept similar to services. You can [import services](#) (see page 75) from many integrated products into CA SOI as full service models. Examine your existing management products for services so you can leverage your investment in these products and avoid having to model a similar service from scratch.

Imported services are synchronized with the source definition when the source changes, and are reconciled into a single service where they have been modeled within multiple domains.

For more information about the connectors that support importing services from their domain manager, see the product-specific *Connector Guide* included with each connector package.

Resource Collection

As you define the services in your enterprise, you also define the specific resources that compose those services. Services can contain everything from applications and databases to servers and network devices. If some resources are omitted from a service model, you cannot ensure that the model is an accurate representation of the service health and availability.

Consider the following as you collect the resources that comprise your services:

- Find out where all required resources are currently managed. Verify that all resources in domain managers for which CA SOI provides connectors are actively managed, so that you can include the resources in service models. Consider a scenario where an application for which CA SOI does not provide a connector manages the resources. In this case, you can build a custom integration with the application using various tools provided.

Note: For more information about implementing provided connectors and building custom integrations, see the *CA Catalyst Implementation Guide* or *Connector Guide* for each specific connector.

- Organize the resources that make up a service into logical subcomponents. Most complex services are sets of subservices or groups. If you can define these sets before modeling the service, you can improve the presentation and performance of your service. You can also reuse these modular components across multiple services.
- Consider the relationships among resources within subcomponents and the service at large. Try to separate subcomponents based on the following high-level relationship types:
 - Resources where the only concern is if a component is providing functionality
 - Resources that are critical to the functioning of a service
 - Resources that have some association or form a group, where the resources in the group can be operated on collectively
 - Resources that are intrinsically dependent on each other

Note: For more information about relationships and how impact is propagated for each relationship, see Relationships and [Propagation Types](#) (see page 17).

- For basic collections of resources or for areas of your enterprise with a high level of volatility, organize resources and note the criteria by which you can group them. Resources that fall into these categories can be good candidates for service discovery policies that automatically create services and relationships based on policy criteria.

Note: For more information, see [Service Discovery](#) (see page 79).

Resource Importance

You must determine the importance of services and the resources in a service. Rely on sources of information such as business continuity plans, disaster recovery plans, and service-level agreements to obtain the services that are most important to your enterprise.

Typically, more information is available about business-critical services, making them the ideal place to start service modeling. Not all services require the same level of complexity. A critical business service may be documented in enough detail for you to create a comprehensive service model, while for a less critical service, a simpler model will suffice. CA SOI provides the flexibility for models of different detail levels, and you can make models more comprehensive as your investment level in the product grows.

Also determine the importance of the resources that make up a service, so that this importance is accurately reflected when issues occur.

Best Practices

As you begin modeling services in CA SOI, consider the following best practices:

- [Determine and start with the services that matter the most to your enterprise](#) (see page 36). The most important services typically contain the most supporting documentation.
- Organize resources into [groups](#) (see page 45) and [subservices](#) (see page 44) wherever possible to increase the modularity of your service. You can reuse groups and subservices across multiple services. Logical groupings make it easier for you to create a service model whose layout, root cause, and service impact is easy to comprehend.
- Start with bottom-up service models, either imported from domain managers or modeled organically, and move to top-down in a phased approach.
- [Take advantage of federated modeling concepts](#) (see page 31), which let you start with relatively simple service models and still obtain root cause information using the distributed intelligence of integrated products.
- Take advantage of [low granularity modeling](#) (see page 46), so that less detailed models can automatically aggregate alerts from related unmodeled CIs and immediately return results similar to those of more comprehensive models.
- Take advantage of [service discovery](#) (see page 79) to automatically add services or create relationships based on criteria that define logical groups of resources and object relationships in your enterprise. Services created by service discovery update dynamically according to changes in your enterprise.
- Gradually move to more comprehensive models as your usage of the product grows.

- Even as your service models become more detailed, keep them as lightweight as possible so that they include only the essential information to support reconciliation, service impact, and root cause.
- As you build the service model, save or validate often, so that potential errors are easier to resolve.

Chapter 4: Building Service Models

This section provides information about how you can build your service models.

This section contains the following topics:

[Service Models](#) (see page 39)

[How to Build a Service Model](#) (see page 39)

[How Service Validation Works](#) (see page 57)

Service Models

Service Models are the basis of all administrative and management functions in CA SOI. Service models represent high-level abstract entities; for example, a web-based retail transaction service, a printing service, and a routing service. Service models help you manage your enterprise from the perspective of business services, providing consolidated, holistic, and realistic views of all IT resources.

Each service model consists of the following entities:

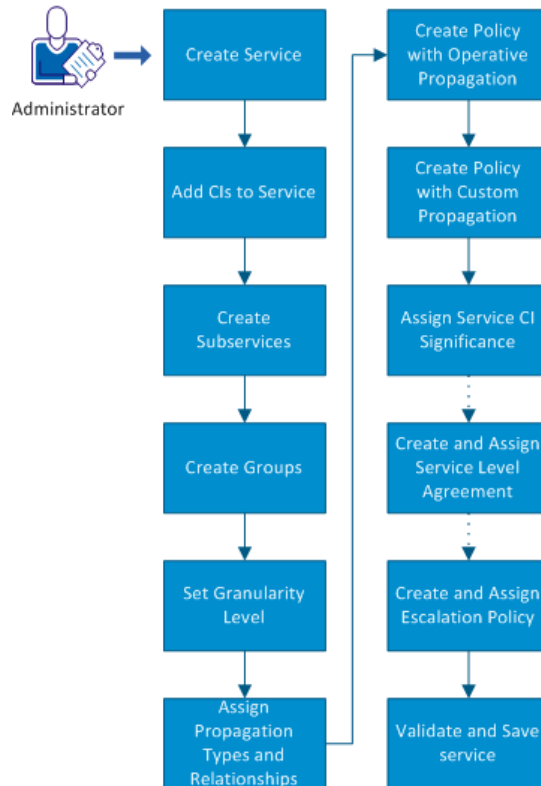
- A CI that represents the service itself.
- Child CIs that represent IT elements that support the service or measure and manage some aspect of its behavior.
- Relationships that determine how CIs interact and depend on one another.
- Set of propagation policies that determine how impact from a CI fault condition propagates to related CIs.

How to Build a Service Model

As an administrator, you build service models in CA SOI. Building service models is a multi-tiered process that lets you define the associated CIs, relationships, and propagation policies, and optional service-level agreements and escalation policies. You can add any CI that an integrated product manages to a service in CA SOI. You can also set the granularity level, which determines whether CA SOI aggregates alerts from child CIs not included in the model.

Use this scenario to guide you through the process:

How to Build a Service Model



1. [Create the Service](#) (see page 41).
 2. [Add CIs to the Service](#) (see page 41).
 3. [Create Subservices](#) (see page 44).
 4. [Create Groups](#) (see page 45).
- Note:** Subservices and groups act together to deliver specific aspects of the service.
5. [Set the Granularity Level](#) (see page 46).
 6. [Assign Propagation Types and Relationships](#) (see page 47).
 7. [Create a Propagation Policy with Operative Propagation](#) (see page 49).
 8. [Create a Propagation Policy with Custom Propagation](#) (see page 50).
 9. [Assign the CI Significance in the Service](#) (see page 53).
 10. (Optional) [Create and Assign a Service-level Agreement](#) (see page 54).
 11. (Optional) [Create and Assign an Escalation Policy](#) (see page 55).
 12. [Validate and Save the Service](#) (see page 55).

Note: You can perform several operations while modeling a service to customize the service display. For more information about these operations, see *Customizing Service Model Display*. You can also associate services with customers to determine how services impact the associated customers. For more information about customers, see the *Administration Guide* and the Customer-centric View of Services section.

Create the Service

You create and edit services in the Service Modeler.

Follow these steps:

1. Open the Operations Console and select Tools, Create New Service.
The Service Modeler opens. An icon labeled New Service appears on the Topology tab. A high-level list of classes appears on the Objects tab in the left pane.
2. (Optional) Click the New Service icon and move it to a different location on the Topology grid.
This may be necessary for services that will contain many resources.
3. Double-click the icon and enter a new name.
The service is renamed.
4. Click Save.
The service is saved.

Add CIs to the Service

You add CIs to services to define the service content. You can add any CI that is managed by a running connector to a service. The Service Modeler provides all CIs available for inclusion in services in the Objects tab. If you do not see an expected CI on the Objects tab, verify that the source domain manager is actively managing the resource.

Note: You can also add resources not managed by a running connector using the Sample connector. This method lets you test connector functionality, and potentially spur discovery and management of created entities in domain managers. For more information about the Sample connector, see the *CA Catalyst Implementation Guide*.

The Service Modeler provides simple drag-and-drop functionality for quickly adding CIs to a service. When you add a CI, it automatically establishes a relationship with its parent CI, whether that is the top-level service CI, a group, or another CI.

Consider the following as you add CIs to a service:

- If you are adding CIs that you want to be a part of a subservice, you must [add them directly in the subservice](#) (see page 44). You can add CIs to a group when you [create the group](#) (see page 45) in the main service.
- Consider whether to make your service model bottom-up or top-down. For bottom-up service models based on specific domains, you can add services imported from a domain manager to the service model as subservices.
- Add only the necessary level of detail. [Federated modeling](#) (see page 31) intelligence lets you create simpler models (for example, a ComputerSystem CI without its associated ports, CPU, memory, and so on). You also may be interested in service impact only up to a specific CI level, or you may want to start with a simple service model and build its level of detail over time.
- You can customize the modeler display to best suit your service size, hierarchy, and other factors. For more information, see Customizing Service Model Display.

Follow these steps:

1. Do any of the following to locate CIs in the Service Modeler:

- Select View, Locator.

The Locator dialog opens. You can search for CIs based on property values. Right-click a CI in the search results and select Add to Modeler to add the CI to the service.

- Enter text by which to filter the available CIs in the Filter field.

The Objects tab displays only the CIs that meet the filter criteria in the containing folders.

- Select the browsing criteria in the Browse By drop-down list.

The Objects tab sorts CIs by the criteria you selected. You can browse by the following classifications:

Top-level Classes

Displays a list of folders sorted by high-level classes such as ComputerSystem, Group, Router, Service, and so on. The CIs in these folders may include child components in subfolders.

All Classes

Displays a list of folders sorted by class name.

Data Source

Displays a list of folders sorted by the source connector name. If a CI exists in three different integrated domain managers, it appears in all three data source folders.

Each connector entry contains a five-digit ID number defined by the USM schema. For a list of connectors and their ID numbers, see the *Connector Guide*.

Monitored Services

Displays a list of existing services. You can drag entire services into the service as subservices or expand the service to find included CIs.

Monitored Groups

Displays a list of existing groups. You can drag entire groups into the service or expand the group to find included CIs.

The number displayed next to each folder only reflects the top-level CIs in the folder. It does not include CIs contained in subfolders.

2. Select a CI to include in the service.

The following tabs below the Objects tab display information about the CI:

Information

Displays a subset of basic CI properties.

USM Properties

Displays the full list of USM properties for the CI.

USM Notebook

Displays a comparison of the USM properties taken from all managed domains and the USM reconciled sheet, which reconciles the properties from all managed domains into a single set of properties.

3. Do any of the following to add the CI to the service:

- Drag the selected CI in the service topology.

The CI is added to the service. If you place the CI in a blank area of the topology pane, it becomes associated to the top-level service CI. To associate the CI with a different CI, drag it on top of that CI.

- Right-click the CI and select Add with sub-components.

The CI and all child CIs are added to the service model associated with the last CI selected on the topology pane. By default, child CIs are not added to the service when you drag a parent item into the topology. If you always want to add child CIs when you drag a CI into the topology, you can change this preference in the Operations Console.

- Select multiple CIs with Shift+click or Ctrl+click and drag them in the topology pane in one of the following ways:
 - Do not lift the left mouse button after making the last selection. Keep the button pressed and drag the CIs.
 - Drag the selected items with the right mouse button.

The CIs are added to the service model associated with the top-level service CI or the CI on which you dragged the items.

All CIs added to a service obtain the default Has Member relationship and aggregate propagation with connected CIs, unless you add them to a group with [automatic policy maintenance](#) (see page 50) enabled.

Create Subservices

Services can have any number of subservices. Subservices let you group CIs and relationships in a modular fashion. Create a subservice for any set of CIs that can operate as a standalone service and that you may use in other services. For example, an Email service may contain subservices for key functions such as Exchange, Active Directory, and so on.

You can reuse created subservices in multiple services, and you can use an existing service as a subservice in other services. Subservices are important entities in top-down service models, where modeled services accurately represent the overall business service topology. Most top-down service models contain multiple subservices.

Follow these steps:

1. Open the Operations Console, right-click the service in the Navigation pane that needs one or more new subservices, and select Create Sub-Service.

The Service Modeler window opens to the same display as when you create a new service.

2. [Add CIs to the service](#) (see page 41).

3. Follow the rest of the applicable steps in How to Build a Service Model.

The subservice is created. When you reopen the parent service, the subservice appears connected to the top-level service CI with a default relationship of Has Member and propagation type of aggregate.

Follow these steps: (existing service as a subservice of another service)

1. Open the Operations Console, right-click the top-level service to which you want to add an existing service as a subservice, and select Tools, Edit Service.

The Service Modeler opens to the selected service's topology.

2. Select Monitored Services from the 'Browse by' drop-down list.

A list of existing services appears.

3. Select the service to include as a subservice and drag it into the Topology pane.

The service is added as a subservice connected to the top-level service CI (if you dragged to an empty spot in the Topology pane or directly onto the service CI) or to the CI on which you dragged the service with a default Has Member relationship and aggregate propagation.

Note: If you add a service CI defined within a domain manager as a subservice, CA SOI does not automatically include all CIs associated with that service in the domain manager unless that service has been fully [imported](#) (see page 75).

Create Groups

Groups are intermediate CI objects that collect CIs and relate them to each other by some role or function. Groups are often used to leverage the same custom propagation policy for a set of CIs, or to funnel several aggregate propagation types into one top-level CI. Group CIs do not commonly have a severity because they are not managed by a domain manager, but the impact of the CIs they contain is considered in how the group CI impacts other CIs and the service.

Follow these steps:

1. [Create a service](#) (see page 41) or [modify a service](#) (see page 67).

The Service Modeler opens.

2. Do one of the following to create a group:

- Select one or more CIs, right-click anywhere on the Topology pane, and select Group, Create.

The group is created as a parent of the selected CIs. The group establishes a default relationship of Has Member and propagation of aggregate with each connected CI.

Note: The main service cannot be included in a group.

- Verify that no CIs are selected, right-click a blank area in the Topology pane, and select Group, Create.

The group is created without relationships to any CIs, including the service CI.

- Right-click an unselected CI and select Group, Create.

The group is created as a child of the right-clicked CI. Use this option to create a group as a direct child of the service CI.

3. Double-click the group and give it a unique name.
4. (Optional) Drag CIs not already included in the service model onto the group CI.

The CIs are added to the group.

5. (Optional) Select one or more CIs that are already included in the service model, right-click the unselected group, and select Group, and one of the following actions:

Add

Creates Has Member relationships and aggregate propagation types from the group to all selected CIs. No existing relationships with other parent CIs are removed. Use this operation to maintain any existing associations the CIs may have in the service model while also adding them to the group.

Move

Creates Has Member relationships and aggregate propagation types from the group to all selected CIs. Existing relationships with other parent CIs are removed. Use this operation to move CIs from one place in the service to the group.

Set the Granularity Level

You can change the [granularity level](#) (see page 29) to determine if you need to include all resources in the model for their alerts to be shown as impacting the service.

Consider the following:

- You can only change the model granularity of a service you are currently editing.
- You can change the default service model granularity for subsequently created services on the Global Settings page of the Administration UI.

Follow these steps:

1. Right-click a service in the Service Modeler.
2. Select Change Granularity and one of the following.

Normal

A normal granularity service model does not aggregate alerts from child CIs not included in the service model.

Low


A low granularity service model aggregates alerts from child CIs not included in the model.

Note: You cannot include an unmodeled CI in any service if the CI is to impact a modeled CI in low granularity mode. This is because the alerts for that CI are attached to an existing model and the low granularity level looks for unmanaged alerts.

3. If a confirmation dialog opens, click Yes.

A low granularity service detects whether unmodeled child CIs exist that are related to modeled CIs and immediately begins propagating impact from those CIs. You can tell when a service is low granularity by the following indicators:


- An 'L' character appears in the Information column of the Services tab.
- Granularity displays as LOW in the service model tooltip on the Topology tab.
- Granularity displays as Low in the Information tab when you select the service.


When CA SOI detects that a low granularity service contains CIs with unmodeled child CIs and associated alerts that are propagating impact, a  icon appears next to the parent CI to indicate this fact. The color of the parent CI itself does not reflect the impact of the alerts; instead, the impact is propagated upwards on the CI's parent, and ultimately, the service.

Assign Propagation Types and Relationships

Each relationship type (and instance) has one specific propagation type. Selecting a relationship type implies a specific propagation type. When you create a service model, the default Has Member relationship (with [aggregate propagation](#) (see page 17)) is established between all services and CIs. If you want to change the relationship type, CA SOI lets you invoke the change relationship action from the toolbar or the shortcut menu.

Follow these steps:

1. Click the Relationship tool  on the Service Modeler toolbar.
A drop-down list appears with all propagation types available for selection.
2. Select a propagation type.
The drop-down list expands to display all relationships that are mapped to the propagation type that you selected.
3. Select a relationship.
The pointer in the Topology view changes to a wand.
4. Click the parent object in the propagation and relationship you want to change, and drag to the child object.
The Relationship Edit dialog opens if you are replacing an existing relationship.
5. Click OK on the Relationship Edit dialog if necessary.

Note: If you need to straighten the line between the two objects, click the *Straighten selected edges* tool .

The new relationship and propagation appear connecting the related items as follows:

- If you have selected the *Advanced* 'Chart display complexity level' option (default), a letter indicates the propagation type (A for aggregate, B for bound, C for custom, O for operative) followed by a colon (:) and the relationship significance value.
 - If you have selected the *Advanced with names* 'Chart display complexity level' option, the relationship name appears followed by a colon (:) and the relationship significance value.
 - For either option, the line connecting each set of related CIs is color-coded to indicate the propagation type (blue for aggregate, burgundy for bound, orange for custom, and pink for operative).
6. Continue changing to the same propagation and relationship, if necessary. For a different propagation and relationship, click the Relationship tool again and select from the drop-down list.
 7. Click another tool when you are finished.

The Relationship tool is deactivated and the previous relationships are replaced.

To change relationships and propagation using the shortcut menu

1. Click a child object in the Service Modeler to select it. To select multiple children of the same parent, hold down the Ctrl key and click them.

Note: Do not select the parent object.


The child objects are surrounded by small boxes.

2. Right-click the parent, select Establish Relationships, and select a propagation type and a mapped relationship.

The Relationship Edit dialog opens if you are replacing an existing relationship.

Note: Alternatively, you can also use the context menu of the relationship to change one or more relationship types. To do so, select and right-click the relationship, then select Change Relationship, *propagation type*, *mapped relationship* from the context menu.

3. Click OK on the Relationship Edit dialog if necessary.

Note: If you need to straighten the line between the two objects, click the *Straighten selected edges* tool .

The new relationship and propagation appear connecting the related items as follows:

- If you have selected the *Advanced* 'Chart display complexity level' option (default), a letter indicates the propagation type (A for aggregate, B for bound, C for custom, O for operative). The tooltip for this letter displays the full propagation and relationship type followed by a colon (:) and the relationship significance value.

- If you have select the *Advanced with names* 'Chart display complexity level' option, the relationship name appears. The tooltip for this name displays the full propagation and relationship type followed by a colon (:) and the relationship significance value.
- For either option, the line connecting each set of related CIs is color-coded to indicate the propagation type (blue for aggregate, burgundy for bound, orange for custom, and pink for operative).
- The relationship line tooltip displays the propagation type, relationship type, significance, and the names of the connected CIs.

Create a Propagation Policy with Operative Propagation

Operative propagation requires you to specify an impact threshold that causes only impact values greater than or equal to the threshold to propagate to the parent item. Operative propagation policy is useful in situations when a dependent CI is only affected by a child item when the child item's impact becomes severe.

By default, every relationship with operative propagation only propagates impact when the impact of the child CI is equal to or greater than 20. You can [change the default operative propagation policy](#) (see page 22) for each relationship that maps to operative propagation. The Service Modeler must be closed to change the default operative policy.

In the Service Modeler, you define operative policies that are specialized for individual relationships.

Follow these steps:

1. [Create a new service](#) (see page 41) or [modify an existing service](#) (see page 67).

The Service Modeler opens.

2. Do one of the following:

- Select and right-click a relationship that has operative propagation and select Operative Policies, Edit. You can select multiple relationships for the same source CI.

Note: You must select the relationship so that the relationship line widens for the Edit option to be available when you right-click.

- Select and right-click a CI that is the source of a relationship that has operative propagation and select Define Policies, Operative Policies, Edit.

The Operative Policy Editor dialog opens. The table displays whether the policy is new or previously defined, the affected source CI, target CI, relationship, and current threshold.

3. Click Set in the Threshold column, enter an impact threshold between 1-40, and press Enter.

The threshold changes.

4. (Optional) Do any of the following if necessary:

- Click Initial to return the Threshold to its initial setting when you opened the dialog.
- Click Default to return the Threshold to its default setting.
- Keep the 'When saving, remove default policy equivalents' check box selected to disregard a custom operative policy equivalent to the defined default policy for that relationship and use the default rather than creating the custom policy.

Note: CIs in maintenance mode are excluded from custom policy calculations related to average or percentage.

The threshold or setting changes accordingly.

5. Click OK.

The policy is associated with the specified relationship. Save the service to save any created policies.

Any custom operative policies that you define appear in the Policies tab in the pane below the service topology. You can right-click existing custom policies to edit or delete them. You can edit and delete multiple selected operative policies at the same time.

Create a Propagation Policy with Custom Propagation

Custom propagation requires you to specify when and how to change the severity of a parent item. The severity of the parent item changes based on the impact value of the children items.

Custom propagation policy is useful in situations when the impact of dependent CIs is complex. For example, a web server farm may have several servers that perform the same role. This redundancy means that a fault condition on one device may not impact the availability of the service. In fact, several servers may need to fail before users experience service degradation.

A sample rule is "If 50% of the servers are down, propagate an impact of Moderately Degraded. If 75% are down, propagate an impact of Severely Degraded." The default custom propagation policy contains the following rules:

- The first rule sets the parent item impact to Slightly Degraded when any of the included items have an impact greater than or equal to Moderately Degraded (20).
- The second rule sets the parent item impact to Moderately Degraded when any of the included items have an impact greater than or equal to Severely Degraded (30).

Note: When you assign custom propagation policy to a CI and the threshold is reached, the target CI attains the specified severity value and an infrastructure alert with the appropriate severity. However, because the alert applies to the CI as it relates to the parent service, the CI icon does not change color as expected, and the severity propagates to a separate impact calculation on the service, which is reflected in the service icon color.

You can [change the default custom propagation policy](#) (see page 22). The Service Modeler must be closed to change the default custom policy.

Note: CIs in maintenance mode are excluded from custom policy calculations related to average or percentage.

In the Service Modeler, you define custom policies that are specialized for individual relationships. A service can have multiple custom propagation policies, and a custom propagation policy between CIs can have up to four rules for propagating impact.

Follow these steps:

1. [Create a new service](#) (see page 41) or [modify an existing service](#) (see page 67).
The Service Modeler opens.
2. (Optional) [Create a group](#) (see page 45) if you have multiple configuration items that will have the same policy.
3. Right-click a parent object or group that has custom propagation and select Define Policies, Custom Policies, Edit.
The Policy Editor dialog opens.

4. (Optional) Select the Automatically Maintained check box to automatically change the relationship to Is Affected By or Is Evolution Of and apply the custom propagation policy when new CIs are added to a group that uses custom propagation. Auto-maintenance applies to relationships added by both manually editing a service and through automation by service discovery or service import.

Note: You can set Automatic Policy Maintenance for the Modeler operation to on by default in the Set Preferences dialog.

5. Select one of the following from the Policy Type drop-down list:

Average

Sets the impact of the parent item based on the average impact values of CIs associated with the policy.

Percentage

Sets the impact of the parent item based on a percentage of CIs that have the impact specified in the rule.

Any

Sets the impact of the parent item when any CIs associated with the policy have the impact specified in the rule.

All

Sets the impact of the parent item when all CIs have the impact specified in the rule.

Default: Any

The text of the rules on the dialog changes to reflect the selected policy type.

6. Complete the following fields and settings for each rule that you want to create:

% of items

(Percentage type only) Defines the percentage of items that must exceed the impact threshold to meet the rule criteria.

Threshold

Defines the impact threshold. The appropriate amount of CIs (as defined by the policy type) must meet or exceed the impact threshold to meet the rule criteria. Define a number between 0 and 40. The impact numbers translate to the following impacts:

- 0: Operational
- 1-10: Slightly Degraded
- 11-20: Moderately Degraded

- 21-30: Severely Degraded
- 31-40: Down

Set Severity

Defines the severity to assign to the parent item if the rule criteria are met.

The text for each rule changes to reflect the new settings. Define as many of the four available rules as the policy requires.

7. Select the entities to which to apply the policy. Use the arrows to move CIs from the Available Configuration Items pane.

All items included in the policy appear in the 'Configuration Items Included In Policy' pane.

8. (Optional) Click the 'Create a new policy' tab and create a different policy for any items that were not included in the current policy.
9. Click OK.

The policy or policies are associated with the specified items. You must save the service to save any created policies.

Any custom policies that you define appear in the Policies tab in the pane below the service topology. You can right-click existing custom policies to edit or delete them. You can edit one custom policy at a time and delete multiple custom policies at the same time.

Note: Only custom policies that you have created appear in the Policies tab, not the defaults.

Assign the CI Significance in the Service

[Significance](#) (see page 26) indicates the importance of a CI to related items in a service, and therefore influences the impact of a CI condition on those related items. It is a value from 1 through 10, where 1 is the least significant and 10 is the most significant. You can set global significance for CI types, and you can set significance for individual CIs and relationships in a service.

Consider the following:

- When you set global significance, it applies to new CIs that you create after the change. Existing CIs retain the previous significance for the type.
- Before you change the significance for a service, verify that all relationships are established. If you change a relationship later, the significance reverts to the default for the CI type.

- When you change significance for an individual CI in a service, the CI's significance value only changes within that service. If the CI belongs to other services, it retains its previous significance value in those services. This behavior helps ensure that you can maintain different significance values for the same CI if it is more or less important to other services.

Follow these steps: (global type significance)

1. Start the Operations Console, and select Tools, Default Significance.
The Global Significance Editor dialog opens, showing the default significance for each CI type.
2. Adjust the slider for the types whose significance you want to change, and click OK.
The global significance changes.

Follow these steps: (individual CI significance)

1. Open the Service Modeler.
2. Right-click a CI or relationship on the Topology tab, select Assign Significance, Set from the shortcut menu, and select a number from 1 to 10.
The previous significance is replaced. The relationship text in the Modeler displays the significance of the relationship between the connected CIs.

Create and Assign a Service-Level Agreement

A *service-level agreement* (SLA) is a contract that specifies the service expectations of internal or external customers. An example is the downtime that is acceptable for various resources.

An SLA can, for example, help ensure that an online shopping site is available and delivering the level of service that internet shoppers expect. If performance is poor, some transactions can be lost, thus reducing profits and discouraging repeat customers.

You can define SLAs when you define or edit service models using the SLA tab. CA SOI uses the measurable provisions that you specify in the SLA to monitor the real-time health of each associated service, and records outage time when the service is down. The recorded time is compared to the SLA thresholds to determine the status of the SLA for a given time period.

For more information about defining, managing, and visualizing SLAs, see [Creating and Working with Service-Level Agreements](#) (see page 111).

Create and Assign an Escalation Policy

After you define service models and their associated CIs, relationships, and propagation policies, you can add escalation policy. *Escalation policy* specifies the automated actions to take in response to fault conditions.

Some of the conditions that require an automated response are as follows:

- Alert impact on the service
- Time since acknowledgment
- Alert attribute values
- Time spent in an alert queue

The following are common actions:

- Sending an email message to an operator or service owner
- Opening a help desk ticket
- Invoking a script or application to help diagnose and remedy the fault condition

The following kinds of escalation policies are available:

Nonglobal

Escalates alerts in one or more specified services or alert queues that meet the policy criteria. For example, you can create nonglobal escalation policy for a payroll service owner who requires a notification when an alert is raised against the payroll service. When the policy considers an alert for escalation, it evaluates the alert against the nonglobal policy first, before any global policy.

Global

Escalates all alerts that meet the policy criteria. For example, you can create global escalation policy for an IT manager who requires notification when *any* service alert is raised in CA SOI.

You can define escalation policy in the Operations Console or when creating or editing a service in the Service Modeler using the Alert Escalation tab.

For more information about defining escalation policy and escalation policy actions, see the *Event and Alert Management Best Practices Guide*.

Validate and Save the Service

You save a service to apply any changes that you make in a Modeler session. The Modeler validates the contents of the service automatically each time you save. You can also validate a service at any time to verify that any operation does not invalidate the service model.

Follow these steps:

1. (Optional) Click Perform Service Validation  at any time.

Note: This operation does not save the service. When you do save, validation is performed automatically.

One of the following occurs:

- The Validation Results dialog opens with a message stating that no problems were detected in the service topology model.

Note: You can set a preference not to display this dialog when the validation is successful. For more information about how to set a preference, see the *Administration Guide*.

- The Service Topology Validation dialog opens with a list of problems in the service model construction that would cause service impact calculation errors, such as a state propagation loop. For more information about the types of potential errors, see [How Service Validation Works](#) (see page 57).

Note: The topology view highlights the objects (that constitute a validation problem) when a specific problem is selected in the list that displays all the validation problems in the service model.

2. [Fix service validation errors](#) (see page 57) if necessary.
3. Click Save or OK in the Service Modeler when you are ready to save the service.

One of the following occurs:

- The Validation Results dialog opens stating that problems with the topology have prevented the save operation. When you click OK on this dialog, the Service Topology Validation dialog opens and displays validation errors. You must [fix service validation errors](#) (see page 57) before you can save the service.


- The Save Service dialog opens stating that the service changes were saved successfully.

Note: You can set a preference to hide this dialog when the save is successful. For more information about how to set a preference, see the *Administration Guide*.

- A dialog opens stating that the service contains noncritical errors. You can select to save the service, but it is placed in the offline/test mode until you fix the errors.

How Service Validation Works

Each step in the service validation process tests one aspect of the service topology. The tests target the structure of the service to verify that a service model is valid and will not cause impact propagation calculation errors.

The Service Modeler performs validation automatically before each save attempt. You can also validate on demand by clicking Perform Service Validation .

Each validation test is assigned a rank, which determines its impact on the subsequent service save process. The ranks and their effect are as follows:

UNKNOWN

Indicates that the test cannot determine the effect on the service.

LOW

Indicates that the model is inefficient, is missing items, or could use cosmetic improvements.

MEDIUM

Indicates that the model structure may affect the operation of the state engine.

HIGH

Indicates that the service cannot be put online in its current structure.

CRITICAL

Indicates that the service cannot be saved.

The tests run in the following order:

1. **CRITICAL:** The topology graph must be connected. The topology must not contain unconnected subgraphs; it must be possible to move between any nodes of the topology using associations, or all CIs must have relationships with other CIs.
2. **HIGH:** The topology graph must not have multiple roots. The state or impact of all CIs should propagate to one node or CI.
3. **CRITICAL:** The root node or CI must not have dependents. The root node is the final target for state or impact propagation.
4. **HIGH:** The topology graph (or subgraphs) must not have cycles formed by relationships. State propagation through relationships must not form loops.
5. **CRITICAL:** Component services of the tested service must not have antecedent CIs. Specifically, no antecedents are prescribed by the tested service's topology; the component services may be complex.
6. **LOW:** All graph nodes must have labels.

Results display in the Service Validation Results dialog if errors occurred. The table in this dialog has the following columns:

Rank

Defines the error rank.

Type

Defines the error type. Types include EDGE and PATH.

Root

Defines whether the problem affects the root service CI.

Object

Defines the affected parent object.

Description

Describes the nature of the problem.

Click each result for the Modeler to highlight the objects and relationship involved in the error. Use this visual information and the Description to resolve problems before saving the service model.

Chapter 5: Customizing Service Model Display

This section describes how you can customize your service model display.

This section contains the following topics:

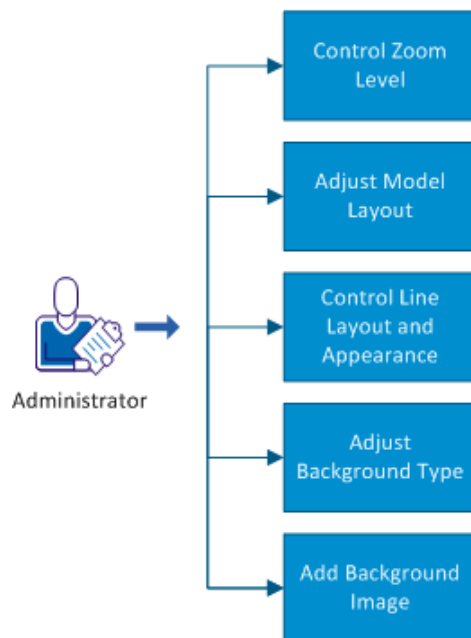
[How to Customize Service Model Display](#) (see page 59)

How to Customize Service Model Display

As an administrator, you can customize the zoom level, model layout, background type and image in the Operations Console to optimize the service model layout.

Use this scenario to guide you through the process:

How to Customize Service Display



- [Control Zoom Level](#) (see page 60)
- [Adjust Model Layout](#) (see page 61)
- [Control Line Layout and Appearance](#) (see page 62)
- [Adjust Background Type](#) (see page 64)
- [Add a Background Image](#) (see page 64)

Note: For a summary of all controls available from the Service Modeler and the Topology view of the Operations Console, see the *Administration Guide*.

Control Zoom Level

Several tools exist in the Service Modeler to control the zoom level of the Topology view. You may require different zoom levels in large service models when you need to edit specific areas or view the entire model in one display.

To control zoom level, click one of the following:



(Interactive Zoom Tool)

Enlarges or reduces the topology when you click the tool and drag it on the screen.



(Marquee Zoom Tool)

Increases the magnification in a specific region when you click the tool and select a region in the topology.



(Zoom Level Control)


Specifies the amount of magnification. Select Auto fit to automatically fit the entire model within the display.

You can also do the following when any tool is selected:

- Hold the Ctrl key and scroll the mouse wheel to quickly zoom in or out.
- Hold Ctrl + Shift and scroll the mouse wheel to more finely adjust the zoom level and resolution.

Adjust Model Layout

Several tools exist to optimize the layout of the service model and how it is organized. Different types of service models require different layouts to best visualize the relationships and hierarchy. You can adjust the layout type to best fit the structure of your service model, and you can adjust the layout mode to specify the level of automated adjustment to perform to the model layout (according to the selected layout type) when you add items.

To adjust the overall service model layout type, click the Apply Automatic Layout  icon and select one of the following:

Circular

Emphasizes clusters that are present in a network's topology.

Grid

Arranges objects in horizontal rows and vertical columns.

Hierarchical

Emphasizes relationships among objects by placing them at different levels, like an organizational chart at a company. This is the default when you build services from scratch.

Orthogonal

Minimizes bend points by arranging objects horizontally and vertically, at 90 degree angles.

Symmetric

Emphasizes symmetries that are present in a service's topology. This is the default for newly discovered services. It is also the fastest and yields the smallest topologies.

The layout adjusts to the selected type.

To adjust the layout mode, right-click an empty area in the Topology view, select Layout Mode, and select one of the following:

Note: Any layout mode adjustment always adheres to the selected layout type.

Automatic

Performs an automatic layout adjustment on the whole model chart after every new item that you add to a model.

Incremental

Performs an automatic layout adjustment after every new item that you add to a model while minimizing the amount of movement.

Manual

Performs no automated layout adjustment. When you drag and drop an item into the model, the item is placed at the drop point, and relationships establish accordingly. When you add an item using the right-click menu or drop an item directly on a parent item, the Modeler uses a fixed location relative to the target parent item.

Control Line Layout and Appearance

You can control how the relationship lines appear that connect resources in a service model. Bent lines are the default, and can improve the presentation. Straight lines can improve the mobility of the nodes in the Topology. You can specify the line type only if the [layout mode](#) (see page 61) is Automatic or Incremental.

Note: Bent and straight lines do not apply to the Topology tab in the Operations Console.

To select bent or straight lines in the service topology, do one of the following:

- Right-click an empty area in the Topology view, select Layout Mode, and select or clear Bend Edges.

The line type for the model is set.

- Select a bent relationship line and click the Straighten Selected Edges tool .

The previously bent relationship line straightens.

- Select View Preferences in the Operations Console, expand Modeler, Edit Mode, and Layout in the Set Preferences dialog, select Edge Bends, and select Yes or No in the Edge Bends drop-down list.

The default line type is set.

- Right-click an empty area in the Topology view and select Route Edges.

Edge bend points are modified so that the lines run orthogonally. Item positions do not change.

- To manipulate bend points, select the bend point (not the relationship line) and drag to modify the bend.

You can also control what information displays for relationships by doing any of the following:

- Click the Chart Complexity Level  tool and select one of the following:

Advanced


Displays a letter in each relationship line that indicates the propagation type.

Advanced with names

Displays the full relationship name in the relationship line.

For both options, the lines are color-coded to indicate the propagation type:

- Aggregate: Blue
- Bound: Burgundy
- Custom: Orange
- Operative: Magenta

- Click the Change relationship visibility tool  and select one of the following:

Specific Propagation Type

Displays only the relationships of the selected propagation type. This may be useful to pinpoint all relationships of a certain propagation, such as custom, to fine-tune their policy.

Show All

Displays all relationship lines.

Hide All

Hides all relationship lines.

Adjust Background Type

You can adjust the background type and pattern of the Topology view.

To adjust the topology background type, right-click an empty area in the topology and select Chart Grid and one of the following:

Type

Defines the pattern type of the topology background. Select from the following options:

None

Displays and maintains no pattern in the topology background.

Line

Displays lines that define the grid areas in the topology. This option presents grid areas as squares with divisions within the squares controlled by the Size setting.

Point

Displays individual points that define the grid areas in the topology. This option presents grid areas as squares, but does not connect the areas with lines. You control the distance between the points with the Size setting.

Invisible

Maintains chart areas that are not visible in the Modeler.

Size

Defines the number of units that populate each area of the grid. This option only has obvious results for the Line and Point grid types.

Add a Background Image

You can add a background image to a service model to enhance the model presentation. For example, if your service can be modeled based on the location of CIs in a specific geographical area, you can add a map of the area as a background image and model the service on the map.

Follow these steps:

1. Copy the image to `SOI_HOME\jsw\bin\sam_topo_images` on the SA Manager system.
2. Open the Operations Console and open the Service Modeler to create or edit a service in which to place the image.

3. Right-click the background of the service topology on the Service Modeler and select Background Image, Select.

The Image File Selection dialog opens.

4. Select the background image from the directory in Step 1 and click Select.

The image appears in the background of the service topology.

5. Click OK.

The changes are saved, and the service's topology includes the background image when you view it on the Operations Console.

Chapter 6: Editing and Managing Services

As an administrator, you can modify and manage your service models. After you create a service, you can edit the model in the Service Modeler or manage or add properties to the service in the Operations Console. You can also associate your services to customers and manage them from the perspective of customers.

This section contains the following topics:

[Modify a Service Model](#) (see page 67)

[Perform Copy, Cut, Paste, Delete, or Remove Operations on a Service in the Console](#) (see page 68)

[View Service Information](#) (see page 70)

[Set Service Priority](#) (see page 72)

[Set the Service Location for Google Earth](#) (see page 72)

[Control User Group Service Access](#) (see page 73)

[Change Multiple CI Relationships Simultaneously](#) (see page 74)

Modify a Service Model

You can modify any created or imported service at any time.

Follow these steps:

1. Open the Operations Console, right-click a service in the Navigation pane, and select Edit Service.

The Service Modeler window opens to the same display as when you create a new service.

2. Make changes to the service model using the applicable steps in How to Build a Service Model.
3. Click Save.

Perform Copy, Cut, Paste, Delete, or Remove Operations on a Service in the Console

You can perform the following operations on a service or subservice in the Navigation pane to adjust the service hierarchy:

Cut

Removes a subservice from the Navigation pane and places it in the paste buffer.

Remove

Detaches a subservice from its parent service. The subservice moves up one level in the tree.

Copy Service

Places a service in the paste buffer and leaves it in the Navigation pane.

Note: You can copy monitored services only.

Paste

Adds a cut or copied service as a subservice to another service.

Note: If you copy and paste a top-level service, the original service is removed from its original top level and pasted as a subservice. However, if you paste a copied subservice, the original subservice appears in its original location and the new location.

Delete

Removes a service from CA SOI permanently. Any subservices are not deleted, but are promoted to the next highest level in the tree in the Navigation pane.

Note: You can also perform all service hierarchy operations from the List tab of the Contents pane for service CIs.

To cut and paste a service

1. Open the Operations Console and right-click a subservice in the Navigation pane, and select Cut.

A confirmation dialog opens.

2. Click Yes.

The subservice is placed in the paste buffer and removed from its parent service. If the service has no other parent, it appears in the top-level Services node. When you paste it in another service, it disappears from the top-level Services folder.

3. Right-click another service under which you want to relocate the cut service, and select Paste.

A Paste dialog opens if the subservice has relationship properties that you customized when it was a part of the previous parent service.

4. Click yes to preserve the previous relationships or no to erase the previous relationships.

The service becomes a subservice of the selected parent service.

To remove a subservice

1. Open the Operations Console, right-click a sub-service in the Navigation pane, and select Remove.

A confirmation dialog opens.

2. Click Yes.

The subservice is detached from the parent service. If a removed subservice does not have other parents, it moves under the top-level Services node in the tree.

To copy and paste a service

1. Open the Operations Console, right-click a service in the Navigation pane, and select Copy.

The service is placed in the paste buffer and remains in the Navigation pane.

Note: The contents of the buffer are erased if a Paste operation does not immediately follow.

2. Right-click another service under which you want the copied service, and select Paste.

The service becomes a subservice of the selected parent service.

To delete a service

1. Open the Operations Console, right-click a service in the Navigation pane, and select Delete.

A confirmation dialog opens.

2. Click Yes.

The service is deleted from CA SOI.

View Service Information

You can view details about a service and manipulate several aspects of service information.

Follow these steps:

1. Start the Operations Console and click the Services tab in the Navigation pane.
A list of services displays. The columns to the right of the service name display the number of alerts of each severity in the service, total number of alerts associated with the service, and the operational mode of the service.
2. Select a service, click the Services tab in the Contents pane, and click the Information tab in the Component Detail pane.
3. Review the following service information on the tab:

Note: Depending on your user group privileges, you may not be able to edit all modifiable properties in this tab.

Service Name

Displays the name of the service at the top of the tab. Click Set to change the service name.

General Information

Displays a list of the following basic service information:

Health

Displays the service health state.

Quality Impact

Displays the quality impact value, which is the inverse of service health.

Risk

Displays the risk of infrastructure alerts on service quality.

Granularity

Displays the current [granularity level](#) (see page 46) as either Normal or Low.

Operational Mode

Displays the current operational mode, either Production or Maintenance. Click Set to manually change the operational mode.

For more information about service maintenance, see the *Administration Guide*.

Priority

Displays the service priority value. Click Set to set the priority.

For more information, see [Set Service Priority](#) (see page 72).

Location

Displays the service location. Click Set to enter or change the location. Google Earth uses this value to display CA SOI services according to physical location.

For more information, see [Set Service Location for Google Earth](#) (see page 72).

Description

Displays the service description. Click Set to enter a new description.

Service Level Agreements

Displays the SLA associated with the service with information such as name, current status, and description.

Maintenance Schedules

Displays maintenance schedules associated with the service. From this table, you can edit an existing maintenance schedule.

Security

Displays the user groups that can access the service. From this table, you can add or remove user group access.

Connectors

Displays the connectors managing the service, if the service was imported.

Associated Escalation Policies

Displays the escalation policies assigned to the service.

More Information:

[Set the Service Location for Google Earth](#) (see page 72)

Set Service Priority

Priority indicates the importance of a service to the business. Priority applies to the Dashboard, where you can sort services by this value. Priority is not used in any impact calculations, and all services have no priority value by default.

Follow these steps:

1. Open the Operations Console and select a service.
The alerts for that service appear in the Contents pane, and general service information appears in the Information tab of the Component Detail pane.
2. Click *set* next to Priority in the General Information area of the Information tab, and select one of the following priorities from the drop-down list:
 - None (0)
 - Low (7)
 - Medium (8)
 - High (9)
 - Critical (10)

Set the Service Location for Google Earth

You can associate a service with a location and then use Google Earth to view it based on the specified location.

If you enter a city, Google Earth shows the service in the center of the city. If you enter a street address (for example, 710 Ashbury, San Francisco, CA) Google Earth shows the service at that location if Google Maps can validate the address. You can include additional information (for example, Building 12 or Floor 3) as long as Google Maps can resolve it as a valid address.

You can view Google Earth with mapped CA SOI services from the Dashboard.

Follow these steps:

1. Open the Operations Console, and click a service in the Navigation pane on the left side.
The alerts for that service appear in the Contents pane, and general service information appears in the Information tab of the Component Detail pane.


2. Click *set* next to Location in the General Information section of the Information tab.
A text box opens.
3. Enter a location and press Enter. The location can be as broad as a country, or as specific as street address with city and state.
The location value appears on the Information tab.

Control User Group Service Access

You can grant or remove user group access to individual services from the Operations Console.


The default user groups have access to all services unless you remove all access from the Users tab. You can only change individual service access for groups that do not have access to all services.

To add user group service access

1. Select a service in the Operations Console.
The alerts for that service appear in the Contents pane, and general service information appears in the Information tab of the Component Detail pane.
2. Scroll to the Security section.
This section displays the user groups that can currently access the service.
3. Click Add .
The Select Users/Groups dialog opens and displays all user groups that do not currently have access to the service.
4. Select a user group and click OK.
The user group appears in the Security table, and its users can now view the service in the Operations Console.

To remove user group service access

1. Select a service in the Operations Console.
The alerts for that service appear in the Contents pane, and general service information appears in the Information tab of the Component Detail pane.
2. Scroll to the Security section.
This section displays the user groups that can currently access the service.

3. Select a user group from the table and click Remove .

The Remove dialog opens.

4. Click Yes to confirm the removal.

Note: A dialog opens if a user group has access to all services by default. You must change this setting in the Users tab before you remove access for a specific service.

The user group disappears from the table and can no longer see the service in the Operations Console.

Change Multiple CI Relationships Simultaneously

You can select multiple CI relationships and change them all simultaneously in the Service Modeler. This way, you can apply the same relationship type to multiple CIs simultaneously; you do not need to individually change each relationship.

Follow these steps:

1. Open the Operations Console and click the Services tab.

A list of services displays.

2. Select the service for which you want to change multiple CI relationships.

3. Right-click the service and select Edit Service from the context menu.

The Service Modeler window opens and displays the service topology in the Topology pane.

4. Select all CIs (excluding the parent CI) for which you want to change the relationship.

5. Right-click the parent CI and select Establish Relationships, *<relationship type>* from the context menu.

The Relationship Edit dialog opens.

6. Click Yes.

A confirmation message appears.

7. Click OK.

The message window closes.

8. Click Save, then click OK in the Service Modeler.

All relationships for the selected CIs are changed and the Service Modeler closes.

Chapter 7: Importing Services

This section describes how you can import services from domain managers.

This section contains the following topics:

[How to Import Services](#) (see page 75)

How to Import Services

As an administrator, you can use CA SOI to import services from one or more domain managers that CA SOI monitors. For example, CA Spectrum or CA CMDB services can become CA SOI services. Importing services lets you quickly populate CA SOI with service models that leverage existing technology. After the import, you can update the services to add more properties or CIs, or you can leave them as is.

You can import services in two ways:

- [Import Services Automatically](#) (see page 75)
- [Import Services Manually](#) (see page 76)

If you remove a service from a domain manager, it is not automatically deleted from CA SOI. If you no longer want to manage the service in CA SOI, delete the service manually in CA SOI. If auto-import is on, services that are deleted from CA SOI but not from the domain manager import again. You can also manually reimport any service that you delete from CA SOI that still exists in the domain manager.

Import Services Automatically

Automatic import type imports all services from the connector when CA SOI starts. Automatic import is useful when domain managers update regularly. However, this type of import is not advisable when domain managers are stable, because automatic import uses a considerable amount of resources.

Follow these steps:

1. Open the Operations Console, and select Tools, Import Services.

The Configure Data Sources dialog opens and displays currently running connectors.

Note: Each connector entry contains a five-digit ID number defined by the USM schema. For a list of connectors and their ID numbers, see the *Connector Guide*.

2. Select the connectors whose services you want to load automatically. They would have *No* in the Auto Import column.

3. Click Auto on the toolbar.
No changes to Yes in the Auto Import column.
4. Click Save or OK.
All services from the connector are automatically imported, and a new import occurs each time CA SOI starts.

Import Services Manually

Manual import type imports specific services that you select. Manual import is recommended if new services were added to the source application during the current session. Manual is the default import type.

Follow these steps:

1. Open the Operations Console, and select Tools, Import Services.
The Configure Data Sources dialog opens and displays currently running connectors.
Note: Each connector entry contains a five-digit ID number defined by the USM schema. For a list of connectors and their ID numbers, see the *Connector Guide*.
2. Select the connector whose services you want to load manually and click Import.
Note: You cannot import services from connectors that have *Offline* in the Connector Status column.
The Import Services dialog opens and lists available services in the selected connector. Services that have already been imported contain a check mark in the Exist in SOI column.
3. (Optional) Enter text in the Filter field to limit the number of services displayed.
The number of listed services decreases.
4. Complete one of the following actions:
 - Click Add All Services to move all services.
 - Select the services you want to import, and click the right arrow.The services move to the right pane of the dialog.
5. Click OK.
The Import Services dialog closes. The Import Status column of the Configure Data Sources dialog displays *Waiting* while the console connects to the source. The same column displays *Importing* while the import is in progress. When the import finishes, the column value changes to *Idle*, and the services appear in the Navigation pane.

Processing Updates for an Imported Service

After you import a service model, connectors send updates as they occur in the domain manager. The following actions occur:

- CIs are added to a service when they are added to the corresponding domain manager service.
- CIs are removed from a service when all relationships that connect CIs to other CIs in the service are removed.

Launch the Source of Imported Services

You can launch the source domain managers of services imported to CA SOI.

Follow these steps:

1. Open the Operations Console, and right-click a service in the Navigation Pane.
A shortcut menu opens.
2. Select Launch, and select the domain manager to open.
The source domain manager interface starts.

Chapter 8: Using Service Discovery

This section describes how to create and manage service models using Service Discovery policies.

This section contains the following topics:

[Service Discovery](#) (see page 79)

[How to Create Dynamic Service Policies](#) (see page 80)

[How to Create Automatic Relationship Policies](#) (see page 85)

[How to Create Unmanaged Relationship Policies](#) (see page 92)

[Topology Warnings](#) (see page 99)

[How to Manage Service Discovery Policies](#) (see page 100)

[Service Discovery Connector Configuration](#) (see page 105)

[How to Use Command Line Service Discovery Operations](#) (see page 107)

Service Discovery

Service Discovery is a CA SOI connector that publishes its data to CA SOI as an ordinary connector running in CA SOI Integration Services. You can use a wizard to define policies that automatically create and maintain services and relationships between source and target CIs according to specified criteria. Service Discovery searches the Persistent Store of all the data that is collected from connectors for CIs that match the policy criteria you define. Service Discovery then creates services and relationships using those CIs.

Service Discovery runs as a part of Integration Services as a connector. The Service Discovery connector appears on the Administration Pages panel among other connectors and the connector can be started or stopped from here.

You can install the Service Discovery connector on the same computer as CA SOI Integration Services and the SA Manager or on a different computer. It is possible to have only one Service Discovery connector in the whole CA SOI solution.

Service Discovery is useful when:

- You want to maintain a collection of CIs that share common qualities.
- Some areas of your enterprise are constantly changing, and you want to maintain a service dynamically without having constantly to add CIs manually.
- Certain relationships occur repeatedly in your enterprise, and manually establishing those relationships would take a significant amount of time.

The following policy types are available:

Note: Service Discovery policies support USM-Core CIs.

[Dynamic service policy](#) (see page 80)

Automatically discovers CIs matching specified USM types and attribute criteria and places them under a specified service.

[Relationship policy](#) (see page 85)

Automatically creates relationships that are based on source and target CI types and match criteria.

[Unmanaged relationship policy](#) (see page 92)

Discovers unmanaged relationships and creates managed relationships (with the same source and target CIs) based on the specified filtering criteria match and whether the source CI is part of the service.

Note: A subservice does not automatically inherit the scope of the parent service. If you want to scope an Automatic Relationship or Unmanaged Relationship policy to a parent service and any of the parent's child services, you must explicitly add the child services to the scope list in the policy.

How to Create Dynamic Service Policies

As an administrator, you can create dynamic service policies that automatically include resources in a new dynamic service based on the criteria that you define. You can use regular expressions and test your expressions with the Regex Tester. The Service Discovery Policy Editor lets you create and manage Service Discovery policies based on the following:

Relationship

Defines the relationship to establish between CIs that meet the policy criteria and the service CI.

Type

Defines the USM type on which to apply the policy.

For example, you can create a Service Discovery policy to match all ComputerSystem CIs within a certain IP address range and add those CIs to a service. When you open the Service Discovery Policy Editor, you can create a policy for one relationship and type. To define multiple policies for one dynamic service, you must go through the wizard multiple times.

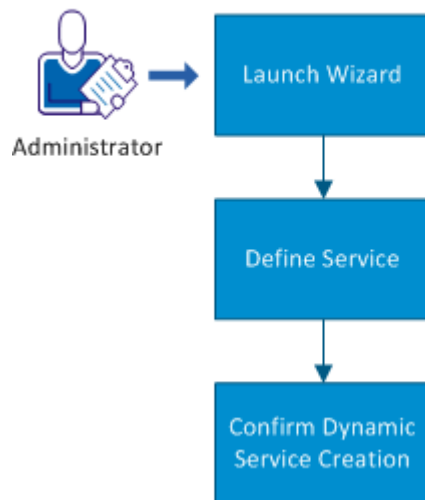
The Service Discovery engine evaluates the Persistent Store for CIs that match the Service CI relationship and attribute criteria. Every CI that matches is then added to the service. The evaluation is continuous, so the service is updated dynamically when matching changes in the Persistent Store occur.

You can create one policy at a time for a relationship. Multiple policies for each relationship are supported, but you must create them in separate operations. Multiple dynamic service policies can contribute to the same service as long as they share the name of the service.

Note: Service Discovery policies support only USM-Core CIs.

Use this scenario to guide you through the process:

How to Create Dynamic Service Policies



1. [Launch the Wizard](#) (see page 82).
2. [Define the Service](#) (see page 82).
3. [Confirm the Dynamic Service Creation](#) (see page 85).

You can also read the [scenario](#) (see page 153) that provides an example for creating dynamic service policies.

Launch the Wizard

You use the Service Discovery Policy Editor to create and manage dynamic services and relationship criteria.

Note: Page fields marked with an asterisk (*) are required.

Follow these steps:

1. Open the Operations Console.
2. Select Tools, Service Discovery Policies.

The Service Discovery Policy Editor opens.

Note: System-wide, only one user can access the Service Discovery Policy Editor at a time.

The Policies tab provides a tree that displays the Dynamic Services, Automatic Relationships, and Unmanaged Relationships available.

3. Right-click Dynamic Services in the Policies tab and select Create.

Define the Service

The wizard guides you through the process of service creation.

Follow these steps:

1. Specify a name for the service in the Service Name field.
2. The Define Service page lets you name the new dynamic service and assign user groups to the service. Assigning user groups for the service lets Service Discovery automatically assign the same user access to the services when it (Service Discovery) creates them. This way, you enhance the security of the service, as only authorized users are able to access the service.

Click Next.

3. On the Relationship Criteria page, select a USM relationship type from the Relationship Type drop-down list. This relationship type defines the relationship between the service CI and CIs that meet the policy criteria.

The Target pane lets you build criteria using Boolean expressions.

4. Select a USM type from the Class drop-down list in the Target pane.

Consider the following:

- The Attribute drop-down list displays a subset of the USM properties available for the selected type. Certain attributes such as MdrProduct attributes and date/time data types are not supported and do not appear in the drop-down.

- When you select a USM type, all available derived types in the USM hierarchy are also included. For example, when you select the ComputerSystem type, its child type VirtualSystem is automatically included in the policy.
 - For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.
5. (Optional) Select a USM attribute, comparison type, and attribute value for the expression. If you want to add all CIs of the selected type without adding attribute criteria, skip to Step 6.

Note: You can use [regular expressions](#) (see page 83) by selecting either Matches regex or Does not match regex from the Comparison Type drop-down list. Regular expressions are not available for all attributes. Click Test Regex to open the [Regex Tester](#) (see page 84) and test the regular expression against a string. For information about using regular expressions in other CA SOI features, see the *Event and Alert Management Best Practices Guide*.

Important! If you do not add attribute criteria, you risk discovering a large number of CIs, which can compromise system performance.

Note: For attribute values, mouse over the field and a tooltip displays the data type required.

6. Click Add.

The criteria are added to the expression, which appears in the logic tree and as an expression in the field at the bottom of the page.

7. (Optional) Repeat Steps 3 - 6 to add additional criteria and logic.

Note: For more information about using advanced logic with multiple target criteria, click the Hints link above the logic tree.

8. Click Next.

The Confirm page opens. A Topology Warning dialog may appear if the Service Discovery editor detects that the policy might cause topology errors or unmanageable services. Evaluate the policy before proceeding to verify whether changes are required. For more information, see [Topology Warnings](#) (see page 99).

Regular Expressions Considerations

Consider the following situations when using Regular Expressions in CA SOI:

Regular Expressions in CA SOI act as a *find*, not as a *match*.

A *find* searches for the pattern across the strings, including the substrings.

A *match* searches for the pattern in the strings only.

Example:

With the following strings: "cart" and "artistic":

- A *find* for the "art" pattern finds "art" in the substrings of both the "cart" and "artistic" strings.
- A *match* for the "art" pattern does not match in "cart" or "artistic" strings because a *match* does not search the substrings.

To perform a *match* in CA SOI, enclose the pattern with "^" and "\$", such as "^art\$". You can use the Regex Tester to verify your expressions before implementing them.

Use the Regular Expression (Regex) Tester

You can use the Regular Expression (Regex) Tester to validate a regular expression before using the expression in CA SOI.

Follow these steps:

1. In a dialog that supports regex, click Test Regex.

The Regex Tester dialog appears.

The Operation field describes the conditions:

- The case-sensitivity
- If the pattern is to match or not match

If you entered an expression in the Attribute Value field, the expression appears in the Regex Pattern for editing.

2. Enter (or edit) the regular expression in the Regex Pattern field.

The Valid? field indicates if the expression you entered is a valid regular expression. The field displays Yes or No.

3. Enter a test string in the Test Text field.

The Found?/Not Found? field indicates if the regular expression finds or does not find the Test Text string, based on the Operation conditions. The field displays True or False.

4. Perform one of the following actions:

- Click Use Pattern to close the RegEx Tester dialog and transfer the Regex Pattern to the Attribute Value field.
- Click Cancel to close the RegEx Tester dialog and leave the Attribute Value field unchanged.

Confirm the Dynamic Service Creation

The Confirm page displays the policy expression for the dynamic service policy criteria you created, displays warnings about policies that could create topology errors or large services, and lets you confirm creation of the new dynamic service policy. The page also displays the user groups that you assign to the service.

Follow these steps:

1. Click Finish to confirm creation of the new dynamic service policy.

The service name appears under the Dynamic Services folder. Expand the service name to view the relationship and type defined in each policy for the service. You can define additional policies for the service through the right-click menu using the same or new relationships.

Note: If you receive a topology warning, you should evaluate to determine whether changes are required. For more information, see [Topology Warnings](#) (see page 99).

2. Click Save or OK when the dynamic service policies are complete.

The Service Discovery engine begins scanning the Persistent Store for CIs that match the policy criteria.

How to Create Automatic Relationship Policies

As an administrator, you use the Service Discovery Policy Editor to create and manage policies that automatically create relationships based on policy criteria.

Automatic relationship policies are based on the following:

Relationship

Defines the type of relationship to create.

Source CI

Defines the CI type to use as the source of the relationship with optional attribute-based criteria.

Target CI

Defines the CI type to use as the target of the relationship with optional attribute-based criteria.

Match Criteria

Defines how the source and target CIs must relate (based on attribute values) for the policy criteria to match and create a relationship between the CIs.

Scope

Defines under which services the relationships are created.

The automatic relationship policy creates a new relationship in the service that automatically adds the target CI to the service. Only the source CI needs to pre-exist in the service.

A CI can become part of the service through various ways; for example:

- Imported service already includes the CI in it.
- CI is manually added to the service.
- Another Service Discovery policy (for example, a dynamic service policy) adds the CI to services.

Note: When a policy specifies a class name (for example, class=ComputerSystem), the policy applies to the CIs of that class and also to its subclasses (as defined by the USM Schema Type hierarchy). For example, in the case of ComputerSystem, the policy would also apply to VirtualSystems.

When the specified source CI appears in a service and the policy criteria are met, Service Discovery adds a relationship to all the specified target CI in the service. You can scope the policies to a subset of service models or apply them to all service models in which the source CI exists.

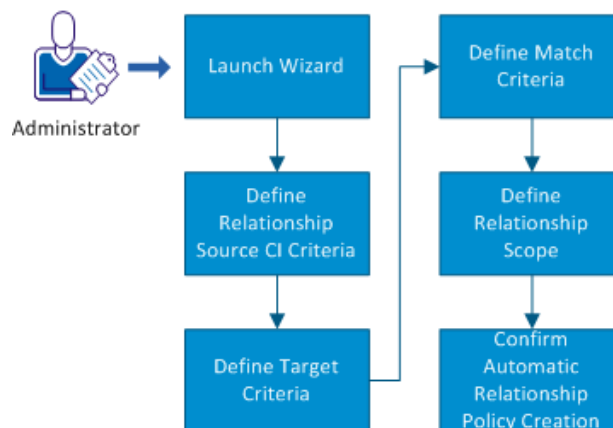
For example, you can create an automatic relationship policy that creates relationships between ComputerSystem CIs and all RunningSoftware contained on the ComputerSystem based on matching ComputerSystem's and RunningSoftware's IP addresses.

You can create one policy at a time for a relationship. Multiple policies for each relationship are supported, but you must create them in separate operations.

Note: Service Discovery policies support only USM-Core CIs.

Use this scenario to guide you through the process:

How to Create Automatic Relationship Policies



1. [Launch the Wizard](#) (see page 87).
2. [Define the Relationship and Source CI Criteria](#) (see page 88).
3. [Define the Target CI Criteria](#) (see page 89).
4. [Define the Match Criteria](#) (see page 90).
5. [Define the Relationship Scope](#) (see page 91).
6. [Confirm the Automatic Relationship Policies Creation](#) (see page 91).

You can also read a [scenario](#) (see page 157) that provides an example for creating automatic relationship policies.

Launch the Wizard

You use the Service Discovery Policy Editor to create and manage automatic relationship policies.

Follow these steps:

1. Open the Operations Console.
2. Select Tools, Service Discovery Policies.

The Service Discovery Policy Editor opens.

Note: System-wide, only one user can access the Service Discovery Policy Editor at a time.

The Policies tab provides a tree that displays the Dynamic Services, Automatic Relationships, and Unmanaged Relationships available.

3. Right-click Automatic Relationships in the Policies tab and select Create.
The wizard opens on the Source Criteria page.

Define the Relationship and Source CI Criteria

The Source Criteria page lets you define the relationship type and the source CI criteria.

Follow these steps:

Note: Page fields marked with an asterisk (*) are required.

1. Select a USM relationship type from the Relationship Type drop-down list.

Once you select a relationship type, the Relationship Type field becomes read-only in subsequent pages and the relationship expression builds as you select the target and match expression.

2. Select the class (USM type) of the source CI from the Class drop-down list in the Source pane.

Consider the following:

- The Attribute drop-down list displays a subset of the USM properties available for the selected type. Certain attributes such as MdrProduct attributes and date/time data types are not supported and do not appear in the drop-down.
- When you select a USM type, all available derived types in the USM hierarchy are also included. For example, when you select the ComputerSystem type, its child type VirtualSystem is automatically included in the policy.
- For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.

3. (Optional) Select a USM attribute, comparison type, and attribute value for the expression. If you want to add all CIs of the selected type without adding attribute criteria, skip to Step 6.

Note: For attribute values, mouseover the field and a tooltip displays the data type required.

4. Click Add.

The criteria are added to the expression, which appears in the logic tree and as an expression in the field at the bottom of the page.

5. (Optional) Repeat Steps 2 - 4 to add additional criteria and logic.

Note: For more information about using advanced logic with multiple source criteria, click the Hints link.

6. Click Next.

Define the Target CI Criteria

The Target Criteria page lets you define criteria for the target CI in the relationship.

Follow these steps:

Note: Page fields marked with an asterisk (*) are required.

1. Select the USM type of the target CI from the Class drop-down list in the Target pane.

Consider the following:

- The Attribute drop-down list displays a subset of the USM properties available for the selected type. Also, certain attributes such as MdrProduct attributes and date/time data types are not supported and do not appear in the drop-down.
- When you select a USM type, all available derived types in the USM hierarchy are also included. For example, when you select the ComputerSystem type, its child type VirtualSystem is automatically included in the policy.
- For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.

2. (Optional) Select a USM attribute, comparison type, and attribute value for the expression. If you want to add all CIs of the selected type without adding attribute criteria, skip to Step 5.

Note: For attribute values, mouseover the field and a tooltip displays the data type required.

3. Click Add.

The criteria are added to the expression, which appear in the logic tree and as an expression in the field at the bottom of the page. CIs of the specific type must meet the attribute-based criteria to match.

4. (Optional) Repeat Steps 2 - 3 to add additional criteria and logic. CIs of the specified type must meet the attribute-based criteria to match.

Note: For more information about using advanced logic with multiple target criteria, click the Hints link.

5. Click Next.

Define the Match Criteria

The Match Criteria page lets you define a comparison between attributes of your source and target.

Follow these steps:

Note: Page fields marked with an asterisk (*) are required.

1. Select a source attribute, comparison type, and target attribute from the drop-down lists as follows:

'Source_CI' Attribute

Defines the USM attribute of the source CI that must have some relation to a target CI attribute

Comparison Type

Defines how the specified source and target CI attributes must relate. The options in this drop-down list change based on the data type of the selected attributes.

'Target_CI' Attribute

Defines the USM attribute of the target CI that must have some relation to a source CI attribute.

Note: For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.

2. Click Add.

The criteria are added to the expression, which appears in the logic tree and as an expression in the field at the bottom of the page.

Note: If you define criteria that conflict with the potential values of the selected attributes (for example, if you select Equal To for two enumerated properties that do not share values), an error message opens.

3. (Optional) Repeat Steps 1 - 3 to add additional conditions and logic.

Note: For more information about using advanced logic with multiple match criteria, click the Hints link.

4. Click Next.

Define the Relationship Scope

The Relationship Scope page lets you define the scope of the relationship, which defines which services are affected when creating the relationships.

Follow these steps:

1. Select a scope:

All services

Creates the relationship in all services of which the source CI is a part.

Specific services

Creates the relationship in the specified services of which the source CI is a part.

2. (Specific services only) Do any of the following:

- Click Add and use the Locator tool to select additional services by optionally filtering services, selecting the services you want to include, clicking Add, and then clicking OK.
- In the Locator tool, select a service from the scope list and click Remove to remove the service from the scope list.
- In the Locator tool, click Clear to clear the service scope list.

Note: If you add a dynamic service to the scope that you later rename, you must manually add the renamed dynamic service to the scope again.

3. Click Next.

The Confirm page opens. A Topology Warning dialog may appear if the Service Discovery editor detects that the policy might cause topology errors or unmanageable services. Evaluate the policy before proceeding to verify whether changes are required. For more information, see [Topology Warnings](#) (see page 99).

Confirm the Automatic Relationship Policies Creation

The Confirm page displays the policy expression for the automatic relationship criteria you created.

Follow these steps:

1. Click Finish to confirm creation of the new automatic relationship policy.

The new policy appears under the Automatic Relationships folder. Expand the semantic type to view the relationship policy. You can easily define additional policies for the semantic by right-clicking the semantic and selecting Add.

Note: If you receive a topology warning, check whether changes are required. For more information, see [Topology Warnings](#) (see page 99).

2. Click OK or Save when the automatic relationship policies are complete.

The Service Discovery engine begins scanning the Persistent Store for CIs that match the policy criteria.

How to Create Unmanaged Relationship Policies

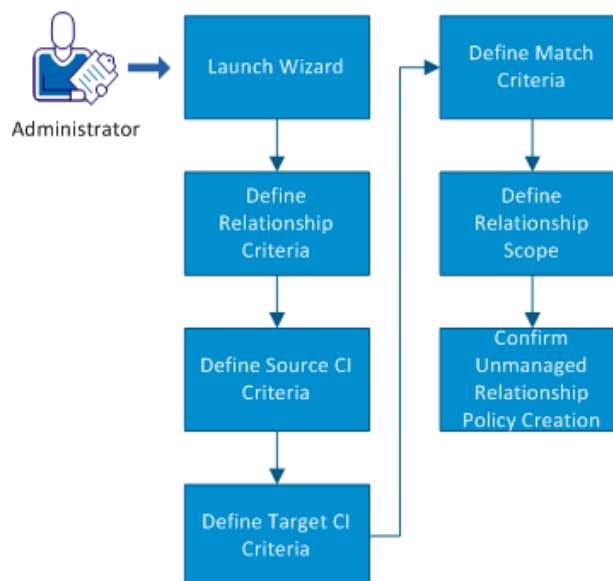
As an administrator, you use the Service Discovery Policy Editor to create an unmanaged relationship policy. An unmanaged relationship is a relationship that exists between two CIs in the CA Catalyst Persistent Store but is not reflected in a CA SOI service model. The unmanaged relationship policy discovers and evaluates unmanaged relationships based on the filtering criteria and whether a source CI is part of a specific or any service. If the filtering criteria match exists, the policy creates a same, but managed, relationship. The newly created managed relationship has the same source and target CIs. This relationship is scoped to all services the source CI is part of. The semantic of the new relationship can be same or different from the original unmanaged relationship depending on the policy definition.

Note: If the source CI is part of multiple services, the policy creates as many relationships between source and target CIs, with each relationship assigned to the appropriate service.

Note: Service Discovery policies support only USM-Core CIs.

Use this scenario to guide you through the process:

How to Create Unmanaged Relationship Policies



1. [Launch the Wizard](#) (see page 93).
2. [Define the Relationship Criteria](#) (see page 94).
3. [Define the Source CI Criteria](#) (see page 95).
4. [Define the Target CI Criteria](#) (see page 96).
5. [Define the Match Criteria](#) (see page 97).
6. [Define the Relationship Scope](#) (see page 98).
7. [Confirm the Unmanaged Relationship Policies Creation](#) (see page 98).

You can also read a [scenario](#) (see page 163) that provides an example for creating unmanaged relationship policies.

Launch the Wizard

Use the Service Discovery Policy Editor to create unmanaged relationship policies.

Follow these steps:

1. Open the Operations Console.
2. Select Tools, Service Discovery Policies.

The Service Discovery Policy Editor opens.

Note: System-wide, only one user can access the Service Discovery Policy Editor at a time.

The Policies tab provides a tree that displays the Dynamic Services, Automatic Relationships, and Unmanaged Relationships.

3. Right-click Unmanaged Relationships in the Policies tab and select Create from the context menu.

The wizard opens on the Relationship Criteria page.

Define the Relationship Criteria

The Relationship Criteria page lets you define the name of the unmanaged relationship policy and filtering criteria that you want to use to discover an unmanaged relationship. The page also lets you specify the type of managed relationship that you want to create after discovering the unmanaged relationship.

Follow these steps:

Note: Page fields marked with an asterisk (*) are required.

1. Specify a name for the policy in the Policy Name field. Ensure that the policy name is different from the existing policy names.
2. Select the Reverse the created relationships option if you want the source and target CIs to swap their positions in the created managed relationship. This implies that the source CI in the unmanaged relationship becomes the target CI in the created managed relationship; similarly, the target CI becomes the source CI.

For example, consider a scenario where a connector publishes unmanaged relationships *IsHostedBy*. These relationships have virtual systems as source CIs and VMWare hosts as target CIs. The service you are populating must have reversed semantic *IsHostFor*, where the sources are VMWare hosts and targets are virtual systems. In such scenarios, enable the Reverse the created relationships option to reverse the source CI and target CI positions in the managed relationships.

Note: In a typical scenario, an unmanaged relationship policy creates a managed copy of the discovered unmanaged relationship. The created managed relationship has the same semantic (if the *Same as Discovered* option is used in Step 3), source CI, and target CI as the original unmanaged relationship. However, if you want to swap the source and target CIs positions in the created managed relationship, use this option.

3. Select a USM relationship type from the Create Relationship of Type drop-down list.

This value specifies the type of managed relationship that you want to create for the discovered unmanaged relationship.

Important! Select *Same as Discovered* if you want to create a managed relationship of the same type as the discovered relationship.

For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.

4. Select the relationship type that you want to discover for identifying the unmanaged relationships.

Note: Select at least one option; you cannot leave this field blank.

5. Use the Add Selected arrow key to move the selected type from Available Types to Selected Types.
6. Click Next.

Define the Source CI Criteria

The Source Criteria page lets you define the source CI criteria for the unmanaged relationship.

Follow these steps:

Note: Page fields marked with an asterisk (*) are required.

1. Select the USM class (type) of the source CI from the Class drop-down list in the Source pane.

Note: Selecting *Entity* as the USM type matches all CI types.

Consider the following:

- The Attribute drop-down list displays a subset of the USM properties available for the selected type. Certain attributes such as MdrProduct attributes and date/time data types are not supported and do not appear in the drop-down.
 - When you select a USM type, all available derived types in the USM hierarchy are also included. For example, when you select the ComputerSystem type, its child type VirtualSystem is automatically included in the policy.
 - For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.
2. (Optional) Select a USM attribute, comparison type, and attribute value for the expression. If you want to add all CIs of the selected type without adding attribute criteria, skip to Step 5.
 3. Click Add.

The criteria are added to the expression, which appears in the logic tree and as an expression in the field at the bottom of the page.

4. (Optional) Repeat the steps to add additional criteria and logic.

Note: For more information about using advanced logic with multiple source criteria, click the Hints link.

5. Click Next.

The Target Criteria page opens.

Define the Target CI Criteria

The Target Criteria page lets you define the filtering criteria for the target CI in the relationship.

Follow these steps:

Note: Page fields marked with an asterisk (*) are required.

1. Select the USM type of the target CI from the Class drop-down list in the Target pane.

Note: Selecting *Entity* as the USM type matches all CI types.

Consider the following:

- The Attribute drop-down list displays a subset of the USM properties available for the selected type. Also, certain attributes such as MdrProduct attributes and date/time data types are not supported and do not appear in the drop-down.
 - When you select a USM type, all available derived types in the USM hierarchy are also included. For example, when you select the ComputerSystem type, its child type VirtualSystem is automatically included in the policy.
 - For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.
2. (Optional) Select a USM attribute, comparison type, and attribute value for the expression. If you want to add all CIs of the selected type without adding attribute criteria, skip to Step 5.
 3. Click Add.

The criteria are added to the expression, which appears in the logic tree and as an expression in the field at the bottom of the page. CIs of the specific type must meet the attribute-based criteria to match.

4. (Optional) Repeat the steps to add additional criteria and logic. CIs of the specified type must meet the attribute-based criteria to match.

Note: For more information about using advanced logic with multiple target criteria, click the Hints link.

5. Click Next.

Define the Match Criteria

The Match Criteria page lets you define a comparison between attributes of your source and target CIs.

Follow these steps:

1. Select a source attribute, comparison type, and target attribute from the drop-down lists as follows:

'Source_CI' Attribute

Defines the USM attribute of the source CI that must have some relation to a target CI attribute.

Comparison Type

Defines how the specified source and target CI attributes must relate. The options in this drop-down list change based on the data type of the selected attributes.

'Target_CI' Attribute

Defines the USM attribute of the target CI that must have some relation to a source CI attribute.

Note: For USM property definitions, see the *USM Schema Documentation* that is provided on the CA SOI Bookshelf.

2. Click Add.

The criteria are added to the expression, which appears in the logic tree and as an expression in the field at the bottom of the page.

Note: If you define criteria that conflict with the potential values of the selected attributes (for example, if you select Equal To for two enumerated properties that do not share values), an error message opens.

3. (Optional) Repeat the steps to add additional conditions and logic.

Note: For more information about using advanced logic with multiple match criteria, click the Hints link.

4. Click Next.

Define the Relationship Scope

The Relationship Scope page lets you define the scope of the relationship, which defines what services are affected when creating the relationships.

Follow these steps:

1. Select a scope:

All services

Creates the relationship in all services of which the source CI is a part.

Specific services

Creates the relationship in the specified services of which the source CI is a part.

2. (Specific services only) Do any of the following:
 - Click Add and use the Locator tool to select additional services by optionally filtering services, selecting the services you want to include, clicking Add, and then clicking OK.
 - In the Locator tool, select a service from the scope list and click Remove to remove the service from the scope list.
 - In the Locator tool, click Clear to clear the service scope list.
3. Click Next.

The Confirm page opens. A Topology Warning dialog may appear if the Service Discovery editor detects that the policy might cause topology errors or unmanageable services. Evaluate the policy before proceeding to verify whether changes are required. For more information, see [Topology Warnings](#) (see page 99).

Confirm the Unmanaged Relationship Policies Creation

The Confirm page displays the unmanaged relationship policy criteria that you created.

Follow these steps:

1. Click Finish to confirm creation of the new unmanaged relationship policy.

The new policy appears under the Unmanaged Relationships folder. Expand the semantic type to view the relationship policy. You can easily define additional policies for the semantic by right-clicking the semantic and selecting Add.

Note: If you receive a topology warning, check whether changes are required. For more information, see [Topology Warnings](#) (see page 99).

2. Click OK or Save when the unmanaged relationship policies are complete.

The Service Discovery engine begins scanning the Persistence Store for relationships matching the policy criteria.

Topology Warnings

Topology warnings indicate that a Service Discovery policy could lead to one of the following topology issues:

Orientation

Occurs when a relationship points from a CI to a top-level service node. All relationships must orient from the root service to its child items. For example, this error could occur if you create an automatic relationship policy with a service as the target CI.

Double relationships

Occurs when more than one relationship using different propagation types exists between two CIs or if equivalent policies are applied to the same service. For example, this error could occur if you create a dynamic service with multiple policies using the same target class and different propagation type, or if you create an automatic relationship with multiple policies using the same source and target class with a different propagation type.

Cycles

Occurs when an item becomes related to itself, two items have separate relationships going in each direction, or when more than two items have a circular relationship structure. For example, this error could occur if you do the following:

- Create a dynamic service policy with a service CI as the target
- Create an automatic relationship policy with the same source and target USM type
- Create an automatic relationship with multiple policies that use the same USM types as source and target. For example, one policy with ComputerSystem as source and BackgroundProcess as target, and another policy with BackgroundProcess as source and ComputerSystem as target.
- Any automatic relationship has the same source CI as another policy's target CI

Unmanageable services

Occurs when a policy creates the potential for an overwhelming number of CIs to be added to a service. For example, this error could occur if you create a dynamic service policy with only a USM type but no target filter criteria, or you create a dynamic service policy using a type that has several children in the USM type hierarchy. For instance, creating a dynamic service policy based on the USM type entity with no target filter conditions would add CIs of all USM types to a service and potentially impact SA Manager performance.

Topology warnings appear in the following locations:

- Confirm page when saving policy
- Details tab in a TOPOLOGY WARNINGS section when you select a policy from the Policies tab.
- Details tab when you select Dynamic Services, Automatic Relationships, Unmanaged Relationships in the Policies tab.
- Policies tab as red service and policy icons
- Policies tab tooltips

You can switch warning display settings in the Set Preferences dialog.

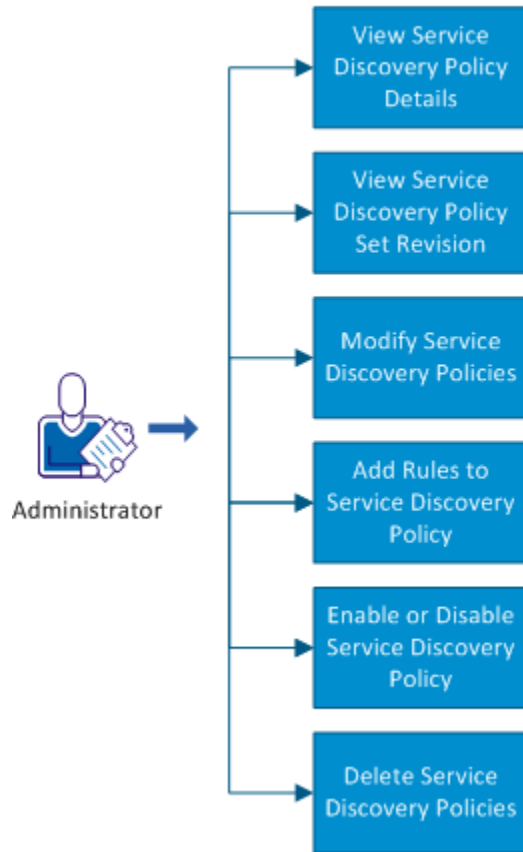
Service Discovery does not prevent the creation of policies that contain warnings. A warning does not mean that the potential problem will definitely occur. If you receive a warning, review the Service Discovery policy to help ensure the policy does not present problems.

How to Manage Service Discovery Policies

As an administrator, you can view, modify, enable, and delete Service Discovery policies. This ability helps you incorporate changed requirements in the policies when required. For example, you can modify policies to add rules to them or delete or disable them if they are not required in the infrastructure.

Use the following scenario to guide you through the process:

How to Manage Service Discovery Policies



- [View Service Discovery Policy Details](#) (see page 102).
- [View Service Discovery Policy Set Revision Information](#) (see page 102).
- [Modify Service Discovery Policies](#) (see page 103).
- [Add Rules to a Service Discovery Policy](#) (see page 103).
- [Enable or Disable Service Discovery Policies](#) (see page 104)
- [Delete Service Discovery Policies](#) (see page 105).

You can perform these tasks in any sequence.

View Service Discovery Policy Details

You can view all Service Discovery policies on the Service Discovery Policy Editor.

Follow these steps:

1. Open the Operations Console.
2. Select Tools, Service Discovery Policies.

The Service Discovery Policy Editor opens.

3. Click a policy in the Policies tab.

The Details tab displays the policy details and also shows [topology warnings](#) (see page 99) that can negatively impact the service topology.

View Service Discovery Policy Set Revision Information

You can view the revision information for the Service Discovery policy set, which informs you of the number of revisions to the policy set, who modified it, and when.

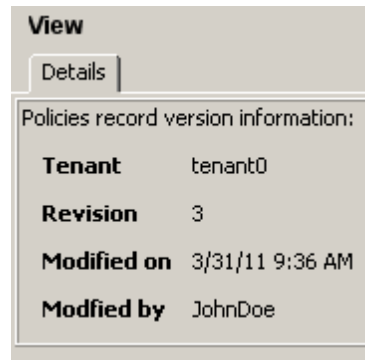
Follow these steps:

1. Select Tools, Service Discovery Policies from the Operations Console.

The Service Discovery Policy Editor opens.

2. Click Service Discovery Policies in the Policies tab.

The policy revision information appears in the Details tab.



Tenant

This field is for a future release of CA SOI.

Revision

Indicates the current revision number of the policy set. Performing a [Redo using redo.rules.bat](#) (see page 110) or a Save on the Service Discovery Policy Editor increases the revision number by one. Performing an [Undo using undo.rules.bat](#) (see page 109) decreases the revision number by one.

Modified on

Indicates the last date and time when the policy set and revision number changed.

Modified by

Indicates either a user ID or cmdLine if the last revision was performed using a [command line operation](#) (see page 107).

Modify Service Discovery Policies

You can modify any Service Discovery policy.

Note: If you rename a dynamic service that is part of the service scope for an automatic relationship policy, you must manually add the renamed dynamic service to the scope again.

Follow these steps:

1. Open the Operations Console.
2. Select Tools, Service Discovery Policies.

The Service Discovery Policy Editor opens.

3. Right-click a policy in the Policies tab and select Modify.

The Service Discovery Policy Editor opens on the Relationship Criteria page (for a dynamic services policy), Source Criteria page (for an automatic relationships policy), and Relationship Criteria page (for an unmanaged relationships policy).

4. Update the conditions as necessary and complete the wizard pages.
5. Save the policies when prompted.

Add a Service Discovery Policy

When you create a policy, you can create one rule at a time for a USM type or relationship type. You can add rules to the same policy if necessary in separate operations.

Follow these steps:

1. Do one of the following in the Policies tab of the Service Discovery Policy Editor:
 - Right-click the dynamic service name and select Add to add a rule to the dynamic service based on a different relationship type from the existing rules.
 - Right-click the relationship folder in a dynamic service and select Add to add a rule to the dynamic service using the selected relationship and based on a different USM type from the existing rules.

- Right-click the automatic relationship name and select Add to add a rule for that relationship type.

The first page of the wizard appears for creating the appropriate rule.

2. Create the new rule.

The rule appears in the appropriate place in the associated policy hierarchy in the Policies tab.

Enable or Disable Service Discovery Policies

You can enable or disable specific Service Discovery policies based on your requirements and control the policy execution schedule.

When you disable a policy, it (disabled policy) does not delete any of the relationships that it has already discovered; this is the default behavior. You can configure this default behavior by configuring the value of the *treatDisabledPoliciesAsDeleted* parameter. This Service Discovery plugin parameter takes *true* or *false* as its value. The default value is *false*, which implies that the relationships that the policy has already discovered are not deleted when the policy is disabled. The value *true* implies that all relationships that the policy has already discovered are deleted when the policy is disabled. You can update the value of the parameter in the Service Discovery plugin properties file `SOI_HOME\ServiceDiscovery\connectivityContext.xml`. Restart Integration Services to make the change effective.

Follow these steps:

1. Open the Operations Console.
2. Select Tools, Service Discovery Policies.

The Service Discovery Policy Editor opens.

3. Select Dynamic Services, Automatic Relationships, or Unmanaged Relationships (as appropriate) in the Policies tab.

A list of related policies appears in the Details tab. For example, if you select Dynamic Services, all dynamic service policies are displayed in the Details tab.

4. Select the appropriate Service Discovery policy that you want to enable or disable in the Details tab.
5. Click the Enables selected policies icon (plus sign with a tick mark) or Disables selected policies icon (minus sign with a tick mark) to enable or disable the policy.

The Enabled column displays the updated status. *Yes* implies that the policy is enabled and *No* signifies that it is disabled.

Note: By default, the policy is enabled when you create it. However, if require a policy to be disabled from the beginning, disable it after creation, but before saving. This way the policy only starts discovering when you enable it manually later.

Delete Service Discovery Policies

You can delete Service Discovery policies to which you have access privileges.

You can select multiple Service Discovery policies and delete them all at the same time from the Service Discovery Policy Editor.

Follow these steps:

1. Open the Operations Console.
2. Select Tools, Service Discovery Policies.

The Service Discovery Policy Editor opens.

3. Select Dynamic Services, Automatic Relationships, or Unmanaged Relationships (as appropriate) in the Policies tab.

A list of related policies appears in the Details tab. For example, if you select Dynamic Services, all dynamic service policies are displayed in the Details tab.

4. Right-click the policy that you want to delete, and select Delete.

A confirmation dialog opens.

5. Confirm the deletion.

Service Discovery Connector Configuration

In the SOI_HOME/ServiceDiscovery directory, you can find the connectivityContext.xml file which contains connector configuration. The file contains three bean definitions:

persistenceDB

Specifies connectivity to Catalyst's Persistent Store database. You can specify driver, URL, database user name, and encrypted password.

ssaDB

Specifies connectivity to the CA SOI database, which is usually the same as the Persistent Store DB. You can specify driver, URL, database user name, and encrypted password.

treatDisabledPoliciesAsDeleted

Specifies how the disabled policies are treated. If set to true, then service discovery deletes all relationships that have been created by the disabled policies so far (while they were enabled). If set to false, discovered relationships are not deleted, but new relationships based on the disabled policies are created.

tickCounts

Specifies the frequency of executing each evaluation algorithm. The 'tick' unit is equal to 5 seconds, thus tick count 6 means 30 seconds.

dynamicGroupTickCount

Frequency for Dynamic Service policy evaluation. This parameter evaluates only recently added CIs (uses timestamps).

unmanagedRelationshipTickCount

Frequency for Unmanaged Relationship policies evaluation. This parameter evaluates only recently added CIs (uses timestamps).

autoRelationshipTickCount

Frequency for Automatic Relationship policies evaluation. This parameter evaluates only recently added CIs (uses timestamps).

This value is 0 by default, because inMemory evaluation of the automatic relationship is the default algorithm. This option make sense if you have SQL Server 2005 or older where inMemory evaluator cannot operate.

inMemoryAutoRelationshipTickCount

Frequency for Automatic Relationship policies evaluation. This is a new, faster evaluator, which requires SQL Server 2008 or newer.

reevaluationTickCount

Frequency for reevaluating existing relationships created by Service Discovery. Deletes already created service discovery relationships which do not match any of the policy.

deleteUnusedSdSheetsTickCount

Frequency for looking for service discovery projections in the notebooks which are not used by any service discovery relationship.

autoRelationshipReevalTickCount

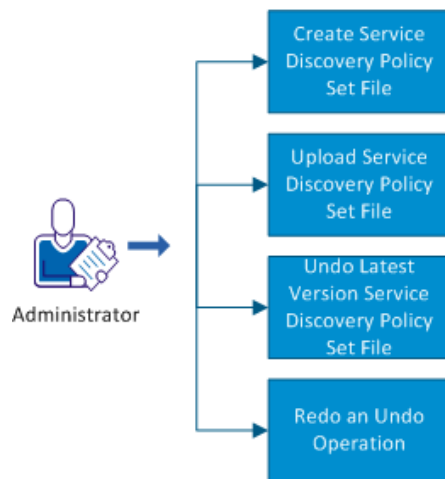
Frequency for Automatic Relationship policies evaluation when a timestamp is not used. It is actually a reset of the evaluation of the automatic relationships.

How to Use Command Line Service Discovery Operations

As an administrator, you can use the command line Service Discovery operations to create and upload service discovery files as well as undo and redo operations. These operations are useful for backing up and restoring service policies.

Use this scenario to guide you through the process:

How to Use Command Line Service Discovery Operations



- [Create a Service Discovery Policy Set File](#) (see page 108).
- [Upload a Service Discovery Policy Set File](#) (see page 108).
- [Undo the Latest Version of a Service Discovery Policy Set](#) (see page 109).
- [Redo an Undo Operation](#) (see page 110).

All .bat files are located in the SOI_HOME\ServiceDiscovery directory on the system where Service Discovery is installed.

Create a Service Discovery Policy Set File

The `read_rules_from_db.bat` file lets you create an `.xml` file that contains the current service policy set definitions from the database. You can use this file as a backup to restore at a later time.

Note: Do not edit the `.xml` file manually. Use the Service Discovery Policy Editor to change your policies.

Follow these steps:

1. Open a command prompt on the SA Manager system.
2. Navigate to `SOI_HOME\ServiceDiscovery`, and run the following command:

```
read_rules_from_db.bat filename.xml
```

filename

Specifies the output `.xml` file you want to create. If a file by the same name already exists, you are prompted to use a different filename and you must run the `.bat` file again with the new filename. You can use a relative or absolute path.

Debug information displays as `read_rules_from_db.bat` creates the file.

File creation completes and the `.xml` file appears either in the directory you specified or the `ServiceDiscovery` directory if you did not specify a path.

Upload a Service Discovery Policy Set File

The `upload_rules_to_db.bat` file lets you upload service policy definitions back into the database from an `.xml` file you created with `read_rules_from_db.bat`.

When you upload the `.xml` file, `upload_rules_to_db.bat` validates the following:

- Class names in policies are valid USM class names
- Properties in the criteria are valid properties of the class
- Operators (StartWith, Contains, Equals, and so on) are valid operator names
- All required properties, such as relationship types, are provided

Uploading a Service Discovery policy set increases the [Revision number](#) (see page 102) by one.

Note: Do not edit the `.xml` file manually. Use the Service Discovery Policy Editor to change your policies.

Follow these steps:

1. Open a command prompt on the SA Manager system.
2. Navigate to SOI_HOME\ServiceDiscovery, and run the following command:

```
upload_rules_to_db.bat filename.xml
```

filename

Specifies the input .xml file created with read_rules_from_db.bat that you want to restore. The upload_rules_to_db.bat file does several semantic checks to avoid loading an invalid file.

Debug information displays as upload_rules_to_db.bat uploads the .xml file to the database.

The upload completes and you are returned to the command line.

3. Close and reopen the Service Discovery Policy Editor so that the policies are refreshed in the Policies tab.

Undo the Latest Version of a Service Discovery Policy Set

The undo_rules.bat file lets you undo the last save of a policy set, whether the save was performed through the Service Discovery Policy Editor or as the result of uploading service policy definitions using upload_rules_to_db.bat. You can undo up to ten saves. Service discovery remembers 10 latest save operations, so the undo can be done 10 times.

Performing an undo decreases the [Revision number](#) (see page 102) by one.

Follow these steps:

1. Open a command prompt on the SA Manager system.
2. Navigate to SOI_HOME\ServiceDiscovery, and run the following command:

```
undo_rules.bat
```

Debug information displays as undo_rules.bat performs the undo.

If the .bat file detects that there is no older version to undo, the operation ends.

The undo completes and you are returned to the command line.

Redo an Undo Operation

The redo_rules.bat file lets you redo an undo you performed using the undo_rules.bat file.

You can perform as many redo operations as undo operations that have been performed. Therefore, if you performed five undo operations, you can perform up to five redo operations. If you attempt to perform more redo operations redo_rules.bat notifies you that the current policy set is the latest.

Performing a redo increases the [Revision number](#) (see page 102) by one.

Follow these steps:

1. Open a command prompt on the SA Manager system.
2. Navigate to SOI_HOME\ServiceDiscovery, and run the following command:

```
redo_rules.bat
```

Debug information displays as redo_rules.bat performs the redo.

If the .bat file detects that the current policy set is the latest, you receive a message and the operation stops.

The redo completes and you are returned to the command line.

Chapter 9: Creating and Working with Service-Level Agreements

This section describes how to create, monitor, and manage service-level agreements for services.

This section contains the following topics:

[How to Create and Work with Service-Level Agreements](#) (see page 111)

How to Create and Work with Service-Level Agreements

As an administrator, you can define, view, edit, attach or detach an SLA with service models. A *service-level agreement* (SLA) is a contract that specifies the service expectations of internal or external customers. An example is the downtime that is acceptable for various resources. You can monitor service performance based on specified thresholds and receive a notification when the SLA approaches or breaches the thresholds.

SLAs let you track service metrics over a specified interval based on specific thresholds. They can help ensure adherence to any quality or availability requirements that an organization must maintain.

SLAs in CA SOI are based on the following attributes:

- Health, Quality, Risk, or Availability metric
- Violation threshold
- SLA time period
- Business hours

Each SLA is based on one of the core CA SOI metrics. An SLA calculates violations that are based on a threshold, spans a defined time period, and is optionally only monitored during defined business hours.

CA SOI uses the measurable provisions that you specify in the SLA to monitor the real-time health of each associated service, and records outage time when the service is down. The recorded time is compared to the SLA thresholds to determine the status of the SLA for a given time period.

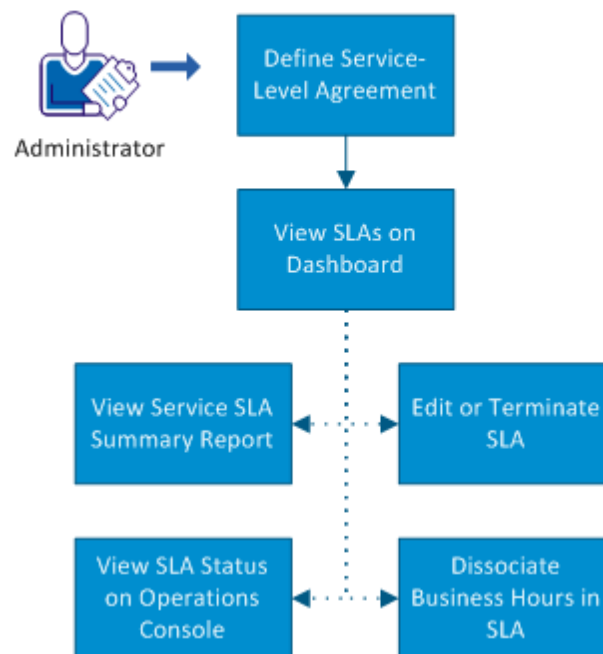
For example, you can define an SLA that tracks the quality of a customer-facing website. If the quality drops below the threshold level, an outage is recorded. If the quality stays below the threshold level past the defined threshold time, a violation is recorded. You can track the SLA status through the Operations Console, Dashboard, and reports to ensure that the website meets the customer quality expectation level.

For an in-depth example, see [SLA Example](#) (see page 119).

When you create an SLA for a service, the SLA evaluation begins after the SLA period start date and when the defined business hours begin, if applicable. The SLA tracks outages and violations; an outage occurs when the service crosses the threshold defined in the SLA. A violation occurs when the service crosses the threshold for the amount of time defined in the SLA.

Use this scenario to guide you through the process:

How to Create and Work with Service-Level Agreements



1. [Define the SLA](#) (see page 113).
2. [View the SLAs on the Dashboard](#) (see page 116).
3. (Optional) [View a Service SLA Summary report](#) (see page 117).
4. (Optional) [View SLA Status on the Operations Console](#) (see page 117).
5. (Optional) [Edit or Terminate an SLA](#) (see page 118).
6. (Optional) [Disassociate the business hours in an SLA](#) (see page 118).

Define a Service-Level Agreement

You define a service-level agreement when creating or editing a service. Each SLA is specific to a service, and a service can only have one SLA.

Follow these steps:

1. Open the Operations Console, and complete one of the following actions:
 - Select Tools, Create New Service.
 - Right-click a service in the Navigation pane and select Edit Service.The Service Modeler opens.
2. Click the SLA tab.
3. Enter a name in the Service Level Objective (SLO) Name field, an optional description in the Description field, and select the Enabled check box to immediately enable the SLA after you save it.
4. Perform the following steps:
 - a. Indicate whether to base the SLA on Health, Quality, Risk, or Availability. For an explanation of these terms, see the *Administration Guide*.
 - b. Select the threshold level. For example, for Health the threshold levels are Minor, Major, Critical, Down, and Unknown. If you select Critical, any period of time when health becomes Critical, Down, or Unknown, it is considered a service outage and is one factor used to determine whether the SLA is violated.
 - c. Select Percent or Seconds and enter a value to determine how long the threshold level must be breached to trigger a violation.

Note: The Seconds value must be greater than 0 and less than 1000000.

5. Click Select for the SLO Period.
6. Perform one of the following steps:
 - Click Create to create a new schedule for the SLA time period. Continue with Step 7.
 - Select an existing schedule and click OK. Continue with Step 9.

A description of the selected period appears in the SLO Period pane.

7. (New SLA period only) Complete the Create Period dialog and click OK.

Start Date

Specifies the date on which to begin calculating the service SLA status. Depending on your selections in the Recurrence pane, the period may not start on the exact start date. For example, if you define a start date of September 15th with a monthly recurrence that starts on the 17th of every month, the period does not start until September 17th.

Recurrence

Specifies how long the SLA period lasts, and when it expires. Select from the following options:

Daily

Specifies that the period recurs after a specified number of days. For example, you can create a period that renews every three days. A daily period starts for the first time on the specified start date.

Weekly

Specifies that the period recurs after a specified number of weeks and begins on a specific day. For example, you can create a period that starts on the first Monday after the specified start date and renews every two weeks.

Monthly

Specifies that the period recurs after a specified number of months and begins on a specific day of the month. For example, you can create a period that starts on the fifth day of the month after the specified start date and renews every three months.

Yearly


Specifies that the period recurs *after* a specified number of years and begins on a specific month and day. For example, you can create a period that starts on January 5 (the first occurrence of this day *after* the start date) and renews every two years.

Note: If the Start Date and the Recurrence date are the same, the recurrence begins the *following* year. If you want a yearly SLO Period recurrence to begin this year, then set a Recurrence date at least one day after the Start Date. For example, if the Start Date is Jan 1, 2013 and the Recurrence date is Jan 1, the first recurrence is Jan 1, 2014. However, if you set the Recurrence Date to Jan 2, the first recurrence happens on Jan 2, 2013.

Note: For all recurrence options, select either No Expiration for the period to recur indefinitely or enter a Schedule End Date when SLA evaluation stops.

8. (New SLA period only) Select the created SLA period and click OK.

A description of the selected period appears in the SLO Period pane.

9. (Optional) Select the Other option in the Business Hours pane to specify business hours to restrict SLA evaluation, and perform one of the following actions:
 - Click Create. Continue with Step 10.
 - Select an available schedule, if any, and click . Continue with Step 11.
10. (New business hours only) Enter the following information in the Create Business Hours dialog and click OK:

Start Date

Specifies when the business hours schedule starts. Depending on your selections in the Recurrence pane, the period may not start on the exact start date. For example, if you define a start date of September 15th with a monthly recurrence that starts every week on Monday, the period does not start until the first Monday after September 15th.

Time

Specifies the daily business hours time interval. These values must be rounded to 30 minute intervals.

Recurrence

Specifies how long the business hours period lasts, and when it expires. Select from the following options:

Daily

Specifies that the period spans every weekday, every weekend day, or every day of the week. For example, you can create a business hours schedule from 9:00-5:00 on Monday through Friday.

Weekly

Specifies that the period starts every week on a specific day.

For all recurrence options, you must either select No Expiration for the period to recur indefinitely or enter a Schedule End Date when the business hours period stops.

Note: You do not have to synchronize the SLA period and business hours recurrence types or recurrence days. For example, you can use a monthly SLA period with weekly business hours. If you define an SLA period that starts on Monday and a business hours schedule that starts on Wednesday, the SLA evaluation begins when the business hours begin.

The created schedule appears in the Current Schedules pane.

Note: If you enter multiple business hours, the product verifies that the hours do not overlap.

11. Click OK.

The service-level agreement is defined.

View SLAs on the Dashboard

The dashboard contains information that you can use to monitor current SLA status. Monitor SLAs from the dashboard as follows:

Current SLA tab

Displays the current state of the SLA for each service. Each SLA has one of the following values in this column:

Compliant

Indicates that the SLA is compliant for the current SLA period.

Violated

Indicates that the SLA threshold is violated for the current SLA period.

Inactive

Indicates that the SLA is inactive, due to the current time being outside of the SLA period or business hours.

If there is no icon in this column, the service does not have a defined SLA.

SLA tab

Displays the following SLA information in the Details of Selected Service pane for the selected service in the Services pane:

SLA Current Status

Displays a pie chart of the SLA status during the current SLA period. This view also displays the following information:

- The SLA type, threshold, and description
- The amount of time in a violated state, if applicable
- The time that the chart was last updated

Click the pie chart to view further details about the current SLA period in the Service SLA Summary report.

SLA History

Displays a bar chart or line chart view of the SLA downtime over the last ten SLA periods. When you scroll over a bar or line, details appear about the nature of the downtime, which can be unplanned (outages and violations), planned maintenance, or unknown.

Note: The bar chart can show a single SLA period, but the line chart requires at least two SLA data points to draw the line.

Click a bar or line to view further details about the represented SLA period in the Service SLA Summary report. The SLA History charts do not display the 'Click to execute the report' tooltip shown in the SLA Current Status view when you scroll over the bar or line. The SLA History part instead shows further SLA period details, but you can still generate a report by clicking the bar or line.

Both views present SLA states as the following color-coded values:

- **Up:** Green
- **Unplanned:** Red

Note: Unplanned maintenance is the total of outage and violation time.

- **Maintenance:** Brown
- **Unknown:** Gray

Note: The SLA charts display the date in yyyy/mm/dd format. For example, 2010/05/08 is May 8th, 2010. The reports generated from the SLA charts use mm/dd/yyyy format (for example, 05/08/2010).

View Service SLA Summary Report

The Service SLA Summary report compiles the SLA data into a series of tables and charts for comprehensive SLA monitoring.

To run the Service SLA Summary report, perform one of the following actions:

- Click the charts in the SLA tab of the Dashboard. This option automatically runs the report for the specific SLA and the SLA time period represented in the chart.
- Click the Reports link and select the Service SLA Summary report. When you run the report from the reporting interface, you must enter settings that define the SLAs to include and the report time period.

Note: For more information about this report and reporting, see the *User Guide*.

View SLAs Status on the Operations Console Data

Monitor the SLA status from the following places in the Operations Console:

Services tab in Contents pane

Displays a list of all services when you select the top-level Services item in the Services tab. The SLA Status column on this tab displays the current status of each service's SLA.

Information tab in Component Detail pane

Displays the following SLA information for the selected service in the Service Level Agreements table:

- Name
- Description

- Current state and last known state
- Last status update
- First violation time

Edit or Terminate an SLA

You can change the name, description, or enabled state of an existing SLA.

You must terminate the existing SLA and create a new one if you need to change any other properties, such as threshold or SLA period. Once SLA evaluation starts, changes to the SLA period, threshold, or business hours may cause an incorrect determination of SLA status.

Follow these steps:

1. In the Service Modeler, click the SLA tab.
2. Perform one of the following actions:
 - Edit the SLA name or description, or clear the Enabled check box to temporarily stop SLA evaluation, and click Save.
 - Click Terminate SLA and click Yes to confirm.

Disassociate Business Hours in a Service-Level Agreement

Before deleting Business Hours schedules, you must disassociate the related SLA.

Follow these steps:

1. In the Service Modeler, click the SLA tab.
2. Select the Other option in the Business Hours pane.
3. Move any schedules from the Current Schedules pane to the Available Schedules pane and click OK.

The SLA is saved and the hours are removed.

SLA Example

Consider a retail website that must be available at all times for customers to place orders. The business enforces concrete performance and availability requirements on the website. Not only must the website be available for a certain amount of time, it must perform above a certain threshold so that customers are not frustrated by a slow response time.

The administrator must ensure that the website response time does not dip below a defined threshold for more than 8 hours in a month.

Assume the following for this example:

- The administrator modeled a service for the website that includes the components that enable the website to perform as expected. The service includes the web servers, order processing applications, product databases, network hosts and routers, and so on. The service also accurately represents the relationships among the CIs and the significance of each component.
- The benchmark to which the website must adhere is response time. The domain manager that manages the service CIs can calculate the response time and translate the metric to a severity that you can associate with service quality in CA SOI.

To [create an SLA](#) (see page 113) that monitors the website response time against defined benchmarks, the administrator enters the SLA properties for the website service similar to the following:

- Base the threshold on Quality in the SLA tab that must equal or exceed a moderately degraded status for 28800 Seconds of the SLA time period.

This threshold assumes that a moderately degraded service quality indicates that the website response time has crossed the response time threshold. It also assumes that any status more severe than moderately degraded also indicates a response time threshold breach.
- Create an SLA period that recurs monthly and starts on the first day of every month.
- Select 24x7 for Business Hours, because the web site must be available for customer use at all times.

After the SLA period begins, the administrator can monitor SLA status to ensure adherence to the response time benchmark. For example, you could schedule a Service SLA Summary report to run every day at a specific time to summarize the downtime, outages, and SLA status for that day.

Chapter 10: Working with Customers

The topics in this section describe how to create and manage customers and their sub-customers.

This section contains the following topics:

[How to Create and Manage Customers](#) (see page 121)

How to Create and Manage Customers

As an administrator, you define customers and associate them with service models to see the impact of service degradation on the customer.

A *customer* in CA SOI is any consumer of a managed service. A customer can represent a specific division of the company such as a product division, region, or city. Customer impact provides IT personnel with an accurate understanding of what fault conditions really mean to a particular customer. For example, the administrator defines a customer to include all services within South America. If a fault occurs on a CI, such as a router, the IT personnel can easily determine how South America is impacted.

Alerts can indicate both service impact and customer impact. To determine the customer impact, an administrator creates a customer and assigns services to that customer.

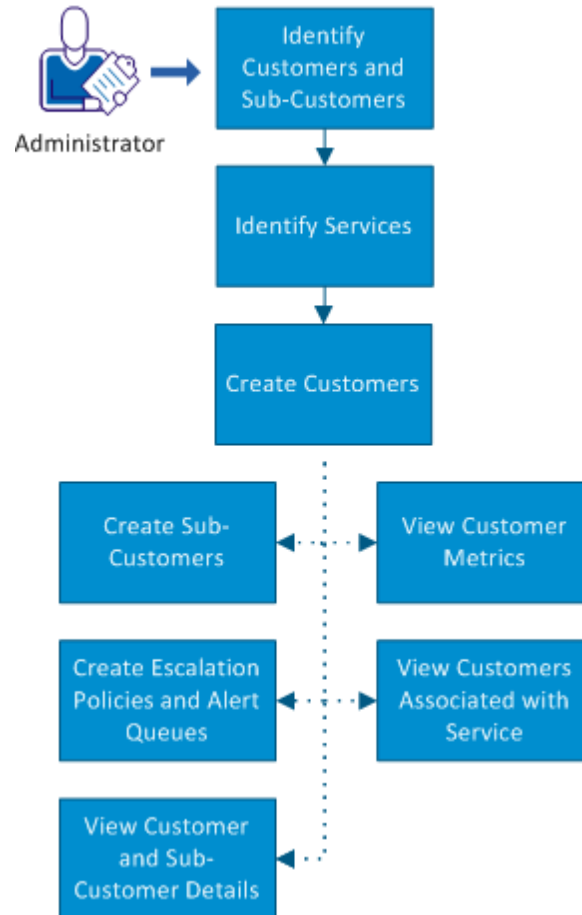
Sub-customers represent smaller divisions of a larger entity customer. For example, you can create a customer that represents a geographical division and sub-customers for smaller regions. You could continue to create additional sub-customers for cities, office buildings, office floors, and so on. You could also create a customer that is your cloud network and the sub-customers are the actual components utilized by those paying for your service. The number of sub-customers and nesting levels depends on how you want to monitor the customer impact.

Administrators can place a priority on a customer, which weighs that customer as higher or lower than other customers.

Administrators can use customer conditions as criteria to define escalation policy and alert queue rules. For example, you can create an escalation policy that sends an email to a specific IT person when the customer impact is Severe. An administrator can also use runtime tokens that display the customer name and severity in escalation actions. For more information about using runtime tokens, see Expandable Runtime Tokens.

Use this scenario to guide you through the process:

How to Create and Manage Customers



1. [Identify the customers and their sub-customers](#) (see page 123)
2. [Identify the services](#) (see page 124)
3. [Create the customers](#) (see page 124)
4. (Optional) [Create the sub-customers](#) (see page 126)
5. (Optional) Perform any of the following tasks:
 - [Configure the customer priority and labels](#) (see page 127).
 - [Create the escalation policy and alert queues](#) (see page 128) using customer-related conditions.
 - [View the customer and the sub-customer information](#) (see page 129).

- [View the customer metrics](#) (see page 130).
- [View the customers that are associated with a particular service](#) (see page 131).

[An example](#) (see page 132) shows you a real-world scenario for creating the customers, sub-customers, priorities, assigning services, and creating escalation policies and alert queues utilizing customers.

Identify Customers, Sub-Customers, and Priorities

First, you identify your customers. Is your customer a region, a company department, an actual buyer of your cloud services, or another entity?

Sub-customers are customers within a larger customer entity. For example, say that your company has divisions in North and South America and each division has the following departments: Accounting, Finance, and Human Resources.

You can identify the following customers (divisions) and sub-customers (departments). You can then monitor the customer impact of any managed services on any division as a whole or on an individual department:

North America

- Accounting
- Finance
- Human Resources

South America

- Accounting
- Finance
- Human Resources

For each customer and sub-customer, you can optionally apply a customer priority.

Although the alert severity determines the customer impact, you can use both customer priority and customer impact to determine escalation policy and alert queues. For escalation policies, if an alert impacts multiple customers, the escalation policies on the higher priority customer are applied first. For alert queues, you can use the customer priority and the customer impact when building rules. For example, you can define an alert queue where alerts of any severity are included if the customer priority is a certain level. Similarly, you can define an alert queue where alerts are included if the customer impact is above a certain threshold, such as Moderate.

Consider an administrator that creates customers for its Infrastructure as a Service (IaaS) services. Each customer represents a virtual machine or network in a cloud that are sold to clients. The administrator determines that a particular virtual machine is the highest importance due to the service level agreement. The administrator creates a customer that is comprised of that virtual machine's services and sets the priority to 10 (the highest priority). The administrator can then configure the escalation policy to trigger first on the highest priority customer. Operators then resolve the highest priority virtual machine issues quickly and they maintain the service level agreement.

Identify Services

Once you have identified your customers and sub-customers, you can easily determine services that are associated with each customer. When you create your customers and sub-customers in CA SOI, you create an association with the services. Therefore, you must [have a service model](#) (see page 33) available from which you can identify the services.

Create Customers


You create customers using a wizard.

Follow these steps:

1. Open the Operations Console and click the Customers tab.
2. Click the Create a New Customer icon.
3. Enter a Customer Name, Customer Identity, and optionally a Description. These values appear on the Information tab in the Contents pane.

Note: The customer identity (customer ID) uniquely identifies a customer. You cannot have duplicate customer IDs, but you can have duplicate customer names.

4. (Optional) Select a customer priority from the Priority drop-down list. You can select a value from 1 through 10; 1 represents the lowest priority and 10 represents the highest. If an alert impacts multiple customers, the escalation policies on the higher priority customer are applied first. You can also use the customer priority and customer impact to determine alert queues.

Note: Click the configure priorities  icon next to the Priority drop-down list to [configure the customer priority](#) (see page 127). You can assign a meaningful label to each priority level; for example, Gold to represent the highest priority. If you do not want to use a specific priority level, you can disable it.

5. Click Next.

You use this screen to assign the services to your new customer.

Consider the following items:

- When you assign a service to a customer, CA SOI automatically assigns the subservices also.
- An asterisk (*) indicates sub-services that are automatically assigned with the selected parent service.
- Assigning a parent service automatically includes all its sub-services. Similarly, removing a parent service automatically removes its sub-services, unless a sub-service is a child to another parent in the Assigned Services list. The sub-services are prefixed with * and you cannot remove the sub-services unless their parent(s) are also removed. Select the Show top level parent services only check box to hide all sub-services and to show only the highest level parent services in the Available Services list.

6. Use the arrows to add or remove services from the Available Services and Assigned Services lists. You can enter a string to filter either list.

7. Click Next.

This screen lets you assign the user groups that have access to the new customer. For more information about how service and role-based security affect customers, see the *Administration Guide*.

Consider the following items:

- The user group must also have access privileges to the services you assigned to the new customer.
- You can also manage user group access to customers in the Users tab.

8. Use the arrows to add or remove user groups from the Available Groups and Allowed Groups lists. You can enter a string to filter either list.

Note: User groups marked with an asterisk have their access set to all customers either by default or by an administrator.

9. Click Next.

10. Verify the new customer information and click Finish.

The new customer displays in the Customers tree.

Create Sub-Customers

Sub-customers represent smaller divisions of a larger entity customer. You create sub-customers under customers.


Follow these steps:

1. Open the Operations Console and click the Customers tab.
2. Select the customer or sub-customer for which you want to create a sub-customer.
3. Click the Create a Sub-Customer icon.

4. Enter a Customer Name, Customer Identity, Customer Priority, and optionally a Description, which appears on the Information tab in the Contents pane.

Note: The Customer identity (customer ID) uniquely identifies a sub-customer. You cannot have duplicate customer IDs.

5. Select a priority from the Priority drop-down list. Customer priority helps to determine the impact of an alert to a customer (customer impact). If an alert impacts multiple customers, the escalation policies on the higher priority customer are applied first.

Note: Use the configure priorities  icon next to the Priority drop-down list to [configure the customer priority](#) (see page 127). You can assign a meaningful label to each priority level; for example, Gold to represent the highest priority. If you do not want to use a specific priority level, you can disable it.

6. Click Next.
7. Use the arrows to add or remove services from the Available Services and Assigned Services lists. You can enter a string to filter either list.

Consider the following items:

- When you assign a service to a customer, CA SOI automatically assigns the subservices also.
- An asterisk (*) indicates sub-services that are automatically assigned with the selected parent service.
- Assigning a parent service automatically includes all its sub-services. Similarly, removing a parent service automatically removes its sub-services, unless a sub-service is a child to another parent in the Assigned Services list. The sub-services are prefixed with * and you cannot remove the sub-services unless their parent(s) are also removed. Select the Show top level parent services only check box to hide all sub-services and to show only the highest level parent services in the Available Services list.

8. Click Next.

This screen lets you assign the user groups that have access to the new sub-customer.

Consider the following items:

- If a parent customer is given a permission for a user group, the child customer also gets the permission. For more information about configuring role-based security, see the *Administration Guide*.
- The user group must also have access privileges to the services you assigned to the new sub-customer.
- You can also manage user group access to customers in the Users tab.

9. Use the arrows to add or remove user groups from the Available Groups and Allowed Groups lists. You can enter a string to filter either list.

Note: User groups marked with an asterisk have their access set to all customers either by default or by an administrator.

10. Click Next.

11. Verify the new sub-customer information and click Finish.

The new sub-customer displays in the Customers tree.


Configure Customer Priorities and Labels

You can optionally apply a priority that can determine escalation policy or alert queues. If an alert impacts multiple customers, the escalation policies on the higher priority customer are applied first. You can configure the customer priority while creating or editing your customer.

You can also assign a meaningful label to each priority level to represent your unique environment. For example, you can use labels such as Gold, Silver, and Bronze to represent the particular customer tiers. You could label customers with the highest priority (say, a value of 10) as "Gold", medium priority (say, a value of 5) as "Silver" and so on.

Additionally, if you do not need certain priority levels, you can disable them. For example, per your organization policy, you set only four priority levels. In this case, you can disable the remaining six priority levels to avoid any confusion when creating alert queue rules or defining escalation policy.

Follow these steps:

1. Open the Define Customer screen in the customer wizard either by [creating a customer](#) (see page 124) or editing a customer.
2. Click the configure priorities  icon next to the Priority drop-down list.
3. Enter appropriate names in the priority level fields (for example, Normal in Priority Level 1 and Moderate in Priority Level 2).

The values that you enter in these fields are represented as options for the corresponding priority levels in the Priority drop-down list.

4. Select the required Disable option for the priority level that you do not want to use.
The disabled priority levels do not display as options in the Priority drop-down list.
5. Click OK.
The customer priority is configured.

More information:

[Create Customers](#) (see page 124)

[Create Sub-Customers](#) (see page 126)

Create Escalation Policies and Alert Queues

You can use the following customer attributes as criteria for escalation policy or alert queues:

- Customer ID
- Customer impact
- Customer name
- Customer priority
- Highest customer impact
- Highest customer priority
- Number of impacted customers

For example, you can create an alert queue with all alerts associated with customers assigned a Customer Priority higher than 8 and a Customer Impact that is Severe.

Similarly, you can create an escalation policy where a specific technician receives an email when the Customer Impact for his or her region (where the region is the customer) is Moderate or higher.

You can also use customer-related runtime tokens in escalation actions. For more information about working with escalation policy, runtime tokens, and alert queues see the *Event and Alert Management Best Practices Guide*.

View Customer and Sub-Customer Details

You can view information about customers and sub-customers, including list of sub-customers, associated services, and associated service alerts.

Follow these steps:

1. Open the Operations Console and click the Customers tab in the Navigation pane.

A list of all customers displays. The columns to the right of the customer name display the number of alerts of each severity for the customer and the total number of alerts for the customer. These alerts are the alerts that are impacting the services that are assigned to the customer.

If a customer is associated with a service and an alert is generated that impacts that service, the alert also impacts the associated customer. An alert can impact multiple customers for the following reasons:

- Because multiple customers are associated with a single service.
- Because an alert impacts multiple services and each service has an associated customer.

2. View the customer tree icon color for any customer to see the overall customer health. Of all the services that are assigned to the customer, the service with the worst health represents the customer health. The customer tree icon color, therefore, shows the color that is based on that service's health.

3. Select a customer (or sub-customer) in the Customers tree and do any of the following tasks:

- Click the Alerts tab to view all the alerts (from the services that are assigned to the selected customer) for the customer. This tab includes information about alert severity, category, summary, count of impacted customers, and so on.

Note: You can also display the number of impacted customers by adding the # Impacted Customers column.

When you remove or add a service to a customer, the associated alert list is changed accordingly.

Note: The alert list for the parent customer shows the aggregate of all alerts from all of its child customers.

- Click the Services tab in the Contents pane to view a list of all services that are associated with the customer. This tab includes information about the service name, health, risk, granularity, and so on.

- Click the List tab in the Contents pane to view a list of sub-customers. This tab includes information about all the sub-customers available under the selected parent customer. The tab displays the customer name, ID, priority, impact (or customer health), quality, risk, and description.

Note: Of all the services that are assigned to the customer, the service with the worst health represents the customer health. Similarly, the service with the worst quality represents the customer quality, and the service with the maximum risk represents the customer risk.

- Click the Information tab in the Contents pane to view general information about customers. This tab includes general information about the customer: name, ID, and priority. This tab also includes current metric information: health, quality impact, risk, and priority. The Security pane shows the user groups that have access to the customer.
 - Click the Customer Impact tab to view the customer impact level. This tab is available to user groups with access privileges only.
4. Review the information. For the Services and Alerts tabs, you can use the additional tabs in the Component Detail pane to get detailed information.

View Customer Metrics

Customer metrics are similar to service metrics; however, customer metrics show the health, quality, and risk that are associated with a customer rather than a service. CA SOI provides the following customer metrics:

customer health

Identifies the worst state that is currently held by either customer quality or customer risk.

customer quality

Identifies the worst customer quality for the services that are assigned to the customer. Quality indicates the level of excellence that consumers of an IT service experience. The highest propagated impact of an associated quality alert determines the customer quality value.

customer risk

Identifies the maximum risk on the customer. This metric indicates the likelihood of delivering the quality of service that is required to support the overall business objectives. The highest propagated impact of an associated risk alert determines the customer risk value.

On the Services tab:

1. Open the Operations Console and click the Services tab.
2. Select a service.
3. Click the Customers tab in the Contents pane.

The Customers tab displays the customer name, identity, description, and customer metrics.

On the Customers tab:

1. Open the Operations Console and click the Customers tab.
2. Select a customer.
3. Click the Information tab in the Contents pane.

The customer metrics display in the Information tab.

View Customers Associated with a Service

You can view all the customers that are associated with a specific service. This information helps you analyze the impact that a particular service has on different customers. When an administrator associates or removes a service from a customer, CA SOI updates the information in the Customers tab accordingly.

Follow these steps:

1. Open the Operations Console and click the Services tab.
2. Select a service for which you want to see the associated customers.
3. Click the Customers tab in the Contents pane.

All customers that are associated with the selected service display in the pane. The pane also provides detailed information about the related customers; for example, name, ID, description, priority, and so on.

Example: Creating and Working with Customers


In this example, Forward Inc. is an MSP that provides finance, operations, and sales IaaS solutions for customers. Forward Inc. has already modeled several services to monitor their network. To ensure the highest level of service for their clients, Forward Inc. now wants to see how alerts on their modeled services impact specific clients. The clients are represented as customers in CA SOI. Forward Inc. Operators are assigned to monitor different services and customers and are provided service and customer access accordingly. The Admin Operator has access to all services and customers.

Forward Inc. provides three service levels: Bronze, Silver, and Gold. Clients who pay for Bronze service are provided the minimum service level and Gold provides the highest service level.


Identify Services

Services Available: Finance, Operations, Sales

The administrator models the three services that are based on the network for each IaaS offering as shown in the following graphic:



The administrator models the Finance and Operations services with two computer system CIs. Similarly, the administrator models the Sales service with three computer system CIs. The modeled services are shown in the following graphics:



Your company's service models are obviously more complex than the models provided in this example. You employ the techniques found in the *Service Modeling Best Practices Guide* such as propagation policy, CI significance, and complex relationships. However, for our demonstrative purposes, we are using simple service models and not delving into the particulars of each service model. Our focus is on the service alerts and how they impact customers of those services.

Identify Customers and Priorities

Customers: Company-A, Company-B

Forward Inc. has two clients that use the IaaS solution services. The administrator plans to create two customers in CA SOI, Company-A and Company-B, to represent these two clients. Company-A has paid a premium for a higher service level, so the administrator will assign a higher customer priority to Company-A.

Company-A uses the Finance and Operations IaaS solutions and Company-B uses the Operations and Sales IaaS solutions, so the administrator will assign the services as follows when creating the customers:

Company-A: Finance and Operations services

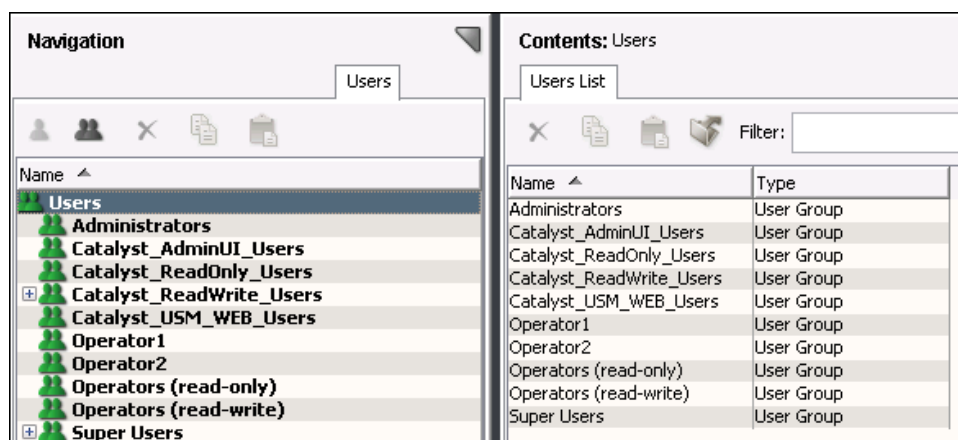
Company-B: Operations and Sales services

Company-A pays for the Gold service level and Company-B pays for the Bronze service level.

An administrator could elect to create sub-customers also. For example, Company-A and Company-B could have sub-customers that represent specific divisions of each client's company. If different divisions used different solution products (or a combination of solution products), an operator could see how alerts impact a specific division. However, for this example, we will use customers only.

Create User Groups

The administrator creates two "Operator" user groups, Operator1 and Operator2 in the Users tab. The users assigned to these user groups are operators that monitor customers. The administrator will assign the customers to Operator1 and Operator2 when using the wizard to create the customers.



The administrator would also assign one or more users to the Operator1 and Operator2 user groups. However, for simplicity in this example, we refer to the users as the Operator1 user and Operator2 user. Remember that these user groups include all users assigned to them.

The administrator also uses the Users tab to ensure that each user group has access to the services they monitor as described in the previous section.

Create Customers, Assign Services, and Assign User Groups

The administrator creates the customers, Company-A and Company-B, and assigns the operator user groups as follows:

Operator1: Company-A customer

Operator2: Company-A customer, Company-B customer

Operator1 will monitor the service for the Company-A customer and Operator2 will monitor the services for both the Company-A and Company-B customers.

Note: Operator1 has permissions to the services (which are Finance and Operation services) of Company-A. Operator2 has permissions to the services (which are Finance, Operation, and Sales services) of Company-A and Company-B. Permissions to the services are given from the Services or Users tab. For more information, see the "Configuring Role-Based Security" section in the *Administration Guide*.

The administrator launches the Create New Customer wizard and completes the Define Customer screen with the Company-A information as shown in the following graphic:

The screenshot shows the 'New Customer' wizard window. On the left, a 'Steps' sidebar lists: 1. Define Customer (highlighted), 2. Assign Services, 3. Assign User Groups, and 4. Confirm. The main area is titled 'Define Customer' and contains the following fields:

- Customer Name ***: Text box containing 'Company-A'
- Customer Identity ***: Text box containing 'A'
- Description**: Empty text box
- Priority**: A dropdown menu currently showing '10'. To its right is a 'Hints...' button.

The customer Priority should reflect the Gold, Silver, and Bronze service levels. Therefore, the administrator will modify the Priority Levels to create more meaningful drop-down list labels. The administrator launches the Configure Customer Priority dialog and sets the following Priority Level labels:

- Priority Level 10 label to Gold
- Priority Level 6 label to Silver
- Priority Level 1 to Bronze

The administrator disables the unused priority levels. The completed dialog is shown in the following graphic:

The screenshot shows the 'Configure Customer Priority' dialog box. It contains the following text: 'You can assign customers a priority level from 1 (lowest) to 10 (highest). Assign a label for each priority level, or disable it. [More details...](#)'

Disable	Priority Level	Label
<input type="checkbox"/>	Priority Level 10	Gold
<input checked="" type="checkbox"/>	Priority Level 9	9
<input checked="" type="checkbox"/>	Priority Level 8	8
<input checked="" type="checkbox"/>	Priority Level 7	7
<input checked="" type="checkbox"/>	Priority Level 6	Silver
<input checked="" type="checkbox"/>	Priority Level 5	5
<input checked="" type="checkbox"/>	Priority Level 4	4
<input checked="" type="checkbox"/>	Priority Level 3	3
<input checked="" type="checkbox"/>	Priority Level 2	2
<input type="checkbox"/>	Priority Level 1	Bronze

At the bottom are 'OK' and 'Cancel' buttons.

The administrator now sees the Company-A customer with the priority label Gold as shown in the following graphic:

The screenshot shows the 'New Customer' window with the 'Define Customer' tab selected. On the left, a 'Steps' sidebar lists: 1. Define Customer (highlighted), 2. Assign Services, 3. Assign User Groups, and 4. Confirm. The main area contains the following fields:

- Customer Name ***: Text box containing 'Company-A'
- Customer Identity ***: Text box containing 'A'
- Description**: Empty text box
- Priority**: A dropdown menu showing 'Gold' and a 'Hints...' link.

The administrator then assigns the Finance and Operations services to Company-A as shown in the following graphic:

The screenshot shows the 'New Customer' window with the 'Assign Services' tab selected. The 'Steps' sidebar now highlights '2. Assign Services'. The main area is divided into two panels:

- Available Services**: A list box containing 'Sales'. Below it is a 'Filter:' text box and 'Displaying 1 of 1'.
- Assigned Services**: A list box containing 'Finance' and 'Operations'. Below it is a 'Filter:' text box and 'Displaying 2 of 2'.

Between the panels are four arrow buttons for moving services: a single right arrow, a single left arrow, a double right arrow, and a double left arrow. At the bottom left, there is a checkbox labeled 'Show top level parent services only' and a note: '* Indicates sub-services that are automatically assigned with the selected parent service. [Details...](#)'

The administrator assigns the Operator1 user group to the Company-A customer in the Assign User Groups screen as shown in the following graphic:

The screenshot shows the 'Assign User Groups' screen within the 'New Customer' window. On the left, a 'Steps' sidebar lists: 1. Define Customer, 2. Assign Services, 3. Assign User Groups (highlighted), and 4. Confirm. The main area is titled 'Assign User Groups'. It features two list boxes: 'Available Groups' on the left and 'Allowed Groups' on the right. The 'Available Groups' list contains: Catalyst_AdminUI_Users, Catalyst_ReadOnly_Users, Catalyst_ReadWrite_Users, Catalyst_USM_WEB_Users, and Operator2. The 'Allowed Groups' list contains: *Administrators, *Operators (read-only), *Operators (read-write), *Super Users, and Operator1 (which is highlighted). Between the lists are four buttons: a right arrow, a left arrow, a double right arrow, and a double left arrow. Below each list box is a 'Filter:' text box and a 'Displaying 5 of 5' status. At the bottom, a note states: '* Indicates the user group has access set to all customers.'

The administrator reviews the Company-A customer information in the Confirm screen and completes the customer creation.

The screenshot shows the 'Confirm' screen within the 'New Customer' window. The 'Steps' sidebar on the left shows: 1. Define Customer, 2. Assign Services, 3. Assign User Groups, and 4. Confirm (highlighted). The main area is titled 'Confirm'. It contains form fields for 'Customer Name' (with the value 'Company-A') and 'Customer Identity' (with the value 'A'). Below these is a 'Description' text area. The 'Priority' is set to 'Gold'. Under 'Added Services', there is a list box with the following items: Name (with a small upward arrow icon), Finance (highlighted), and Operations.

The administrator uses the same process to create the Company-B customer as shown in the Confirm screen. The administrator assigns the Operations and Sales services and sets the Priority to Bronze. The Confirm screen is shown in the following graphic:

New Customer

Steps

1. Define Customer
2. Assign Services
3. Assign User Groups
- 4. Confirm**

Confirm

Customer Name **Customer Identity**

Description

Priority Bronze

Added Services

Name
Operations
Sales

View Alerts on Customers

The administrator clicks the Customers tab and selects the parent Customers node in the customer tree, then clicks the Alerts tab in the Contents pane. The Alerts tab displays all alerts on the Company-A and Company-B customers as shown in the following graphic:

Navigation

Customers

Contents: Customers

Alerts

Filter:

Filtered By: Maintenance

Severity	Date/Time	Name	Class	Category	Summary	Service Impact	# Impacted Servi...
Major	Jun 27, 2012 2:11:40 PM EDT	Operations	Service		Service is moderately degraded due to 1 active r...	Moderate	1
Critical	Jun 27, 2012 2:12:51 PM EDT	Sales	Service		Service is severely degraded due to 1 active roo...	Severe	2
Critical	Jun 27, 2012 2:36:45 PM EDT	Server3	Computer S...		Service is stopped	Moderate	2
Critical	Jun 27, 2012 2:36:44 PM EDT	Server5	Computer S...		Service is stopped - Again	Moderate	1
Critical	Jun 27, 2012 2:36:44 PM EDT	Server5	Computer S...		Service is stopped - From event notification	Moderate	1
Critical	Jun 27, 2012 2:17:20 PM EDT	Server6	Computer S...	Risk	Low Memory - nearing threshold	Severe	1

The administrator then selects the Company-A customer in the Navigation pane and clicks the Information tab to view the health status. The health status is Major as shown in the following graphic:

Navigation

Customers

Name	Σ	▼	▼	▼	▼
Customers	6	5	1		
Company-B	6	5	1		
Company-A	2	1	1		

Contents: Company-A

Information

Company-A

General Information

Customer Name: Company-A

Customer Identity: A

Health: Major

Quality Impact: None

Risk: None

Description

Priority: Gold

Security

Filter: Displaying 5 of 5

Name	Type	Source Type
Administrators	UserGroup	All Access
Operator1	UserGroup	Local
Operators (read-only)	UserGroup	All Access
Operators (read-write)	UserGroup	All Access
Super Users	UserGroup	All Access

Remember that the customers available to a user depend on the user group the user is assigned to and that user group's customer assignment. Because the administrator has access to all customers and services, the administrator sees both Company-A and Company-B customers on the Operations Console. If an Operator1 user logs in to the Operations Console, the user sees only Company-A. The administrator assigned both Company-A and Company-B to the Operator2 user group, so an Operator2 user sees both the Company-A and Company-B customers.

The administrator then views the Information tab for the Company-B customer, which shows a health status of Critical as shown in the following graphic:

Navigation

Customers

Name	Σ	▼	▼	▼	▼
Customers	6	5	1		
Company-B	6	5	1		
Company-A	2	1	1		

Contents: Company-B

Information

Company-B

General Information

Customer Name: Company-B

Customer Identity: B

Health: Critical

Quality Impact: None

Risk: Severe

Description

Priority: Bronze

Security

Filter: Displaying 4 of 4

Name	Type	Source Type
Administrators	UserGroup	All Access
Operators (read-only)	UserGroup	All Access
Operators (read-write)	UserGroup	All Access
Super Users	UserGroup	All Access

Create Alert Queues Using Customer Attributes

The administrator wants to create an alert queue named "Gold Customers" that shows the alerts on all Gold service level (priority) customers. The administrator launches the New Alert Queue wizard and defines the queue criteria where the Customer Priority attribute is equal to the Gold level. The completed Define Queue Criteria screen looks like the following graphic:

The screenshot shows the 'New Alert Queue' wizard window. On the left, a 'Steps' sidebar lists: 1. Define Queue Criteria (highlighted), 2. Assign Escalation Policies, 3. Assign User Groups, and 4. Confirm. The main area is titled 'Define Queue Criteria'. It contains a 'Queue Name *' field with 'Gold Customers', a 'Queue Priority' slider set to 5, and a 'Hints...' link. Below is a 'Description' text area. The 'Queue Criteria' section has an 'Attribute' dropdown, a 'Comparison Type' dropdown, an 'Ignore Case' checkbox, and an 'Attribute Value' field. 'Add', 'Apply', and 'Clear' buttons are to the right. A 'Hints...' link is below. A logic tree shows an 'AND' condition with 'Customer Priority Equal To "Gold"'. On the right, buttons for 'New AND', 'New OR', 'AND/OR', 'Cut', 'Copy', 'Paste', and 'Clear' are visible.

The administrator clicks the Alert Queues tab on the Operations Console, clicks the Alerts tab in the Contents pane, then clicks the Customer Impact tab in the Component Detail pane. The Operations Console appears like the following graphic:

The screenshot displays the Operations Console interface. On the left, the **Navigation** pane shows the **Alert Queues** tab selected, with a tree view containing **Alert Queues**, **Gold Customers**, and **Default**. The **Contents** pane on the right shows the **Alerts** tab selected, displaying a table of alerts filtered by **Maintenance**. Below this, the **Component Detail** pane shows the **Customer Impact** tab selected, displaying a table of customer impact data.

Name	Σ	▼	▼	▼
Alert Queues	7	5	2	
Gold Customers	2	1	1	
Default	5	4	1	

Severity	Date/Time	Name	Class	Category	Summary	Service Impact
Major	Jun 27, 2012 2:11:40 PM EDT	Operations	Service		Service is moderately degraded due to 1 active r...	Moderate
Critical	Jun 27, 2012 2:36:45 PM EDT	Server3	Computer S...		Service is stopped	Moderate

Component Detail: Server3 of type Computer System

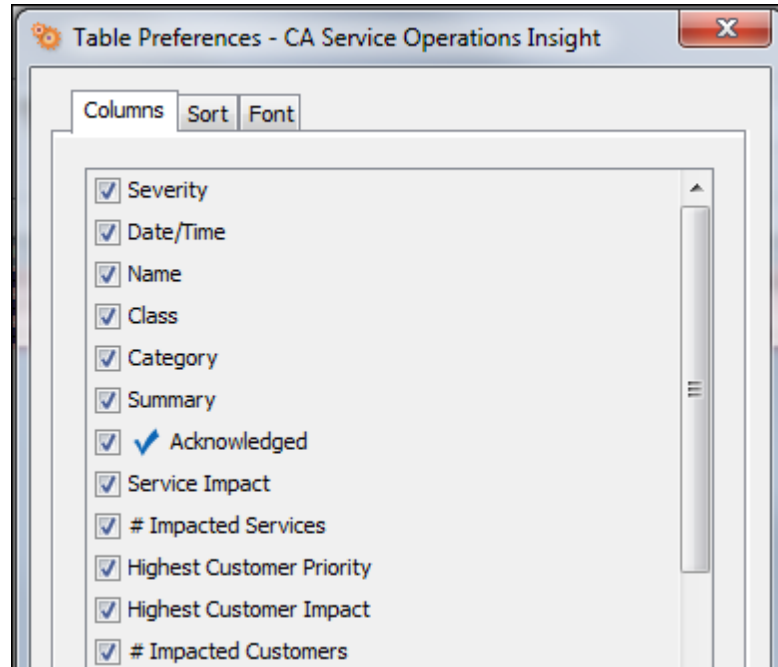
Alert Details | Information | Root Cause | **Service Impact** | Customer Impact | USM Properties | USM Notebook | SOL Properties

Filter: [] Displaying 2 of 2

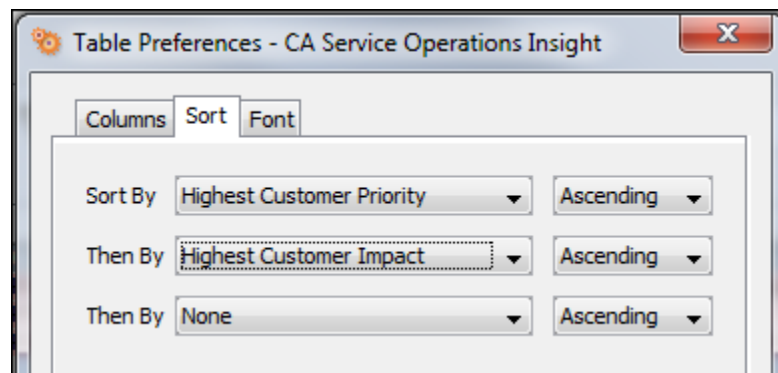
Name	Customer Identity	Customer Impact	Priority	Description
Company-A	A	Moderate	Gold	
Company-B	B	Moderate	Bronze	

In this example, only the Company-A customer has a Gold customer priority, so only its alerts appear in the new Gold Customers alert queue. If the administrator adds a customer, say Company-C, with a Gold customer priority, its alerts dynamically appear in the queue.

The administrator adds additional columns to the Alerts tab by right-clicking a column heading and selecting Highest Customer Priority and Highest Customer Impact. The completed dialog is shown in the following graphic:



The administrator can then sort the Alerts tab by these columns also. For example, the following graphic shows that the administrator is sorting the Alerts tab first by Highest Customer Priority and then by Highest Customer Impact.



This sort would show the highest customer impact on Gold service level customers.

Create Escalation Policy Using Customer Attributes

The administrator also elects to use the customer attributes to create escalation policy. The administrator launches the Alert Escalation Policy Editor and uses the following attributes: Customer Priority Equal to "Gold" AND Customer Impact Greater Than or Equal To "Moderate".

The completed Attributes tab is shown in the following graphic:

The screenshot shows the 'Alert Escalation Policy Editor - CA Service Operations Insight' window. The 'Attributes' tab is selected, showing a list of attributes and their values. The 'Name' field is 'Create Ticket for Gold Customer Severe Health' and the 'Policy is currently' status is 'Enabled'. The 'Description' field is empty. The 'Policy Type' is 'Global (Applies to all alerts)'. The 'Alert Selection' tab is also visible, showing a list of attributes and their values. The 'Attributes' tab shows a list of attributes and their values. The 'Attribute' is 'Customer Impact', the 'Comparison Type' is 'Greater Than Or Equal To', and the 'Attribute Value' is 'Moderate'. The 'Hints...' section shows a list of conditions: 'Customer Priority Equal To "Gold"' and 'Customer Impact Greater Than Or Equal To "Moderate"'. The 'Hints...' section also includes buttons for 'New AND', 'New OR', 'AND/OR', 'Cut', 'Copy', 'Paste', and 'Clear'. The 'Show Policy Summary >>' button is at the bottom left. The 'OK' and 'Cancel' buttons are at the bottom right.

Alert Escalation Policy Editor - CA Service Operations Insight

Policy Definition | Policy Actions | Service Assignment | Alert Queue Assignment

Name * Create Ticket for Gold Customer Severe Health Policy is currently ☒ Enabled ☐ Disabled

Description

Policy Type ☒ Global (Applies to all alerts) ☐ Non-Global (Applies to all alerts of assigned service or queue) [Hints...](#)

Alert Selection | Time | Attributes | Escalation Schedule

Attribute Customer Impact

Comparison Type Greater Than Or Equal To ☐ Ignore Case

Attribute Value Moderate [Add](#) [Apply](#) [Clear](#)

[Hints...](#)

- AND
 - Customer Priority Equal To "Gold"
 - Customer Impact Greater Than Or Equal To "Moderate"

[New AND](#)
[New OR](#)
[AND/OR](#)
[Cut](#)
[Copy](#)
[Paste](#)
[Clear](#)

[Show Policy Summary >>](#)

* indicates a required field

[OK](#) [Cancel](#)

The administrator then creates a Policy Action that automatically creates a help desk ticket when the policy conditions are met.

Appendix A: Service Modeling Examples and Scenarios

This section includes service modeling examples and scenarios. Examples of real-world scenarios and service models that illustrate the service modeling process are provided.

This section contains the following topics:

[Service Modeling Example 1 - Finance Service](#) (see page 145)

[Service Modeling Example 2- Shopping Cart Service](#) (see page 149)

[Service Modeling Example 3 - Dynamic Service Policy](#) (see page 153)

[Service Modeling Example 4 - Automatic Relationships Policy](#) (see page 157)

[Service Modeling Example 5 - Unmanaged Relationships Policy](#) (see page 163)

Service Modeling Example 1 - Finance Service

This scenario involves a large company with a vast infrastructure that must stay online to maintain the continuity of business processes. The company's IT manager is tasked with creating a service model representing the infrastructure components that comprise the company's finance department.

Finance Service Resources

The finance department resources must remain available so that the company maintains an accurate record of financial activity. Resources that affect the finance department include the following:

- A finance server hosts critical finance applications
- Database servers and databases maintain financial records
- A subnetwork under which the finance resources run

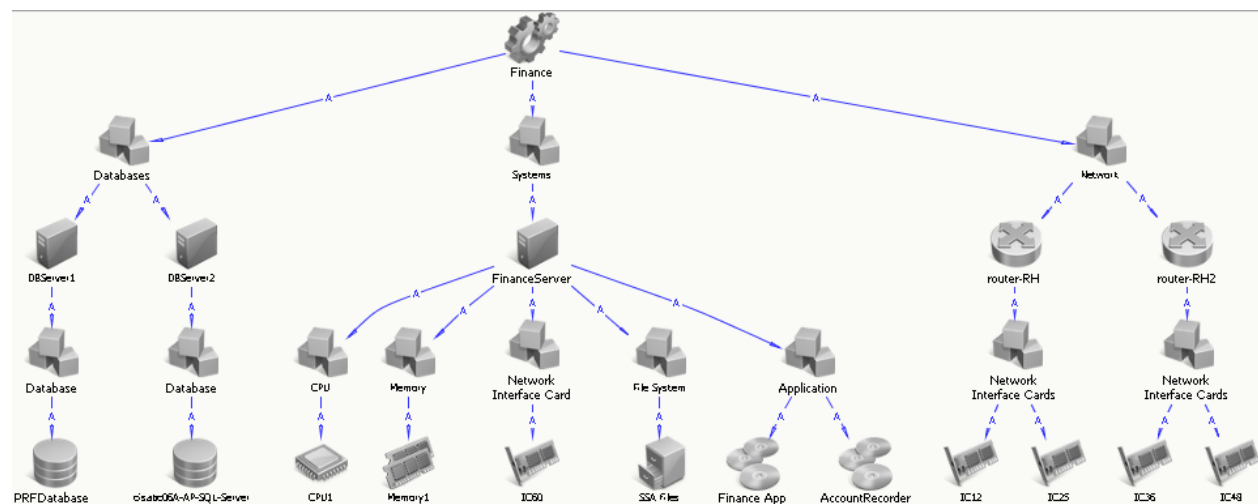
The following domain managers manage these resources:

- CA NSM and CA Spectrum manage the finance server.
- CA Insight DPM manages the finance database.
- CA Spectrum and CA eHealth manage the finance subnetwork.
- CA Application Performance Management manages the finance applications.

All important finance resources are currently managed, but no domain manager provides a consolidated view of all resources affecting the finance department infrastructure. Therefore, the domain managers cannot accurately represent service status and the root cause of problems affecting the service. A service model in CA SOI can leverage the intelligence of these domain managers to accomplish what a single domain manager cannot.

Finance Service Model

The IT manager models the Finance service in CA SOI as shown in the following graphic:



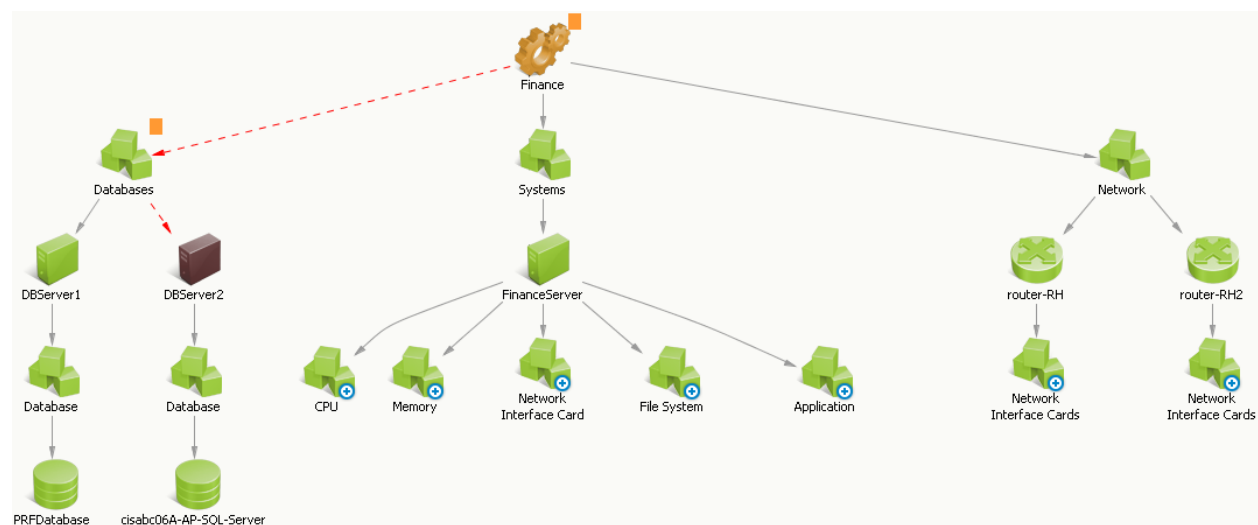
The Finance service has the following characteristics:

- It is a bottom-up service model arranged by domain. [Logical groups](#) (see page 45) represent the network, systems, and database resources.
- The service model uses logical groups to assemble CIs of the same type (for example, Application and Network Interface Cards) and create a logical, easy to understand representation of all resources.
- The service is modeled organically with no services imported from the source domain managers. If any of the source domain managers contained a service representation similar to one of the top-level groups, the IT manager could have imported that service and added it to the Finance service as a subservice.

- All CIs and groups use [aggregate propagation](#) (see page 17), so that impact propagates directly from related items.
- Employees enter financial records using the Finance App and AccountRecorder applications. If these applications fail, the entire finance department cannot function. Therefore, the IT manager [increased the significance](#) (see page 53) of the application CIs from the default of 4 to 8. Because the FinanceServer hosts the applications, the significance of the FinanceServer CI has also been increased from the default of 5 to 9. These changes adjust the significance to reflect that the finance applications are the most important aspects of the Finance service, other than the supporting network.

Example Finance Service Escalation Flow

The modeled Finance service provides the necessary information to quickly detect and resolve issues that cause service degradation. Consider this simple service condition, displayed from the Topology view of the Operations Console as shown in the following graphic:



The conditions appear and are resolved as follows:

1. Two alerts for DBServer2 appear in the Operations Console as shown in the following graphic:


Severity	Date/Time	Name	Class	Category	Summary	Service Impact	# Impacted Servi...	Source
Down	Nov 1, 2010 2:50:06 PM CDT	DBServer2	Computer S...	Risk	Misconfigured buffers	Moderate	2	CA:09998_LCDV
Critical	Nov 1, 2010 2:50:06 PM CDT	DBServer2	Computer S...	Quality	Low memory	Moderate	2	CA:09998_LCDV

These infrastructure alerts originated from the source domain manager (in this case, CA Insight DPM) and indicate Down and Critical alerts on the DBServer2 CI. If configured, escalation policies could perform specified actions on the alert, such as opening a help desk ticket or sending an email to a technician.

2. The SA Manager calculates the impact on the Finance service, changes the Finance service status, and generates a service alert as shown in the following graphic:

Severity	Date/Time	Name	Class	Category	Summary	Service Impact	# Impacted Servi...	Source
Major	Nov 4, 2010 10:41:25 AM CDT	Finance	Service		Service is moderately degraded due to 1 active r...	Moderate	2	CA Spectrum(R)

The Major severity results from the [service impact](#) (see page 27) value of 20 as shown in the following graphic:



CA SOI calculates the impact based on the severity of the root cause alert and the significance of the alerted CI. In this case, the root cause alert is the Down alert because it is a higher severity than the Critical alert, which is confirmed in the Root Cause tab. If other areas of the service had a similar severity with a higher CI significance (FinanceApp, for example), the service impact would be greater, and this condition would instead aggregate to the service level as the root cause condition.

- Service stakeholders can see a graphical view of all service conditions from the Dashboard as shown in the following graphic:

Services	Priority	Current SLA	Health	Quality	Risk	Availability [24 hours]	Operational Mode	Launch To
Finance	Unspecified		Down	Critical	Down	100%*	Production	Action

Notice that service quality and risk are both affected, because the Down alert belongs to the Risk category, and the Critical alert belongs to the Quality category. Health is a reflection of the worst state held by quality or risk.

- To resolve the problem, the assigned technician right-clicks the alert to drill down into the source domain manager (in this case, CA Insight DPM) to learn more about the problems, if necessary.
- The assigned technician reconfigures the database server buffers, resolves the lack of free memory, and clears the infrastructure alerts.
- The alerts disappear from the Operations Console and the service condition returns to normal.

Service Modeling Example 2- Shopping Cart Service

This scenario involves a retail company that contracts a managed service provider to host its online ordering application and associated hardware resources. The managed service provider must create a Shopping Cart service to represent the infrastructure components of the retail company's online ordering application.

Shopping Cart Service Resources

The online ordering resources must remain available at all times and maintain specific quality levels, so that the retail company guarantees optimal customer experience and the managed service provider meets contractual obligations. Resources that affect the Shopping Cart service include the following:

- A web server farm that hosts the online ordering applications
- A database cluster that maintains shopping cart records

The following domain managers manage these resources:

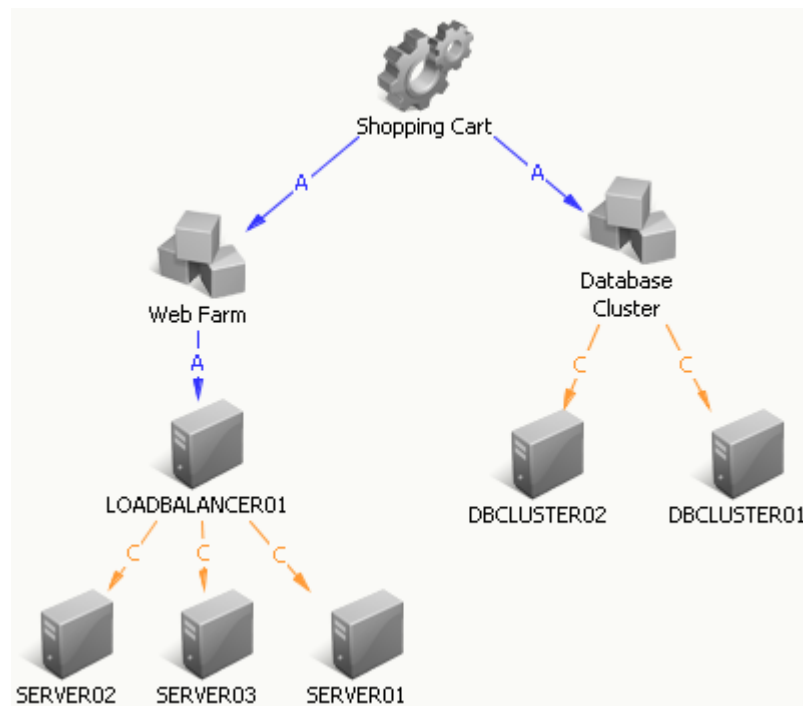
- CA NSM, CA Spectrum, and CA eHealth manage the availability and health of the web server farm.
- CA Application Performance Management manages the performance of the underlying applications.
- CA CMDB and CA Application Performance Management maintain a record of the service and its underlying components.

- CA Application Configuration Manager monitors the configuration compliance of all components.
- CA Service Desk is the integrated help desk application.
- CA Insight DPM manages the availability and health of the database cluster.

The important online ordering resources are currently managed, and CA CMDB maintains a record of the complete service. However, CA SOI can combine the service view with the management information, detect the root cause of problems, and provide a quick path to resolution through links to the source domain managers and the integrated help desk application.

Shopping Cart Service Model

The managed service provider [imports the Shopping Cart service](#) (see page 75) from CA CMDB. The service appears in CA SOI as shown in the following graphic:

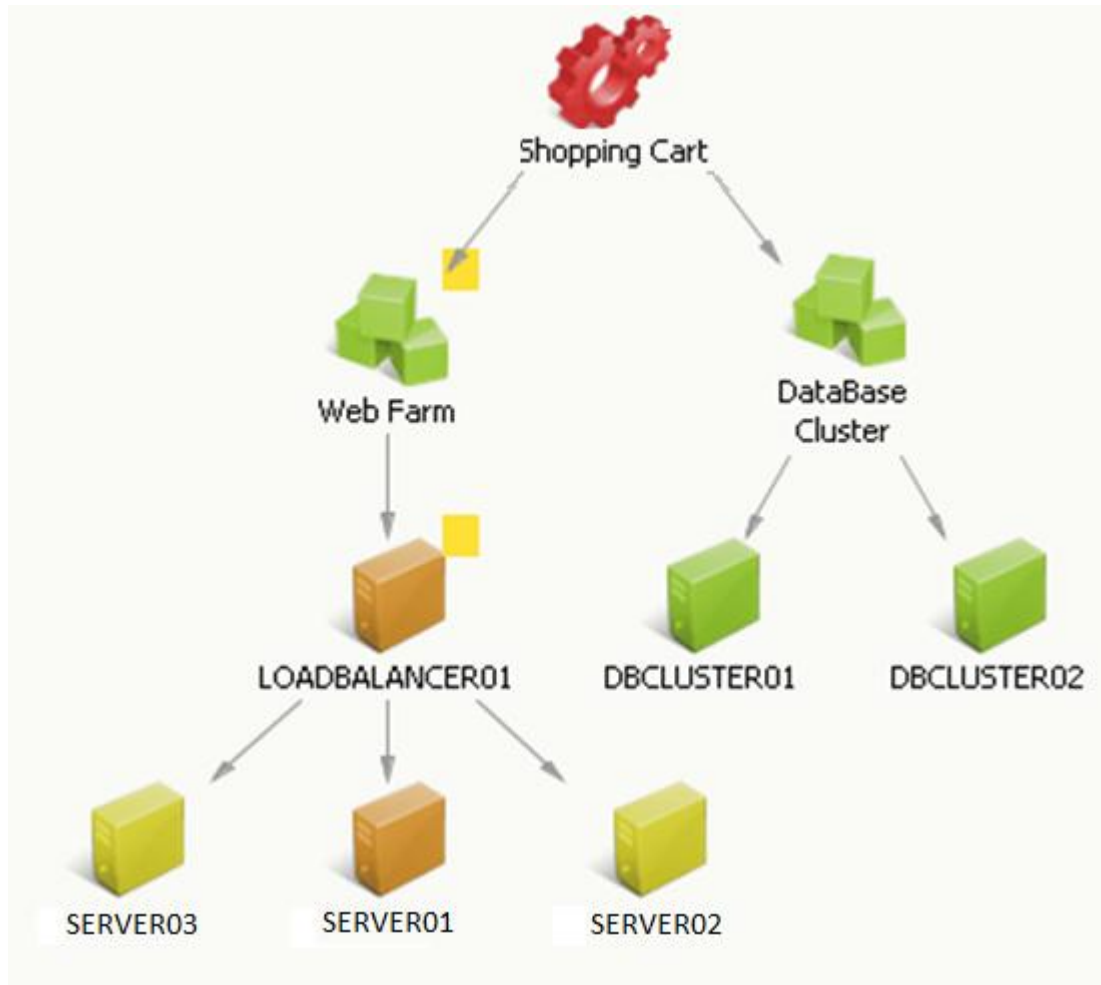


The Shopping Cart service has the following characteristics:

- [Logical groups](#) (see page 45) represent the two main resource domains: the web server farm and the database cluster. The groups use [aggregate propagation](#) (see page 17) to the service as shown by the letter A on the arrows connecting the CIs.
- The CIs in the web server farm and database cluster use [custom propagation](#) (see page 19) as shown by the letter C on the arrows connecting the CIs. Each group of CIs has an associated [custom propagation policy](#) (see page 50) that specifies the conditions under which to propagate impact to the related CI (LOADBALANCER01) or group (Database Cluster). Custom propagation policy takes into account that if one server in a cluster or web farm fails, this should not have a severe impact on the service, as long as the other servers are running at peak performance.
- The managed service provider has created [escalation policy](#) (see page 55) that opens a CA Service Desk ticket when a service alert occurs.
- CA Application Performance Management and CA Application Configuration Manager track important service-level metrics and compliance information for the service and associated CIs. If necessary, the managed service provider could also create an [SLA](#) (see page 54) for the service that tracks performance against service health, quality, risk, or availability.

Example Shopping Cart Service Escalation Flow

The Shopping Cart service in CA SOI provides the necessary information to quickly detect and resolve issues that cause service degradation. Consider these simple service conditions, displayed from the Topology view of the Operations Console as shown in the following graphic:



These conditions appear and are resolved as follows:

1. Alerts appear in the Operations Console for the Web Farm servers as shown in the following graphic:

Severity	Name	Ticket ID	Category	Description
Major	SERVER01		Risk	HIGH CPU UTILIZATION
Minor	SERVER02		Risk	Low CPU Utilization. The utilization of 5% for CPU instance 0x015777 named 'CPU_Ti
Minor	SERVER03		Risk	Low CPU Utilization. The utilization of 7% for CPU instance 0x064321 named 'CPU_Ti
Major	LOADBALANCER		Risk	Configuration for device: Loadbalancer01 is outside of standard compliance

SERVER01 has a major severity alert for high CPU utilization, while SERVER02 and SERVER03 have minor severity alerts for unusually low CPU utilization. LOADBALANCER01 has a major severity alert from CA Application Configuration Manager for a compliance violation.

2. The impact of the web farm alerts does not propagate, because it is not high enough to meet the conditions defined in the [custom propagation policy](#) (see page 50). However, these alerts do affect the service risk value.
3. An alert appears for the Shopping Cart service CI as shown in the following graphic:

▼ Critical	Shopping Cart	Quality	Transaction times have breached a critical threshold on Shopping Cart Service
------------	---------------	---------	---

This alert from CA Application Performance Management indicates a critical threshold breach in the transaction time for online ordering.

4. The critical alert causes a similar service alert as shown in the following graphic:

▼ Critical	Shopping Cart	60	Service is severely degraded due to one active critical alarm
------------	---------------	----	---

This service alert changes the Shopping Cart service condition to critical.

5. [Escalation policy](#) (see page 55) automatically opens a ticket in CA Service Desk in response to the service alert. The assigned technician can directly access the ticket in CA Service Desk by clicking the link on the number in the Ticket ID column.
6. The assigned technician consults CA SOI, CA Service Desk, and the source domain manager to resolve the problem before the managed service provider violates contractual obligations.

Service Modeling Example 3 - Dynamic Service Policy

This scenario involves a system administrator who wants to automatically place relevant CIs system-wide under services. The system administrator has two sets of conditions and will create a policy for each:

- All hardware in a specific region
- All virtual systems that exist and those added in the future

Policy - All Running Hardware with Specific Starting IP Address

Because one of the company's enterprise regions is located in a specific subnet, the system administrator wants to create dynamic service policy to collect the hardware resources from that region and add to a service. This provides the system administrator with a consolidated view of that region. The system administrator could then model the relationships among the resources or set up automatic relationships.

Because all hardware resources in the region begin with the same IPV4 address (10.0.21.x), the system administrator creates a dynamic service policy for all IPV4 addresses starting with 10.0.21.

The system administrator [creates a new dynamic service policy](#) (see page 80) using the Service Discovery Policy Editor wizard:

Define Service

The system administrator completes the Define Service page as follows:

Service Name: RunningHardwareService

Relationship Type: Has Access To

The system administrator then completes the Target fields:

Class: Running Hardware

Attribute: Device IPV4Address With Domain

Note that all USM properties available are available except the following:

- domain properties (MdrProduct, MdrProductInstance, MdrElementID)
- date and time properties

Comparison Type: Starts With

Attribute Value: 10.0.21

The system administrator clicks Add and the criteria appear as shown in the following graphic:

Relationship Criteria

Service Name * RunningHardwareService

Relationship Type * Has Access To

Target

Class * Running Hardware

Attribute Device IPv4Address With Domain

Comparison Type Starts With ☒ Ignore Case

Attribute Value 10.0.21 **Add**

[Hints...](#)

AND

'Device IPv4Address With Domain' Starts With "10.0.21"

If the system administrator wanted to put a dynamic service CI with a different class, the system administrator can define another policy which has the same service name and relationship type, but a different class.

Confirm

The dynamic services policy appears on the Confirm page as follows:

```
Dynamic Service Policy 'RunningHardwareService'
WITH
    relationship 'Has Access To'
FOR
    Running Hardware with properties
    (
        'Device IPv4Address With Domain' starts with (ignore case) "10.0.21."
    )
```

Policy - All Virtual Systems by a Specific Vendor

The company is converting its data center to virtual resources over time. The system administrator wants a service container for all existing virtual machines and virtual machines that come online in the future.

The system administrator creates a dynamic service policy that automatically puts all virtual systems by vendor VMware, Inc. under a service.

The system administrator [creates a new dynamic service policy](#) (see page 80) using the Service Discovery Policy Editor wizard:

Define Service

The system administrator completes the Define Service page as follows:

Service Name: VirtualSystemsService

Relationship Type: Has Access To

The system administrator then completes the Target fields:

Class: Virtual System

Attribute: Vendor

Comparison Type: Equal To

Attribute Value: VMware, Inc.

The system administrator clicks Add and the criteria appear as shown in the following graphic:

Relationship Criteria

Service Name * VirtualSystemsService

Relationship Type * Has Access To

Target

Class * Virtual System

Attribute Vendor

Comparison Type Equal To ☐ Ignore Case

Attribute Value VMware, Inc. **Add**

[Hints...](#)

AND

'Vendor' Equal To "VMware, Inc."

Confirm

The dynamic services policy appears on the Confirm page as follows:

```
Dynamic Service Policy 'VirtualSystemsService'
WITH
  relationship 'Has Access To'
FOR
  Virtual System with properties
  (
    'Vendor' equals "VMware, Inc."
  )
```

Service Modeling Example 4 - Automatic Relationships Policy

This scenario involves a system administrator whose company is in the early stages of adopting a service-oriented infrastructure and wants to define general rules to automate universal granular relationships that the service administrator would otherwise have to create manually.

The system administrator has two sets of criteria for creating relationships and will create a policy for each:

- All computer systems and software that runs on them
- All database instances and their tablespaces

Policy - All Computer Systems and Running Software

The company is in the process of hiring more employees and so they are constantly adding new computer systems bundled with various software packages.

The system administrator wants to automatically create relationships between computer systems and the software that runs on them. When new computer systems come online, the system administrator wants the relationship creation automated also.

The system administrator [creates a new automatic relationship policy](#) (see page 85) using the Service Discovery Policy Editor wizard:

Source Criteria

The system administrator completes the Source Criteria page with the following values:

Relationship Type: Has Access To

Class: ComputerSystem

The system administrator creates the first source criteria:

Attribute: Primary IPV4Address

Comparison Type: Is Set

The system administrator clicks Add then creates the second source criteria:

Attribute: Primary Dns Name

Comparison Type: Is Set

The system administrator clicks Add and the resulting source criteria appear as shown in the following graphic:

Source Criteria

Relationship Type * Has Access To

Source

Class * Computer System

Attribute Primary Dns Name

Comparison Type Is Set ☐ Ignore Case

Attribute Value

[Hints...](#)

AND

- 'Primary IPV4Address' Is Set
- 'Primary Dns Name' Is Set

Target Criteria

The system administrator clicks Next and completes the Target Criteria page as follows:

Class: RunningSoftware

The system administrator creates the first target criteria:

Attribute: Device IPV4Address

Comparison Type: Is Set

The system administrator clicks Add then creates the second target criteria:

Attribute: Device Dns Name

Comparison Type: Is Set

The system administrator clicks Add and the resulting target criteria appear as shown in the following graphic:

Target Criteria

Relationship Type Has Access To

Target

Class * Running Software

Attribute Device Dns Name

Comparison Type Is Set ☐ Ignore Case

Attribute Value

[Hints...](#)

AND

- 'Device IPV4Address' Is Set
- 'Device Dns Name' Is Set

Match Criteria

The system administrator creates the first criteria:

'Computer System Attribute: Primary IPV4Address

Comparison Type: Equal To

'Running Software' Attribute: Device IPV4Address

The system administrator clicks Add and creates the second criteria:

'Computer System Attribute: Primary Dns Name

Comparison Type: Equal To

'Running Software' Attribute: Device Dns Name

The system administrator clicks Add and the resulting match criteria appear as shown in the following graphic:

Match Criteria

Relationship 'Has Access To' between 'Computer System' and 'Running Software'

Criteria Selection

'Computer System' Attribute Primary Dns Name

Comparison Type Equal To ☒ Ignore Case

'Running Software' Attribute Device Dns Name

[Hints...](#)

AND

- 'Primary IPV4Address' Equal To 'Device IPV4Address'
- 'Primary Dns Name' Equal To 'Device Dns Name'

Relationship Scope

The system administrator selects All services so that all services with matching ComputerSystem CIs are considered for creating the relationship.

Confirm

The automatic relationship policy appears on the Confirm page as follows:

```
Automatic Relationship Policy 'Has Access To'
BETWEEN
  Computer System with properties
  (
    'Primary IPV4Address' is set AND
    'Primary Dns Name' is set
  )
AND
  Running Software with properties
  (
    'Device IPV4Address' is set AND
    'Device Dns Name' is set
  )
WHEN
  (
    'Computer System.Primary IPV4Address' equals (ignore case) 'Running
Software.Device IPV4Address' AND
    'Computer System.Primary Dns Name' equals (ignore case) 'Running
Software.Device Dns Name'
  )
SCOPED TO
  (
    All services
  )
```


Policy - All Database Instances and Tablespaces

The company is adding new enterprise software that requires adding multiple databases. The company is also planning on adding additional databases in the future.

The system administrator wants to automatically create relationships between database instances and their respective tablespaces. When a new database comes online, the system administrator wants the relationship creation automated also.

The system administrator [creates a new automatic relationship policy](#) (see page 85) using the Service Discovery Policy Editor wizard:

Source Criteria

The system administrator completes the Source Criteria page with the following values:

Relationship Type: Has Access To

Class: Database Instance

Attribute: DBInstanceName

Comparison Type: Is Set

The system administrator clicks Add and the resulting source criteria appear as shown in the following graphic:

Source Criteria

Relationship Type * Has Access To

Source

Class * Database Instance

Attribute DBInstance Name

Comparison Type Is Set ☐ Ignore Case

Attribute Value

[Hints...](#)

AND

- 'DBInstance Name' Is Set

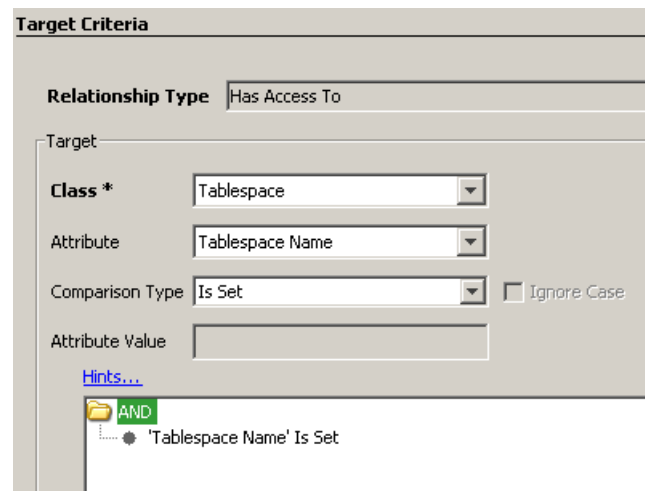
Target Criteria

Class: Tablespace

Attribute: Tablespace Name

Comparison Type: Is Set

The system administrator clicks Add and the resulting source target criteria appears as shown in the following graphic:



The 'Target Criteria' dialog box is shown. It has a title bar 'Target Criteria'. Below the title bar is a section 'Relationship Type' with a dropdown menu set to 'Has Access To'. Below this is a section 'Target' with four fields: 'Class *' with a dropdown set to 'Tablespace', 'Attribute' with a dropdown set to 'Tablespace Name', 'Comparison Type' with a dropdown set to 'Is Set', and 'Attribute Value' with an empty text box. To the right of the 'Comparison Type' dropdown is a checkbox labeled 'Ignore Case' which is unchecked. Below these fields is a blue link 'Hints...'. At the bottom of the dialog is a tree view showing a folder icon with the text 'AND' and a single criterion: 'Tablespace Name' Is Set.

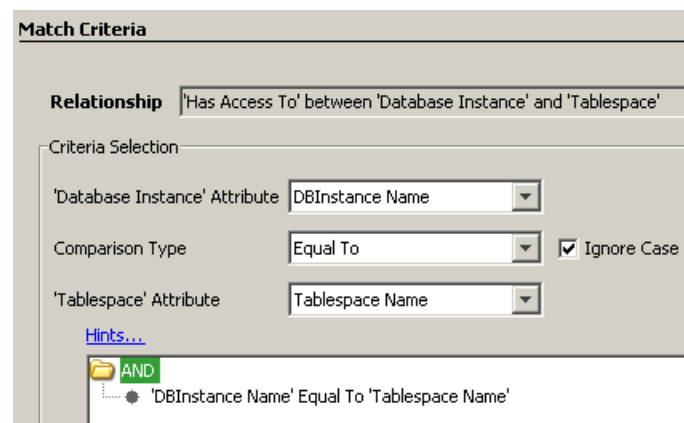
Match Criteria

'Database Instance' Attribute: DBInstance Name

Comparison Type: Equal To

'Tablespace' Attribute: Tablespace Name

The system administrator clicks Add and the resulting match criteria appears as shown in the following graphic:



The 'Match Criteria' dialog box is shown. It has a title bar 'Match Criteria'. Below the title bar is a section 'Relationship' with a text field containing the text "'Has Access To' between 'Database Instance' and 'Tablespace'". Below this is a section 'Criteria Selection' with three fields: ''Database Instance' Attribute' with a dropdown set to 'DBInstance Name', 'Comparison Type' with a dropdown set to 'Equal To', and ''Tablespace' Attribute' with a dropdown set to 'Tablespace Name'. To the right of the 'Comparison Type' dropdown is a checkbox labeled 'Ignore Case' which is checked. Below these fields is a blue link 'Hints...'. At the bottom of the dialog is a tree view showing a folder icon with the text 'AND' and a single criterion: 'DBInstance Name' Equal To 'Tablespace Name'.

Relationship Scope

The system administrator selects All services so that all services are considered for relationship creation.

Confirm

The automatic relationship policy appears on the Confirm page as follows:

Automatic Relationship Policy 'Has Access To'

BETWEEN

Database Instance with properties

```
(  
    'DBInstance Name' is set  
)
```

AND

Tablespace with properties

```
(  
    'Tablespace Name' is set  
)
```

WHEN

```
(  
    'Database Instance.DBInstance Name' equals (ignore case)  
    'Tablespace.Tablespace Name'  
)
```

SCOPED TO

```
(  
    All services  
)
```

Service Modeling Example 5 - Unmanaged Relationships Policy

This scenario, *Add Hosted Systems*, involves a system administrator whose company's IT infrastructure includes a domain manager that publishes information about operating systems. The operating systems are Linux-based operating systems. They are either *hosting* operating systems (which implies they run a virtual computer or computers) or *hosted* operating systems (which implies they run on a virtual system). The domain manager also publishes *IsHostFor* relationship between *hosting* and *hosted* operating systems. The system administrator wants to create an unmanaged relationship policy. The policy ensures adding a *hosting* operating system to a service also adds all *hosted* operating systems.

The system administrator creates a new unmanaged relationship policy using the Service Discovery Policy Editor wizard:

Relationship Criteria

The system administrator completes the Relationship Criteria page with the following values:

Policy Name: Add Hosted Systems

Create Relationship of Type: Same as Discovered

Discover Relationships of Type: IsHostFor

The system administrator clicks Add and the resulting relationship criteria appear as shown in the following graphic:

The screenshot shows the 'Relationship Criteria' page of the Service Discovery Policy Editor wizard. The page has a title bar 'Relationship Criteria'. Below the title bar, there are several fields and controls:

- Policy Name ***: A text box containing 'Add Hosted Systems'.
- ☐ **Reverse the created relationships**: An unchecked checkbox.
- Create Relationship of Type ***: A dropdown menu showing 'Same as Discovered'.
- [Hints...](#): A blue hyperlink.
- Discover Relationships of Type:**: A section with two panes:
 - Available Types**: A list box containing 'HasDetail', 'HasMember', and 'HasRequirementFor'.
 - Selected Types**: A list box containing 'IsHostFor'.

Source Criteria

The system administrator completes the Source Criteria page with the following values:

Class: Operating System

Attribute: OSType

Comparison Type: Equal To

Attribute Value: Linux

The system administrator clicks Add and the resulting source criteria appear as shown in the following graphic:

The screenshot shows a dialog box titled "Source Criteria". It has a "Source" section with the following fields:

- Class ***: Operating System
- Attribute**: OSType
- Comparison Type**: Equal To
- Attribute Value**: Linux

There is an "Ignore Case" checkbox which is unchecked. An "Add" button is located to the right of the "Attribute Value" field. Below the "Add" button is a "Hints..." link. Below the "Hints..." link is a tree view showing a folder icon, an "AND" label, and a single criterion: "OSType' Equal To 'Linux'". At the bottom of the dialog, the logical expression is displayed: `((('OSType' equals "Linux")))`.

Target Criteria

The system administrator completes the Target Criteria page with the following values:

Class: Operating System

Attribute: OSType

Comparison Type: Equal To

Attribute Value: Linux

The system administrator clicks Add and the resulting target criteria appear as shown in the following graphic:

Target Criteria

Target

Class * Operating System

Attribute OSType

Comparison Type Equal To ☐ Ignore Case

Attribute Value Linux **Add**

[Hints...](#)

AND

'OSType' Equal To "Linux"

(("OSType" equals "Linux"))

Note: This scenario does not require any match criteria.

Relationship Scope

The system administrator selects All services so that all services with matching OperatingSystem CIs are considered for creating the relationship.

Confirm

The unmanaged relationship policy appears on the Confirm page as follows:

```
Unmanaged Relationship Policy 'Add Hosted Systems' creates
relationships with
SEMANTIC
    Same as Discovered
BETWEEN
    Operating System with properties
    (
        'OSType' equals "Linux"
    )
```

```
AND
  Operating System with properties
  (
    'OSType' equals "Linux"
  )
SCOPED TO
  (
    All services
  )
BASED ON SEMANTICS
  (
    IsHostFor
  )
```


Glossary

alert

An *alert* is a message on the Operations Console that reports a fault condition that is associated with a resource or service.

alert escalation

Alert escalation is the ability to enact some escalating action as a result of an alert. Actions include opening a help desk ticket, running a command, sending an email, and so on.

alert filter

An *alert filter* limits the number and alert types that are shown on the Operations Console.

alert queue

Alert queues are user-defined alert groups. CA SOI auto-assigns alerts to a particular alert queue based on user-defined policy, which can include alert content and associated CIs.

bidirectional connectors

A *bidirectional connector* supports both inbound and outbound operations.

CI

See *configuration item*.

configuration item (CI)

A *configuration item (CI)* is a managed resource such as a printer, software application, or database. Configuration items support services. A synonym for a configuration item is a *resource*.

Connector

A *connector* is software that provides the interface for the data exchange between the CA Catalyst infrastructure and a domain manager.

domain manager

A *domain manager* is a management application that provides information to CA Catalyst and CA SOI using a connector.

enrichment

An *enrichment* adds information to an event or alert from an outside source. An outside source can include a database, method, or script. You can add an enrichment to event policies to enrich matching events with important information.

escalation policy

Escalation policy specifies the automated actions to take in response to fault conditions.

event

An *event* is a message that indicates an occurrence or change in your enterprise. Events can indicate a negative occurrence or object state change. CA SOI Event Management lets you view and manage events that are received from all connectors.

Event Management

Event Management provides a processing layer between raw connector USM alert data and alert management. You can influence the data that makes it into the Operations Console as alerts and how those alerts appear. Available Event Management actions include event correlation, filtering, creation, and enrichment.

event policy

An *event policy* is a combination of event search patterns and an action to perform when the patterns match. You can deploy event policies on all or specific sources, and available actions include filtering, event creation, and enrichment.

impact

Impact indicates how much a CI affects a service and related CIs.

inbound to connector operations

Inbound to connector operations (also referred to as "southbound") use records in the CA Catalyst Persistent Store and the CA Catalyst Synchronizer to create, update, or delete items in the source domain manager.

JDBC

Java Database Connectivity is a programming interface that lets Java applications access an SQL database.

JMS (Java Message Service)

The *Java Message Service (JMS)* is a messaging standard that lets application components that are based on the Java 2 Platform Enterprise Edition (J2EE) create, send, receive, and read messages.

Mid-tier connector

The *Mid-tier connector* is an intermediate processing layer through which you can deploy cross-domain event policy for automated processing on events from all connectors. All events flow through the Mid-tier connector before reaching the Operations Console as alerts.

outbound from connector operations

Outbound from connector operations obtain data (such as services, CIs, topology, alerts, and status) from the source domain manager.

priority

Priority indicates the importance of a service to the business.

relationship

Relationships in a service model show how CIs are linked to form the service topology.

resource

A *resource* is a managed resource such as a printer, software application, or database. Resources support services. A synonym for a resource is a configuration item (CI).

root cause alert

A *root cause alert* is an alert that CA SOI determines after analyzing the alerts associated with a service which is based on one of the following criteria:

1. A triggered root cause rule determining the alert that is the true root cause of the service degradation which is based on relationships and topology.
2. The alert with the highest impact if no root cause rules have been triggered.

service

A *service* typically consists of several CIs, which are grouped to represent entities like web server farms or clusters. Services can also contain *subservices*, which are subordinate service models. Service models typically represent high-level abstract entities like a web-based retail transaction service, an application server service, or a source control service. You can define any service type with CA SOI as long as one of the integrated domain managers monitors the service components.

service model

A *service model* is a definition of a service or other entity in your enterprise. It is a logical grouping of resources, associations, dependencies, and policies.

service-level agreement

A *service-level agreement* (SLA) is a contract that specifies the service expectations of internal or external customers. An example is the downtime that is acceptable for various resources.

significance

Significance indicates the importance of a CI.

subservice

A *subservice* is used to indicate a subordinate service model.

synchronization

Synchronization is the CA Catalyst capability of updating source domain manager data due to CI changes. Changes are a result of reconciliation or other changes to data in the Persistent Store, including CI creation and deletion.

Unified Service Model (USM)

The *Unified Service Model (USM)* is the semantic schema that is used as the CA Catalyst and CA SOI infrastructure.